

The stability of Delaunay triangulations

Jean-Daniel Boissonnat

Ramsay Dyer

Arijit Ghosh

December 18, 2014

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 quantify the magnitude of the perturbations under which the Delaunay triangulation remains unchanged.

Contents

1	Introduction	1
2	Background	2
2.1	Sampling parameters and perturbations	2
2.2	Simplices	2
2.3	Complexes	6
3	Parameterized genericity	8
3.1	The Delaunay complex	8
3.1.1	The Delaunay complex in other metrics	10
3.2	Protection	10
3.2.1	Local Delaunay triangulation	11
3.2.2	Safe interior simplices	13
3.3	Thickness from protection	15
4	Delaunay stability	17
4.1	Perturbations and circumcentres	17
4.1.1	Almost circumcentres	17
4.1.2	Circumcentres and metric perturbations	18
4.1.3	Circumcentres and point perturbations	21
4.2	Perturbations and protection	21
4.3	Perturbations and Delaunay stability	23
4.3.1	Point perturbations	23
4.3.2	Metric perturbation	24
5	Conclusions	25
	Bibliography	25

1 Introduction

One of the central properties of Delaunay complexes, which was demonstrated when they were introduced [Del34], is that under a very mild assumption they are embedded, i.e., they define a triangulation of Euclidean space. The required assumption is that there are not too many cospherical points; the points are “generic”. The assumption is not considered limiting because, as Delaunay showed, an arbitrarily small affine perturbation can transform any given point set into one that is generic.

Given the assumption of a generic point set, we are assured that the Delaunay complex defines a triangulation, but a couple of issues arise when working with these triangulations. One is that the Delaunay triangulation can be highly sensitive to the exact location of the points. For example, the Delaunay triangulation of a point set might be different if a coordinate transform is first performed using floating point arithmetic.

Another problem concerns the geometric quality of the simplices in the triangulation. We define the *thickness* of a simplex as a number proportional to the ratio of the smallest altitude to the longest edge length of the simplex, and we demonstrate why this is a useful measure of the geometric quality of the simplex. For points in the plane, if there is an upper bound on the ratio of the radius of a Delaunay ball to the length of the shortest edge of the corresponding triangle, then there is a lower bound on the thickness of any Delaunay triangle. However, when there are three or more spatial dimensions, the thickness of Delaunay simplices may become arbitrarily small in spite of any bound on the circumradius to shortest edge length.

Both of these issues are shown to be related to points being close to a degenerate (non-generic) configuration. We parameterize Delaunay’s original definition of genericity, saying that a point set $P \subset \mathbb{R}^m$ is δ -generic if every m -simplex in the Delaunay complex has a Delaunay ball that is at a distance greater than δ to the remaining points in P . We show that a bound on δ leads to a bound on the thickness of the Delaunay simplices, and also that the Delaunay complex itself is stable with respect to perturbations of the points or of the metric, provided the perturbation is small enough with respect to δ in a way that we quantify. In a companion paper [BDG13a], we develop a perturbation algorithm to produce δ -generic point sets.

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].

Our interest in the problem of near-degeneracy in Delaunay complexes stems from work on triangulating Riemannian manifolds. An established technique is to compute the triangulation locally at each point in an approximating Euclidean metric, and then perform manipulations to ensure that the local triangulations fit together consistently [BWY11, BG11]. The reason the manipulations are necessary is exactly the problem of the instability of the Delaunay triangulation, and sometimes this is most conveniently described as an instability with respect to a perturbation of the local Euclidean metric.

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 triangulation without explicit requirements on the properties of the Voronoi diagram [ES97], in contrast to a common practice 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 -genericity 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. We conclude with some remarks on the construction and application of δ -generic point sets.

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. We will consider other metrics besides the Euclidean one. A generic metric is denoted d , and the associated open and closed balls are $B(c; r)$, and $\overline{B}(c; r)$. Generally, we denote the topological closure of a set X by \overline{X} , the interior by $\text{int } X$, and the boundary by ∂X . The convex hull is denoted $\text{conv}(X)$, and the affine hull is $\text{aff}(X)$.

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

$$\sin \angle(U, V) = \sup_{\substack{u \in U \\ \|u\|=1}} \|u - \pi_V u\|, \quad (1)$$

where π_V is the orthogonal projection onto 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 $P \subset \mathbb{R}^m$, 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 $d_{\mathbb{R}^m}(x, P) < \epsilon$ 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 $d_{\mathbb{R}^m}(x, P \cup \partial \text{conv}(P)) < \epsilon$ for all $x \in \text{conv}(P)$. Equivalently, P is an ϵ -sample set if it satisfies a sampling radius ϵ for

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

The set P is λ -*separated* if $d_{\mathbb{R}^m}(p, q) > \lambda$ for all $p, q \in P$. 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 $P \subset \mathbb{R}^m$ given by a function $\zeta : P \rightarrow \mathbb{R}^m$. 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 $\tilde{P} = \zeta(P)$, and let \tilde{p} represent $\zeta(p) \in \tilde{P}$. We will only be considering ρ -perturbations where ρ is less than half the sparsity of P , so $\zeta : P \rightarrow \tilde{P}$ is a bijection.

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

2.2 Simplices

Given a set of $j + 1$ points $\{p_0, \dots, p_j\} \subset P \subset \mathbb{R}^m$, a (geometric) j -*simplex* $\sigma = [p_0, \dots, p_j]$ is defined by the convex hull: $\sigma = \text{conv}(\{p_0, \dots, p_j\})$. 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 $\text{int } \sigma = \sigma \setminus \partial\sigma$. Again this is generally different from the topological interior. In particular, a 0-simplex p is equal to its interior: it has no boundary. 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 opposite* 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 $\dim \text{aff}(\sigma) < j$. 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.

A *circumscribing ball* for a simplex σ is any m -dimensional ball that contains the vertices of σ on its boundary. If σ admits a circumscribing ball, 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 general a degenerate simplex may not have a circumcentre and circumradius, but in the context of the Euclidean Delaunay complexes we will work with, the degenerate simplices we may encounter do have these properties. 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 $\text{aff}(\sigma)$ and intersects it at the circumcentre of σ . Its dimension is $m - \dim \text{aff}(\sigma)$.

