Equating the witness and restricted Delaunay complexes

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh, Steve Oudot

Equating the witness and restricted Delaunay complexes

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh, and Steve Y. Oudot

2011

Abstract

It is a well-known fact that the restricted Delaunay and witness complexes may differ when the landmark and witness sets are located on submanifolds of $\mathbb{R}^d$ of dimension 3 or more. Currently, the only known way of overcoming this issue consists of building some crude superset of the witness complex, and applying a greedy sliver exudation technique on this superset. Unfortunately, the construction time of the superset depends exponentially on the ambient dimension, which makes the witness complex based approach to manifold reconstruction impractical. This work provides an analysis of the reasons why the restricted Delaunay and witness complexes fail to include each other. From this a new set of conditions naturally arises under which the two complexes are equal.

1   Introduction

Various subcomplexes of the Delaunay triangulation have been used with success for approximating k-submanifolds of $\mathbb{R}^d$ from finite collections of sample points. Perhaps the most popular one in small dimensions ($k \in \{1, 2\}$ and $d \in \{2, 3\}$) is the so-called *restricted Delaunay complex*, defined as the subcomplex spanned by those Delaunay simplices whose dual Voronoi faces intersect the manifold. Its main attraction is that it can be shown to be a faithful approximation to the manifold underlying the data points, in terms of topology (ambient isotopy), of geometry (Hausdorff proximity), and of differential quantities (normals, curvatures, etc), all this under sampling assumptions which only constrain the sampling density [1, 2, 8, 3]. These qualities explain its success in the context of curve and surface meshing or reconstruction, where it is used either as a data structure for the algorithms, or as a mathematical tool for their analysis, or both — see [10] for a survey.

The story becomes quite different when the data is sitting in higher dimensions, where two major bottlenecks appear:

(i) The nice structural properties mentioned above no longer hold when the dimension k of the submanifold is 3 or more. In particular, normals may become arbitrarily wrong [12], and more importantly the topological type of the complex may deviate significantly from the one of the manifold [7]. These shortcomings bring into question the usefulness of the complex as a theoretical tool.

(ii) It is not known how to compute the restricted Delaunay complex without computing the full-dimensional Delaunay triangulation or at least its restriction to some local d-dimensional neighborhood. The resulting construction time incurs an exponential dependence on the ambient dimension d, which makes the complex a prohibitively costly data structure in practice. Another issue is the high degree of the algebraic operations involved in the construction of

Delaunay triangulations. Specifically, we need to evaluate signs of polynomials of degree $d+2$ in the input variables, which is prohibitive for large d.

To address problem (i), Cheng *et al.* [12] suggested to use *weighted* Delaunay triangulations. The intuition underlying their approach is simple: when the restricted Delaunay complex contains badly shaped simplices, called *slivers*, its behavior in their vicinity may be arbitrarily bad: wrong normals, wrong local homology, and so on. By carefully assigning weights to the data points, one can remove all the slivers from the restricted Delaunay complex and thus have it recover its good structural properties. This idea was carried on in subsequent work [7, 5], and it is now considered a fairly common technique in Delaunay-based manifold reconstruction.

Yet, the question of computing a good set of weights given the input point cloud remains. This question is closely connected to problem (ii) above, since determining which simplices are the slivers to be removed requires that the restricted Delaunay complex be computed first. To address this issue, it has been proposed to build some superset of the restricted Delaunay complex, from which the slivers are removed [7, 12]. After the operation the superset becomes equal to the restricted Delaunay complex, and thus it shares its nice properties. Unfortunately, the supersets proposed so far were pretty crude, and their construction times depended exponentially on the ambient dimension d, which made the approach intractable in practice.

To circumvent the building time issue, Boissonnat and Ghosh [5] proposed to use a different sub-complex of the Delaunay triangulation, called the *tangential complex*, whose construction reduces to computing local Delaunay triangulations in (approximations of) the k-dimensional tangent spaces of the manifold at the sample points. Once these local triangulations have been computed, the tangential complex is assembled by gluing them together. Consistency issues between the local triangulations may appear, which are solved once again by a careful weight assignment over the set of data points to remove slivers. The benefit is that the slivers to be removed are determined directly from the complex, not from some superset, so the complexity of the sliver removal phase reduces to a linear dependence on the ambient dimension d, while keeping an exponential dependence on the intrinsic dimension k. This makes the approach tractable under the common assumption that the data points live on a manifold with small intrinsic dimension, embedded in some potentially very high-dimensional space. The obtained complex is not the restricted Delaunay complex in general, and the question of whether the latter can be effectively retrieved remains open. Also, the algebraic degree of the polynomials involved in the construction of the tangential complex is $k+2$, which is a neat improvement over $d+2$ but still limits the practical use of the tangential complex to manifolds of small dimensions.