The *altitude* of a vertex p in σ is $D(p, \sigma) = d_{\mathbb{R}^m}(p, \text{aff}(\sigma_p))$. 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 $\Delta(\tau) \leq \Delta(\sigma)$.

Our definition of thickness is essentially the same as that employed by Munkres [Mun68]. Munkres defined the thickness of σ^j as $\frac{r(\sigma^j)}{\Delta(\sigma^j)}$, where $r(\sigma^j)$ is the radius of the largest contained ball centred at the barycentre. This definition of thickness turns out to be equal to $\frac{j}{j+1} \Upsilon(\sigma^j)$.

Whitney [Whi57] employed a volume-based measure of simplex quality, and variations on this, typically referred to as *fatness*, have been popular in works on higher dimensional Delaunay-based meshing [CDE⁺00, Li03, BWY11]. We find a direct bound on the altitudes to be more convenient, because it yields a cleaner and tighter connection between the geometry and the linear algebra of simplices. Typically, a bound on some geometric displacement related to a simplex is obtained by bounding the inverse of a matrix associated with the simplex, and thickness is well suited for this task.

As a motivating example, consider the problem of bounding the angle between the affine hull of a simplex and an affine space that lies close to all the vertices of the simplex. Such a bound is relevant when meshing submanifolds of Euclidean space, for example, where it is desired that the affine hulls of the simplices are in agreement with the nearby tangent spaces of the manifold.

Whitney [Whi57, p. 127] obtained such a bound, which manifestly depends on the quality of the simplex. Using thickness as a quality measure we obtain a sharper result:

Lemma 2.1 (Whitney angle bound) Suppose σ is a j -simplex whose vertices all lie within a distance η from a k -dimensional affine space, $H \subset \mathbb{R}^m$, with $k \geq j$. Then

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

The idea of the proof is to express the unit vector u in Equation (1) in terms of a basis for $\text{aff}(\sigma)$ given by the edges of σ that emanate from some arbitrarily chosen vertex. The projection \tilde{u} of u into H can then be expressed in terms of the projected basis vectors, using the same vector of coefficients. Since the vertices of σ all lie close to H , the projected basis vectors do not differ significantly from the originals, so bounding the magnitude of the difference between u and \tilde{u} comes down to bounding the magnitude of the vector of coefficients of the unit vector u . This bound depends on how well-conditioned the basis is, and this is closely related to the thickness of σ .

These observations can be conveniently expressed and made concrete in terms of the singular values of a matrix. An excellent introduction to singular values can be found in the book by Trefethen and Bau [TB97, Ch. 4 & 5], but for our purposes we are primarily concerned with the largest and the smallest singular values, which we now describe.

We denote the i^{th} singular value of a matrix A by $s_i(A)$. The singular values are non-negative and ordered by magnitude. The largest singular value can be defined as $s_1(A) = \sup_{\|x\|=1} \|Ax\|$; it is the magnitude of the largest vector in the range of the unit sphere. The first singular value also defines the operator norm: $\|A\| = s_1(A)$. The standard observation that a bound on the norms of the columns of A yields a bound on $\|A\|$ is obtained by a short calculation.

Lemma 2.2 If $\eta > 0$ is the least upper bound on the norms of the columns of an $m \times j$ matrix A , then $\eta \leq \|A\| \leq \sqrt{j}\eta$.

We will also be interested in obtaining a lower bound on the smallest singular value which, for an $m \times j$ matrix A with $j \leq m$, may be defined as $s_j(A) = \inf_{\|x\|=1} \|Ax\|$.

From the given definitions, one can verify that if A is an invertible $m \times m$ matrix, then $s_1(A^{-1}) = s_m(A)^{-1}$, but it is convenient to also accommodate non-square matrices, corresponding to simplices that are not full dimensional. If A is an $m \times j$ matrix of rank $j \leq m$, then the *pseudo-inverse* $A^\dagger = (A^\top A)^{-1} A^\top$ is the unique left inverse of A whose kernel is the orthogonal complement of the column space of A . We have the following general observation:

Lemma 2.3 If A is an $m \times j$ matrix of rank $j \leq m$, then $s_i(A^\dagger) = s_{j-i+1}(A)^{-1}$.

The columns of A form a basis for the column space of A . The pseudo-inverse can also be described in terms of the *dual basis*. If we denote the columns of A by $\{a_i\}$, then the i^{th} dual vector, w_i , is the unique vector in the column space of A such that $w_i^\top a_i = 1$ and $w_i^\top a_j = 0$ if $i \neq j$. Then A^\dagger is the $j \times m$ matrix whose i^{th} row is w_i^\top .

By exploiting a close connection between the altitudes of a simplex and the vectors dual to a basis defined by the simplex, we obtain the following key lemma that relates the thickness of a simplex to the smallest singular value of an associated matrix:

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 the i^{th} row of P^\dagger is given by w_i^\top , where w_i is orthogonal to $\text{aff}(\sigma_{p_i})$, and

$$\|w_i\| = D(p_i, \sigma)^{-1}.$$




Figure 1: Choosing p_0 as the origin, the edges emanating from p_0 in $\sigma = [p_0, \dots, p_j]$ form a basis for $\text{aff}(\sigma)$. The proof of Lemma 2.4 demonstrates that the dual basis $\{w_i\}$ consists of vectors that are orthogonal to the facets, and with magnitude equal to the inverse of the corresponding altitude.

We have the following bound on the smallest singular value of P :

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

Proof By the definition of P^\dagger , it follows that w_i belongs to the column space of P , and it is orthogonal to all $(p_{i'} - p_0)$ for $i' \neq i$. Let $u_i = w_i / \|w_i\|$. By the definition of w_i , we have $w_i^\top (p_i - p_0) = 1 = \|w_i\| u_i^\top (p_i - p_0)$. By the definition of the altitude of a vertex, we have $u_i^\top (p_i - p_0) = D(p_i, \sigma)$. Thus $\|w_i\| = D(p_i, \sigma)^{-1}$. Since

$$\max_{1 \leq i \leq j} D(p_i, \sigma)^{-1} = \left(\min_{1 \leq i \leq j} D(p_i, \sigma) \right)^{-1} = (j \Upsilon(\sigma) \Delta(\sigma))^{-1},$$

Lemma 2.2, yields

$$s_1(P^\dagger) \leq (\sqrt{j} \Upsilon(\sigma) \Delta(\sigma))^{-1},$$

because $s_i(A^\top) = s_i(A)$ for any matrix A . The stated bound on $s_j(P)$ follows from Lemma 2.3. \square

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 2.2, $s_1(P) \leq \sqrt{j} \Delta(\sigma)$, 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 . Roughly speaking, thickness imparts a kind of stability on the geometric properties of a simplex. This is exactly what is required when we want to show that a small change in a simplex will not yield a large change in some geometric quantity of interest. For example, we will use Lemma 2.4 in the demonstration of Lemma 4.1, which is the technical lemma related to the stability of the space of circumcentres of a simplex. Lemma 2.4 also facilitates a concise demonstration of Whitney's angle bound:

Proof of Lemma 2.1 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 $\text{aff}(\sigma)$. Let W be the k -dimensional subspace parallel to H , and let $\pi : \mathbb{R}^m \rightarrow W$ be the orthogonal projection onto W .

We desire an upper bound on $\|u - \pi u\|$ for all unit vectors $u \in U$. 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

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

W is at a distance less than η from H , because $p_0 \in W$ and $d_{\mathbb{R}^m}(p_i, H) \leq \eta$ for all $0 \leq i \leq j$. It follows that $\|v_i - \pi v_i\| \leq 2\eta$, and Lemma 2.2 yields

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

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 $K \subset 2^P$ 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 $\iota : P \hookrightarrow \mathbb{R}^m$. (We assume ι is injective on P , and so do not distinguish between P and $\iota(P)$.) For each abstract simplex $\sigma \in K$, we have an associated geometric simplex $\text{conv}(\iota(\sigma))$, 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 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 $\sigma \cap \tilde{\sigma} \neq \emptyset$, then $\tau = \sigma \cap \tilde{\sigma}$ is a simplex and $\tau \leq \sigma$ and $\tau \leq \tilde{\sigma}$. Observe that the simplices in a geometric simplicial complex are necessarily non-degenerate. 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 *dimension of a complex* K is the largest dimension of the simplices in K . We say that K is an m -complex, to mean that it is of dimension m . The complex K is a *pure m -complex* if it is an m -complex, and every simplex in K is the face of an m -simplex.

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 continuous map $\iota : |K| \rightarrow \mathbb{R}^m$, 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 $\text{star}(K')$ when there is no risk of ambiguity, otherwise we also specify the parent complex, as in $\text{star}(K'; K)$. A simple example of the star of a complex is depicted in Figure 2.




Figure 2: The star of a subcomplex $K' \subset K$ is the subcomplex $\text{star}(K') \subset K$ that consists all the simplices that share a face with K' (this includes all of K' itself), and all the faces of these simplices. Here we show an embedded 2-complex, with all 2-simplices shaded. The subcomplex K' consists of the two indicated triangles, and their faces (blue). The simplices of $\text{star}(K')$ are shown in bold (red and blue). The other simplices do not belong to $\text{star}(K')$ (black).

A *triangulation* of $P \subset \mathbb{R}^m$ is an embedded complex K with vertices P such that $|K| = \text{conv}(P)$. 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 2.5 (Triangulation at a point) A complex K is a *triangulation at* $p \in \mathbb{R}^m$ if:

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

In a general complex Condition 4 above is not a local property, however in the case of Delaunay complexes that interests us here, local conditions are sufficient to verify the condition, as we will show in Section 3.2.1. Observe also that Condition 4 also precludes intersections with degenerate simplices, since such a simplex would have a face that violates the condition.

If σ is a simplex with vertices in P , then any map $\zeta : P \rightarrow \tilde{P} \subset \mathbb{R}^m$ defines a simplex $\zeta(\sigma)$ whose vertices in \tilde{P} are the images of vertices of σ . If K is a complex on P , and \tilde{K} is a complex on \tilde{P} , then ζ induces a *simplicial map* $K \rightarrow \tilde{K}$ 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 $K \stackrel{\zeta}{\cong} \tilde{K}$ if we wish to emphasise that the correspondence is given by ζ .

A simplicial map $\zeta : K \rightarrow \tilde{K}$ defines a continuous map $\zeta : |K| \rightarrow |\tilde{K}|$, by barycentric interpolation on each simplex $\sigma \in K$. We observe the following consequence of Brouwer's invariance of domain:

Lemma 2.6 Suppose K is a complex with vertices $P \subset \mathbb{R}^m$, and \tilde{K} a complex with vertices $\tilde{P} \subset \mathbb{R}^m$. Suppose also that K is a triangulation at $p \in P$, and that $\zeta : P \rightarrow \tilde{P}$ induces an injective simplicial map $\text{star}(p) \rightarrow \text{star}(\zeta(p))$. If $\text{star}(\zeta(p))$ is embedded, then

$$\zeta(\text{star}(p)) = \text{star}(\zeta(p)),$$

and $\zeta(p)$ is an interior point of \tilde{P} .

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

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

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

3 Parameterized genericity

In this section we examine the Delaunay complex of $P \subset \mathbb{R}^m$, 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 $\partial B \cap P$. The *Delaunay complex* is the set of Delaunay simplices, and is denoted $\text{Del}(P)$.

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

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

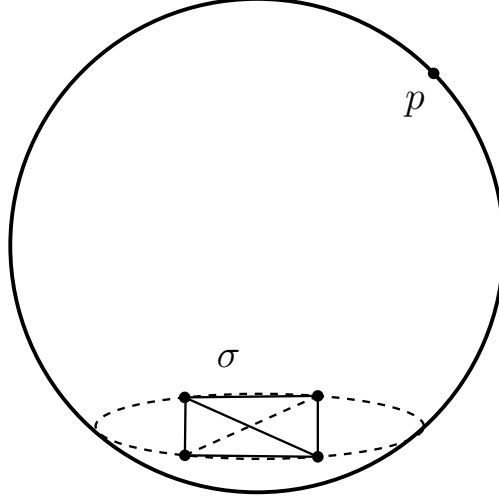


Figure 3: Lemma 3.2: If the affine hull of σ is not full dimensional, then a Delaunay ball has freedom to expand, and σ must be the face of a higher dimensional Delaunay simplex.