Enter the witness complex. In light of the apparent hardness of manifold reconstruction, researchers have turned their focus to the somewhat easier problem of inferring some topological invariants of the manifold without explicitly reconstructing it. Their belief was that more lightweight data structures would be appropriate for this simpler task, and it is in this context that Vin de Silva introduced the *witness complex* [15]. Given a point cloud W, his idea was to carefully select a subset L of landmarks on top of which the complex would be built, and to use the remaining data points to drive the complex construction. More precisely, a point $w \in W$ is called a *witness* for a simplex $\sigma \in 2^L$ if no point of L is closer to w than w is to the vertices of σ, i.e. if there is a ball centered at w that includes the vertices of σ, but no other points from L. The witness complex is then the largest abstract simplicial complex that can be assembled using only witnessed simplices. The geometric test for being a witness can be viewed as a simplified version of the classical Delaunay predicate, and its great advantage is to require a mere comparison of (squared) distances

(hence the evaluation of the signs of polynomials of degree 2 in the input variables). As a result, witness complexes can be built in arbitrary metric spaces, and the construction time is bound to the size of the input point cloud rather than to the dimension d of the ambient space.

Since its introduction, the witness complex has attracted interest [4, 7, 9, 11, 14, 20], which can be explained by its close connection to the Delaunay triangulation and restricted Delaunay complex. In his seminal paper [15], de Silva showed that the witness complex is always a subcomplex of the Delaunay triangulation $\text{Del}(L)$, provided that the data points lie in some Euclidean space or more generally in some Riemannian manifold of constant sectional curvature. With applications to reconstruction in mind, Attali *et al.* [4] and Guibas and Oudot [20] considered the case where the data points lie on or close to some k-submanifold of $\mathbb{R}^d$, and they showed that the witness complex is equal to the restricted Delaunay complex when $k = 1$, and a subset of it when $k = 2$. Unfortunately, the case of 3-manifolds is once again problematic, and it is now a well-known fact that the restricted Delaunay and witness complexes may differ significantly (no respective inclusion, different topological types, etc) when $k \geq 3$ [22]. To overcome this issue, Boissonnat, Guibas and Oudot [7] resorted to the sliver removal technique described above on some superset of the witness complex, whose construction incurs an exponential dependence on d. The state of affairs as of now is that the complexity of witness complex based manifold reconstruction is exponential in d, and whether it can be made only polynomial in d (while still exponential in k) remains an open question.

2   Background

We will be working primarily within the context of the ambient space $\mathbb{R}^d$, and distances are given by the standard norm $\|\cdot\|$. The distance between a point p and a set $X \subset \mathbb{R}^d$, is the infimum of the distances between p and the points of X, and is denoted $d_{\mathbb{R}^d}(p, X)$. We refer to the distance between two points a and b as $\|b - a\|$ or $d_{\mathbb{R}^d}(a, b)$ as convenient. A ball $B_{\mathbb{R}^d}(c; r) = \{x \mid \|x - c\| < r\}$ is open, and $\overline{B}_{\mathbb{R}^d}(c; r)$ is its topological closure.

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

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 an $m \times j$ matrix of rank $j \leq m$, 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 $\mathbb{R}^d$, 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. We will make use of the following observation:

Lemma 2.3 Suppose U and V are subspaces of $\mathbb{R}^d$ with $\dim U \leq \dim V$. Then

$$\angle(U, V) = \angle(V^\perp, U^\perp),$$

where $U^\perp$ and $V^\perp$ are the orthogonal complements of U and V in $\mathbb{R}^d$.

Proof   Suppose $\angle(U, V) = \alpha$. Let $v_* \in V^\perp$ be a unit vector. There are unit vectors $u \in U$, and $u_* \in U^\perp$ such that $v_* = au + bu_*$. We will show that $\angle(v_*, u_*) \leq \alpha$. First note that this angle is complementary to $\angle(v_*, u)$, i.e.,

$$\angle(v_*, u_*) = \frac{\pi}{2} - \angle(v_*, u). \tag{1}$$

There is a unit vector $v \in V$ such that $\angle(u, v) = \alpha_0 \leq \alpha$. Viewing angles between unit vectors as distances on the unit sphere, we exploit the triangle inequality: $\angle(v_*, v) \leq \angle(v_*, u) + \angle(u, v)$, from whence

$$\angle(v_*, u) \geq \frac{\pi}{2} - \alpha_0.$$

Using this expression in Equation (1), we find

$$\angle(v_*, u_*) \leq \alpha_0 \leq \alpha,$$

which implies, since v_* was chosen arbitrarily, that $\angle(V^\perp, U^\perp) \leq \angle(U, V)$.