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

Proof Suppose $\dim \text{aff}(\sigma) < m$. 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 $\text{aff}(\sigma)$. 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 $\text{aff}(\sigma) = \text{aff}(\mathbf{P})$, and thus $\dim \text{aff}(\mathbf{P}) < m$. 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 $\text{conv}(\mathbf{P})$, then there must be non-empty balls centred on ℓ at either side of c . If $p \in \mathbf{P} \setminus \text{aff}(\sigma)$ 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 $\tau \in \text{Del}(\mathbf{P})$ is a j -simplex that is not a boundary simplex, and $q \in \mathbf{P} \setminus \tau$, 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 $\text{conv}(\mathbf{P})$, 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 $P \subset U \subset \mathbb{R}^m$ and $d : U \times U \rightarrow \mathbb{R}$ is a metric, then we define the Delaunay complex $\text{Del}_d(P)$ 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 $\text{Del}_d(P)$ consists of all the simplices which are circumscribed by a Delaunay ball with respect to the metric d . The simplices of $\text{Del}_d(P)$ are, possibly degenerate, geometric simplices in \mathbb{R}^m . As for $\text{Del}(P)$, $\text{Del}_d(P)$ has the combinatorial structure of an abstract simplicial complex, but unlike $\text{Del}(P)$, $\text{Del}_d(P)$ 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 $q \in P \setminus \sigma$. 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 $\delta \geq 0$.




Figure 4: A Delaunay simplex σ is δ -protected if it has a Delaunay ball $B_{\mathbb{R}^m}(c; r)$ such that $\overline{B_{\mathbb{R}^m}(c; r + \delta)} \cap (P \setminus \sigma) = \emptyset$.

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

Observe that we have employed a strict inequality in the definition of δ -protection. In particular, a δ -generic point set is generic even when $\delta = 0$. In order for the quantity δ to be meaningful, it should be considered with respect to a sampling radius ϵ for P . We will always assume that $\delta \leq \epsilon$.

In his seminal work, Delaunay [Del34] demonstrated that if there is no empty ball with $m + 2$ points from P on its boundary, then $\text{Del}(P)$ is realized as a simplicial complex in \mathbb{R}^m . In other words $\text{Del}(P)$ is an embedded complex, and in fact it is a triangulation of P , 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 $\text{Del}(\mathbf{P})$, as shown by Lemma 3.5 below.

In particular, if \mathbf{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 $\tilde{\delta} > 0$.

3.2.1 Local Delaunay triangulation

Delaunay avoided boundary complications by assuming a periodic point set, but we are particularly interested in the case where the point sets come from local patches of a well-sampled compact manifold without boundary. Periodic boundary conditions are not appropriate in this setting, but this is not a problem because, as we show here, Delaunay's argument applies locally.

Delaunay's proof that the Delaunay complex of a generic periodic point set is a triangulation of \mathbb{R}^m consists of two observations. First it is observed that if two Delaunay simplices intersect, then they intersect in a common face. This shows that $\text{Del}(\mathbf{P})$ is embedded. The argument is not complicated by the presence of boundary points:

Lemma 3.5 (Embedded star) Suppose $\text{aff}(\mathbf{P}) = \mathbb{R}^m$ and $p \in \mathbf{P}$. If all the m -simplices in $\text{star}(p; \text{Del}(\mathbf{P}))$ are protected, then $\text{star}(p; \text{Del}(\mathbf{P}))$ is embedded, and it is a pure m -complex.

Proof We first observe that the m -simplices in $\text{star}(p)$ are not degenerate. If σ^m is degenerate, then by Lemma 3.2, there is a simplex τ with $\text{aff}(\tau) = \mathbb{R}^m$, and $\sigma^m < \tau$. We have $\tau \in \text{star}(p)$, since $p \in \tau$. An affinely independent set of $m + 1$ vertices from τ defines a non-degenerate m -simplex $\tilde{\sigma}^m < \tau$, and since its unique circumball is also a Delaunay ball for τ , it cannot be protected, a contradiction.

Now suppose that $\sigma, \tau \in \text{star}(p)$ and $\sigma \cap \tau \neq \emptyset$. We need to show that they intersect in a common face. By Lemma 3.2, we may assume that σ and τ are m -simplices. Assume $\sigma \neq \tau$, and let B_1 and B_2 be the Delaunay balls for σ and τ . Then $\text{aff}(\partial B_1 \cap \partial B_2)$ defines an $(m - 1)$ -flat, H . Since B_1 and B_2 are empty balls, 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 $\partial B_1 \cap \partial B_2$. \square

The second observation Delaunay made is that, in the case of a periodic (infinite) point set, every $(m - 1)$ -simplex is the face of two m -simplices (Lemma 3.2). The implication here is that $\text{Del}(\mathbf{P})$ cannot have a boundary, and therefore must cover \mathbb{R}^m . Here we flesh out the argument for our purposes: If an embedded finite complex contains m -simplices then its topological boundary must contain $(m - 1)$ -simplices. We first observe that the topological boundary of an embedded complex is defined by a subcomplex:

Lemma 3.6 (Boundary complex) If K is an embedded (finite) complex in \mathbb{R}^m , then the topological boundary of $|K| \subset \mathbb{R}^m$ is defined by a subcomplex: $\partial|K| = |\text{bd}(K)|$, where the subcomplex $\text{bd}(K) \subset K$ is called the *boundary complex* of K .

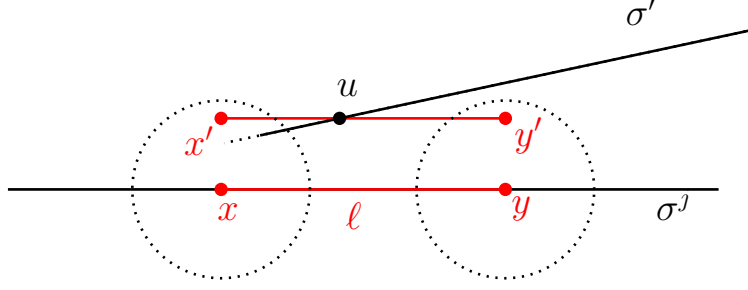


Figure 5: Diagram for the proof of Lemma 3.6.

Proof Since K is finite, $\partial|K|$ is contained in $|K|$. Suppose $x \in \partial|K|$. Then $x \in \text{int } \sigma^j$ for some $\sigma^j \in K$. We wish to show that $\sigma^j \subset \partial|K|$. Suppose to the contrary that $y \in \text{int } \sigma^j$, but y does not belong to $\partial|K|$. This means that $y \in \text{int } |K|$.

Consider the segment $\ell = [x, y] \subset \text{int } \sigma^j$. Let $Z \subset K$ be the subcomplex consisting of those simplices that do not contain σ^j . Let

$$r_1 = \min_{\sigma \in Z} d_{\mathbb{R}^m}(\ell, \sigma).$$

Choosing $r \leq r_1$, and $x' \in B_{\mathbb{R}^m}(x; r) \setminus |K|$, let $y' = y + (x' - x)$. Since $y \in \text{int } |K|$, we may assume that r is small enough so that $y' \in \text{int } |K|$.

Consider the segment $\ell' = [x', y']$. By construction, $\ell' \cap |Z| = \emptyset$. However, consider the point $u \in \text{int } \ell'$ that is the point in $\ell' \cap |K|$ that is closest to x' . The point u lies in the interior of some simplex $\sigma' \in K$, but we cannot have $\sigma^j \leq \sigma'$. Indeed if this were the case, x' would lie in $\text{aff}(\sigma')$, and so $u \in \partial\sigma'$, contradicting the assumption that $u \in \text{int } \sigma'$.

But this means that $\sigma' \in Z$, which contradicts the fact that $\ell' \cap |Z| = \emptyset$. Therefore we must have $y \in \partial|K|$ for all $y \in \text{int } \sigma^j$.

Finally, observe that if $\tau < \sigma^j$, then $\tau \subset \partial|K|$, since $\partial|K|$ is closed. \square

Lemma 3.7 (Pure boundary complex) If K is a (finite) pure m -complex embedded in \mathbb{R}^m , then its boundary complex is a pure $(m - 1)$ -complex.

Proof Since K is finite, $\text{bd}(K)$ is nonempty; it contains at least the vertices in $\partial\text{conv}(|K|)$. We will show that if $\sigma^j \in \text{bd}(K)$, is a j -simplex, with $0 \leq j < m - 1$, then there is a $\sigma^k \in \text{bd}(K)$ with $\sigma^j < \sigma^k$. The result then follows, since $\text{bd}(K)$ cannot contain m -simplices, because K is embedded.

Suppose $\sigma^j \in \text{bd}(K)$, and $x \in \text{int } \sigma^j$. Let $Z \subset K$ be the subcomplex consisting of simplices that do not contain σ^j , and let

$$r = \min_{\sigma \in Z} d_{\mathbb{R}^m}(x, \sigma).$$

Let $B = B_{\mathbb{R}^m}(x; r)$, and choose $y \in B \setminus |K|$. Let F be the $(m - j)$ -dimensional affine space orthogonal to $\text{aff}(\sigma^j)$ and containing y , and let $S^{m-j-1} = F \cap \partial B_{\mathbb{R}^m}(x; r')$, where $r' = \|x - y\|$. See Figure 6.

Since K is pure, there is an m -simplex σ^m with $\sigma^j < \sigma^m$. We have $\sigma^m \cap S^{m-j-1} \neq \emptyset$. Indeed, choose $w \in \text{int } \sigma^m$, and $u \in \sigma^j$ different from x , and observe that the plane Q defined by x, w, u intersects $B \cap \sigma^m$ in a semi-disk, by construction of B . By the construction of S^{m-j-1} , it must intersect this semidisk.

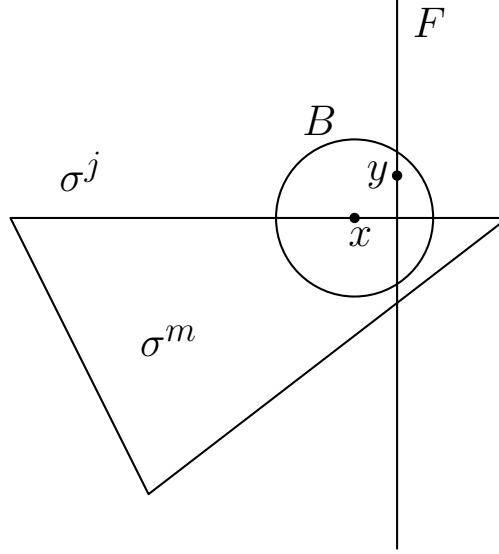


Figure 6: Diagram for the proof of Lemma 3.7.

Let $z \in S^{m-j-1}$ be a point that minimises the geodesic distance in S^{m-j-1} to y . Then $z \in \partial|K|$. Thus $z \in \text{int } \sigma^k$ for some $\sigma^k \in \text{bd}(K)$, and since $z \in B$, σ^k cannot belong to Z . Thus $\sigma^j \leq \sigma^k$, but since $S^{m-j-1} \cap \sigma^j = \emptyset$, we have $\sigma^j < \sigma^k$. \square

From Lemma 3.2, Lemma 3.5, and Lemma 3.7, one can verify that if P is generic then $\partial\text{conv}(P) = \partial|\text{Del}(P)|$, and thus obtain the standard result that $\text{Del}(P)$ is a triangulation of P . However, we are interested localizing the result, without the assumption that the entire point set is generic. We have the following local version of Delaunay's triangulation result:

Lemma 3.8 (Local Delaunay triangulation) If $p \in P$ is an interior point, and the Delaunay m -simplices incident to p are protected, then $\text{Del}(P)$ is a triangulation at p .

Proof By Lemma 3.5, $\text{star}(p)$ is a pure m -complex, and it is embedded. It follows then from Lemma 3.7, that the boundary complex $\text{bd}(\text{star}(p))$ is a pure $(m-1)$ -complex. Thus p cannot belong to $\text{bd}(\text{star}(p))$. Indeed, it follows from Lemma 3.2 that any $(m-1)$ -simplex $\sigma \in \text{star}(p)$ is the face of at least two m -simplices in $\text{Del}(P)$, and if $p \in \sigma$, then both of these m -simplices belong to $\text{star}(p)$, and are embedded, with intersection σ . Thus p cannot belong to an $(m-1)$ -simplex in $\text{bd}(\text{star}(p))$, and therefore $p \in \text{int } |\text{star}(p)|$.

It remains to verify Condition 4 of Definition 2.5. The argument is similar to the proof of Lemma 3.5: Suppose $x \in (\text{int } \tau) \cap \sigma$ for $\sigma \in \text{star}(p)$. We may assume that σ is an m -simplex. Then consider the Delaunay balls B_1 for σ and B_2 for τ . If $B_1 = B_2$, then, since σ is protected, τ must be a face of σ , and so belong to $\text{star}(p)$. Assume then that $B_1 \neq B_2$, and let H be the $(m-1)$ -flat defined by $\text{aff}(\partial B_1 \cap \partial B_2)$. Since B_1 is empty, $x \in \text{int } \tau$ cannot lie in the open half-space defined by H and containing σ . Since $x \in \sigma$ also, it must lie in H , and therefore all vertices of τ lie in $H \cap \partial B_2 = H \cap \partial B_1$, and so τ is a face of σ . \square

3.2.2 Safe interior simplices

We wish to consider the properties of Delaunay triangulations in regions which are comfortably in the interior of $\text{conv}(P)$, and avoid the complications that arise as we approach the boundary of the

point set. We introduce some terminology to facilitate this.

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

Lemma 3.9 If P is an ϵ -sample set, and $\sigma \in \text{Del}(P)$ has a vertex p such that $d_{\mathbb{R}^m}(p, \partial\text{conv}(P)) \geq 2\epsilon$, 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 $d_{\mathbb{R}^m}(x, \partial\text{conv}(P)) < \epsilon$. But then $d_{\mathbb{R}^m}(p, \partial\text{conv}(P)) \leq d_{\mathbb{R}^m}(p, x) + d_{\mathbb{R}^m}(x, \partial\text{conv}(P)) < 2\epsilon$, contradicting the hypothesis on p .

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

This suggests the following:

Definition 3.10 (Deep interior points) Suppose $P \subset \mathbb{R}^m$ is an ϵ -sample set. The subset $P_I \subset P$ consisting of all $p \in P$ with $d_{\mathbb{R}^m}(p, \partial\text{conv}(P)) \geq 4\epsilon$ is the set of *deep interior points*.

By Lemma 3.9, all the simplices that include a deep 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 deep interior points:




Figure 7: If P is δ -generic for P_J , then the safe interior simplices are the simplices in $\text{star}(P_J)$. Here P_J consists of the two large vertices (blue). They must be at least 4ϵ from $\partial\text{conv}(P)$ (which is not depicted in the figure). The safe interior simplices are shaded. All the simplices in $\text{star}(\text{star}(P_J))$ are δ -protected. These simplices have bold outlines (red), but are not necessarily shaded.

Definition 3.11 (δ -generic for P_J) The set $P \subset \mathbb{R}^m$ is δ -generic for P_J if $P_J \subseteq P_I$ and all the m -simplices in $\text{star}(\text{star}(P_J; \text{Del}(P)))$ are δ -protected. The *safe interior simplices* are the simplices in $\text{star}(P_J; \text{Del}(P))$.

It follows that, if P is δ -generic for P_J , with sampling radius ϵ , and τ is a safe 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 8(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 8(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 satisfied.

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 $\cos \alpha = \frac{d_{\mathbb{R}^m}(a, q)}{2r} \geq \frac{\delta}{2\epsilon}$, which means that $\alpha \leq \arccos \frac{\delta}{2\epsilon} \leq \frac{\pi}{2}$. Similarly, with $\beta = \angle cab$, we have $\beta \leq \arccos \frac{\lambda}{2\epsilon} \leq \frac{\pi}{2}$. Thus $\frac{\pi}{2} \leq \gamma = \alpha + \beta \leq \gamma'$, where

$$\gamma' = \arccos \frac{\delta}{2\epsilon} + \arccos \frac{\lambda}{2\epsilon}.$$

Since $\sin \gamma \geq \sin \gamma'$, when $\frac{\pi}{2} \leq \gamma \leq \gamma' \leq \pi$, we have

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

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

Since $\text{aff}(\tau) \subset H$, 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 . \square

We thus obtain a bound on the thickness of the safe interior simplices when P is δ -generic for P_J . Since Lemma 3.13 yields a lower bound of $\frac{\sqrt{3}\delta^2}{2\epsilon}$ on the altitudes of the safe interior simplices, and since $\Delta(\sigma) \leq 2\epsilon$, we have that $\Upsilon(\sigma) \geq \frac{\sqrt{3}\delta^2 m}{4\epsilon^2}$ for all safe interior σ . If $\delta = \nu_0 \epsilon$, we obtain a constant thickness bound.

Theorem 3.14 (Thickness from protection) If $P \subset \mathbb{R}^m$ is δ -generic for P_J with $\delta = \nu_0 \epsilon$, where ϵ is a sampling radius for P , then the safe interior simplices are Υ_0 -thick, with

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

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

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

then there is a point $c \in N(\sigma)$ 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 (2) 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

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

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:

$$P^\top C(\sigma) = b. \quad (3)$$

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

$$P^\top x = b - x_a. \quad (4)$$

From Equations (3) and (4) we have

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

From Equation (2), it follows that $\|x_a\| \leq \frac{\sqrt{m}\xi^2}{2}$, and from Lemmas 2.3 and 2.4 we have $\|P^{-1}\| \leq (\sqrt{m}\Upsilon(\sigma)\Delta(\sigma))^{-1}$, 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 $\text{aff}(\sigma)$. We observe that \hat{x} also must satisfy Equation (2), 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

$$\left| \|p_i - x\| - \|p_j - x\| \right| \leq \xi \quad \text{for all } i, j \leq k. \quad (5)$$

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 for some $\xi > 0$ the point $x \in \mathbb{R}^m$ is a ξ -centre such that

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

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. The situation is depicted schematically in Figure 9. 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:




Figure 9: The maps f_e and f (described in the main text) map circumcentres to the origin. Since $f_e^{-1}(0)$ contains a unique point, and $f_e(\partial B)$ lies far from the origin, a consideration of the degree of the mappings, together with the fact that $f_e(\partial B)$ and $f(\partial B)$ are close, reveals that $f^{-1}(0)$ cannot be empty, and thus B must contain a circumcentre of σ^m with respect to the metric d .

Lemma 4.4 (Circumcentres: metric perturbation) Let $U \subset \mathbb{R}^m$, and let $d : U \times U \rightarrow \mathbb{R}$ 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 $\sigma = [p_0, \dots, p_m] \subset U$ 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 $f_e : \overline{B} \rightarrow \mathbb{R}^m$ given by

$$f_e(x) = (d_{\mathbb{R}^m}(x, p_1) - d_{\mathbb{R}^m}(x, p_0), \dots, d_{\mathbb{R}^m}(x, p_m) - d_{\mathbb{R}^m}(x, p_0))^{\top}. \quad (6)$$

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 ,

$$f(x) = (d(x, p_1) - d(x, p_0), \dots, d(x, p_m) - d(x, p_0))^{\top}, \quad (7)$$

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 $F : \overline{B} \times [0, 1] \rightarrow \mathbb{R}^m$ between $f_e(x) = F(x, 0)$ and $f(x) = F(x, 1)$ with $F(x, t) \neq 0$ for all $x \in \partial\overline{B}$ and $t \in [0, 1]$.

Proof We define the homotopy $F : \overline{B} \times [0, 1] \rightarrow \mathbb{R}^m$ 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 (6).

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

$$J(f_e)_0 = -\frac{1}{R(\sigma)} P^\top, \tag{8}$$

where as usual P is the matrix whose columns are $p_i - p_0$. Since $\text{vol}(\sigma^m) = \frac{|\det(P)|}{m!}$, Equation (8) implies

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

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 $g : \overline{B} \rightarrow \mathbb{R}^m$ at a regular point $p \in g(B)$ is defined by

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

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 $h : \overline{B} \rightarrow \mathbb{R}^m$ is homotopic to g by a homotopy $H : \overline{B} \times [0, 1] \rightarrow \mathbb{R}^m$ 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 $\tilde{\sigma} = [\tilde{p}_0, \dots, \tilde{p}_m]$ 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 (7) by the function

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

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

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

4.2 Perturbations and protection

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

Definition 4.9 (Secure simplex) A simplex $\sigma \in \text{Del}(P)$ 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 safe interior simplices Lemma 3.12 and Theorem 3.14 allow us to consolidate some of these parameters with the ratio δ/ϵ :

Lemma 4.10 (Safe 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 safe interior m -simplices are secure, with $\mu_0 = \nu_0$, and $\Upsilon_0 = \frac{\sqrt{3}\nu_0^2}{4m}$.

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

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

then $\zeta(\sigma) = \tilde{\sigma} \in \text{Del}(\tilde{P})$ 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 $\sigma \in \text{Del}(\mathbf{P})$, 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 $\tilde{\mathbf{P}}$.

Let $q \in \mathbf{P}$ 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 $\text{conv}(\mathbf{P})$ and $d : U \times U \rightarrow \mathbb{R}$ 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 $\sigma \in \text{Del}(\mathbf{P})$ 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 $\text{Del}_d(\mathbf{P})$, 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 $\sigma \in \text{Del}(\mathbf{P})$, 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 $q \in \mathbf{P}$ 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. \square

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

A consequence of Delaunay’s triangulation result is that if a perturbation does not destroy any m -simplices in the Delaunay complex of a generic point set, then no new simplices are created either, and the complex is unchanged. More precisely we have:

Lemma 4.13 Suppose $P \subset \mathbb{R}^m$ is a generic sample set, and $Q \subseteq P$ is a subset of interior points. If $\zeta : P \rightarrow \tilde{P}$ is a perturbation such that $\zeta(\text{star}(Q; \text{Del}(P))) \subseteq \text{star}(\zeta(Q); \text{Del}(\tilde{P}))$, and every m -simplex $\tilde{\sigma}^m \in \zeta(\text{star}(Q))$ is protected in $\text{Del}(\tilde{P})$, then $\zeta(\text{star}(Q)) = \text{star}(\zeta(Q))$.

Proof Let $p \in Q$. By Lemma 3.5, $\text{star}(\zeta(p))$ is embedded, and by Lemma 3.8, $\text{Del}(P)$ is a triangulation at p . Since $\zeta : P \rightarrow \tilde{P}$ is injective, it follows that the simplicial map induced by ζ must be injective, and the result follows from Lemma 2.6. \square

Lemma 4.11 establishes bounds on a ρ -perturbation $\zeta : P \rightarrow \tilde{P}$ which will guarantee that if $Q \subset P$, and the simplices in $\text{star}(Q)$ are secure, then $\zeta(\text{star}(Q)) \subseteq \text{Del}(\tilde{P})$. Lemma 4.11 also guarantees that, if ρ is small enough, the m -simplices in $\text{star}(\zeta(Q); \text{Del}(\tilde{P}))$ will be protected. Thus if Q consists only of interior points of P , Lemma 4.13 applies. We have the following stability theorem for protected Delaunay triangulations:

Theorem 4.14 (Stability under point perturbation) Suppose $P \subset \mathbb{R}^m$ and $Q \subseteq P$ is a subset of interior points such that every m -simplex in $\text{star}(Q)$ is secure. If $\zeta : P \rightarrow \tilde{P}$ is a ρ -perturbation, with

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

then

$$\text{star}(Q; \text{Del}(P)) \stackrel{\zeta}{\cong} \text{star}(\zeta(Q); \text{Del}(\tilde{P})).$$

The ρ -relaxed Delaunay complex for P was defined by de Silva [dS08] by the criterion that $\sigma \in \text{Del}^\rho(P)$ 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 $q \in P$. Thus the simplices in $\text{Del}^\rho(P)$ 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 $P \subset \mathbb{R}^m$ and $Q \subseteq P$ is a set of interior points such that every m -simplex in $\text{star}(Q)$ is secure. If

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

then

$$\text{star}(Q; \text{Del}^\rho(P)) = \text{star}(Q; \text{Del}(P)).$$

Proof Suppose that $\sigma \in \text{star}(\mathbf{Q}; \text{Del}^\rho(\mathbf{P}))$. 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 $\zeta : \mathbf{P} \rightarrow \tilde{\mathbf{P}}$, and σ has become $\tilde{\sigma} \in \text{star}(\zeta(\mathbf{Q}); \text{Del}(\tilde{\mathbf{P}}))$. By Theorem 4.14, $\text{star}(\zeta(\mathbf{Q}); \text{Del}(\tilde{\mathbf{P}})) \stackrel{\zeta}{\cong} \text{star}(\mathbf{Q}; \text{Del}(\mathbf{P}))$, and therefore $\sigma \in \text{star}(\mathbf{Q}; \text{Del}(\mathbf{P}))$. \square

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 $\text{star}(\mathbf{Q}; \text{Del}_d(\mathbf{P}))$ that do not already exist in $\text{star}(\mathbf{Q}; \text{Del}(\mathbf{P}))$.

Lemma 4.16 Suppose $\text{conv}(\mathbf{P}) \subseteq U \subset \mathbb{R}^m$ and $d : U \times U \rightarrow \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $\mathbf{Q} \subseteq \mathbf{P}$ is a set of interior points such that every m -simplex $\sigma \in \text{star}(\mathbf{Q})$ 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

$$\text{star}(\mathbf{Q}; \text{Del}_d(\mathbf{P})) \subseteq \text{star}(\mathbf{Q}; \text{Del}(\mathbf{P})).$$

Proof Let $B(c; r)$ be a Delaunay ball for simplex $\sigma \in \text{star}(\mathbf{Q}; \text{Del}_d(\mathbf{P}))$. Then $d(c, p) \leq d(c, q)$ for all $p \in \sigma$, and $q \in \mathbf{P}$. 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 $q \in \mathbf{P}$, and therefore $\sigma \in \text{Del}^{2\rho}(\mathbf{P})$. The result now follows from Corollary 4.15. \square

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 $\text{conv}(\mathbf{P}) \subseteq U \subset \mathbb{R}^m$ and $d : U \times U \rightarrow \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $\mathbf{Q} \subseteq \mathbf{P}$ is a set of interior points such that every m -simplex $\sigma \in \text{star}(\mathbf{Q})$ 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

$$\text{star}(\mathbf{Q}; \text{Del}_d(\mathbf{P})) = \text{star}(\mathbf{Q}; \text{Del}(\mathbf{P})).$$

Using Lemma 4.10, and recognizing that the safe 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 \mathbf{P} is δ -generic for \mathbf{P}_J , with sampling radius ϵ and $\delta = \nu_0 \epsilon$. Suppose also that $\text{conv}(\mathbf{P}) \subseteq U$, and $d : U \times U \rightarrow \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. If

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

then

$$\text{star}(\mathbf{P}_J; \text{Del}_d(\mathbf{P})) = \text{star}(\mathbf{P}_J; \text{Del}(\mathbf{P})).$$

5 Conclusions

We have quantified the close relationship between the genericity of a point set, the quality of the simplices in the Delaunay complex, and the stability of the Delaunay complex under perturbation. The problem of poorly shaped simplices in a higher dimensional Delaunay complex can be seen as a manifestation of point sets that are close to being degenerate. The introduction of thickness as a geometric quality measure for simplices facilitated the stability calculations, which develop around a consideration of the circumcentres of a simplex in the presence of a perturbation.

We considered a point set $P \subset \mathbb{R}^m$ meeting a sampling radius ϵ and showed a constant bound on the thickness of the Delaunay simplices provided P is δ -generic with $\delta = \nu_0 \epsilon$ for some constant ν_0 . The question then arises: What is the least upper bound on the feasible ν_0 as a function of the dimension m ?

In a companion paper [BDG13a], we develop a perturbation algorithm which produces a δ -generic point set from a given ϵ -sample set. Since the triangulation results of the current work are localised, we can extend the perturbation algorithm to construct Delaunay triangulations of abstract Riemannian manifolds that are not necessarily embedded in an ambient space, as we have shown in subsequent work [BDG13b]. The idea is that a manifold can be locally well approximated by Euclidean space, so we fit together local Euclidean Delaunay patches where the Euclidean metric varies slightly between patches. This is where the stability of the Delaunay patches is important. In this setting we can also accommodate variations in the sampling radius between neighbouring patches. Thus the algorithm is able to triangulate sample sets whose sampling radius is defined by a Lipschitz sizing function.

Acknowledgements

The authors would like to thank David Cohen-Steiner for suggesting the proof technique employed in Lemma 4.4. Comments and suggestions from the reviewers considerably improved the final document; the authors gratefully acknowledge this contribution.

This work was partially supported by the CG Learning project. The project CG Learning acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open grant number: 255827.

The research has also received funding from the European Research Council (ERC) under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement No. 339025 GUDHI (Algorithmic Foundations of Geometry Understanding in Higher Dimensions).

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. 1
- [BDG13a] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay stability via perturbations. Research Report RR-8275, INRIA, 2013. 1, 25
- [BDG13b] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay triangulation of manifolds. Research Report RR-8389, INRIA, 2013. 25
- [BG11] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. Technical Report RR-7142 v.3, INRIA, 2011. (To appear in DCG). 1

- [BS04] D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: nearest neighbor relations for imprecise points. In *SODA*, pages 410–419, 2004. 1
- [BWY11] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic Delaunay mesh generation. Technical Report RR-7712, INRIA, 2011. 1, 3
- [CDE⁺00] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S. H Teng. Sliver exudation. *Journal of the ACM*, 47(5):883–904, 2000. 3
- [CDR05] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA*, pages 1018–1027, 2005. 1
- [CG12] G. D. Cañas and S. J. Gortler. Duals of orphan-free anisotropic Voronoi diagrams are embedded meshes. In *SoCG*, pages 219–228, New York, NY, USA, 2012. ACM. 1
- [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. 20
- [Del34] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934. 1, 10, 11, 23
- [dS08] V. de Silva. A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata*, 135:39–64, 2008. 23
- [Dug66] J. Dugundji. *Topology*. Allyn and Bacon, Inc., Boston, 1966. 8
- [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. 1
- [ES97] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Int. J. Comput. Geometry Appl.*, 7(4):365–378, 1997. 1
- [FKMS05] S. Funke, C. Klein, K. Mehlhorn, and S. Schmitt. Controlled perturbation for Delaunay triangulations. In *SODA*, pages 1047–1056, 2005. 1
- [Li03] X.-Y. Li. Generating well-shaped d -dimensional Delaunay meshes. *Theoretical Computer Science*, 296(1):145–165, 2003. 3
- [LL00] G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *SoCG*, pages 341–349, 2000. 1
- [LS03] F. Labelle and J. R. Shewchuk. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SoCG*, pages 191–200, 2003. 1
- [Mun68] J. R. Munkres. *Elementary differential topology*. Princeton University press, second edition, 1968. 3
- [Mun84] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, 1984. 6
- [TB97] L.N. Trefethen and D. Bau. *Numerical linear algebra*. Society for Industrial Mathematics, 1997. 4
- [Whi57] H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957. 3