Since $\dim V^\perp \leq \dim U^\perp$, and the orthogonal complement is a symmetric relation on subspaces, the same argument yields the reverse inequality. $\qquad\square$

2.1   Sampling parameters

Given two subsets W and X of $\mathbb{R}^d$, the symmetric *Hausdorff distance* between them is defined by

$$d_H(W, X) = \max\{\sup_{w \in W} d_{\mathbb{R}^d}(w, X), \sup_{x \in X} d_{\mathbb{R}^d}(x, W)\}.$$

If W is a finite set, we say it is an ϵ-*sample* of X if $d_H(W, X) < \epsilon$. We say that ϵ is the *sampling radius*, because any ball of radius ϵ centred on X must contain a point of W. Note that W is not required to be a subset of X. We say that W is β-*sparse* if for all $w, w' \in W$, $d_{\mathbb{R}^d}(w, w') \geq \beta$.

Our main structural result concerns sampling on $\mathbb{M}$, a differentiable compact k-dimensional submanifold of $\mathbb{R}^d$. For $p \in \mathbb{M}$, $T_p\mathbb{M}$ denotes the tangent space at p, which we usually identify with a k-flat in the ambient space. The normal space, $N_p\mathbb{M}$, is the orthogonal complement of $T_p\mathbb{M}$ in $T_p\mathbb{R}^d$, and we likewise treat it as the affine space of dimension $d - k$ orthogonal to $T_p\mathbb{M} \subset \mathbb{R}^d$.

A ball $B = B_{\mathbb{R}^d}(c; r)$ is a *medial ball* at p if it is tangent to $\mathbb{M}$ at p, and $B \cap \mathbb{M} = \emptyset$, and it is maximal in the sense that any ball containing B and centred on the line containing p and c must either intersect $\mathbb{M}$ or coincide with B. The *local reach* at p is the infimum of the radii of the medial balls at p, and the *reach* of $\mathbb{M}$, denoted $\mathrm{rch}(\mathbb{M})$, is the infimum of the local reach over all points of $\mathbb{M}$. The reach of $\mathbb{M}$ imposes an upper bound on the sampling radius for our results; we require $\mathbb{M}$ to have positive reach. This comes via our use of the following observation from Giesen and Wagner [19, Lemma 6]:

Lemma 2.4 If $\hat{p}, \hat{c} \in \mathbb{M}$ and $\|\hat{p} - \hat{c}\| \leq r < \mathrm{rch}(\mathbb{M})$, then $d_{\mathbb{R}^d}(\hat{p}, T_{\hat{c}}\mathbb{M}) \leq \frac{r^2}{2\mathrm{rch}(\mathbb{M})}$.

2.2 Simplices and complexes

Although the structures of interest we will introduce in Section 2.3 are defined as abstract simplicial complexes, we will be working primarily with geometric simplices $\sigma \subset \mathbb{R}^d$. Given $L \subset \mathbb{R}^d$, and an abstract simplex $\sigma = \{p_0, \ldots, p_j\}$ whose vertices belong to L, the inclusion $L \hookrightarrow \mathbb{R}^d$ induces a natural mapping $\varphi : \sigma \to \mathbb{R}^d$ that is piecewise linear on a geometric realization of σ. In general, φ may not be an embedding, i.e., the vertices of σ may not be affinely independent in $\mathbb{R}^d$, however our sampling criteria will explicitly impose a thickness condition that allows us to consider $\sigma \subset \mathbb{R}^d$ as a geometric simplex, $[p_0, \ldots, p_j]$, without degeneracy. After the introduction of the abstract simplicial complexes in Section 2.3, σ will be considered as a geometric simplex.

The affine hull of a j-simplex σ is denoted $\mathrm{aff}(\sigma)$, and we also make use of the $(d-j)$-dimensional affine space $N(\sigma)$ composed of the centres of the balls that circumscribe σ. This space is orthogonal to $\mathrm{aff}(\sigma)$ and intersects it at the *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 $R(\sigma)$.

Other important geometric properties of σ include its diameter (i.e., its longest edge), $\Delta(\sigma)$; its shortest edge, $L(\sigma)$; and the ratio of its circumradius and shortest edge, $\Phi(\sigma) = R(\sigma)/L(\sigma)$.

We write $\tau \leq \sigma$ to indicate that τ is a (not necessarily proper) face of σ. For p a vertex of σ, we denote by σ_p the face opposite p. The *altitude* of p in σ is $D(p, \sigma) = d_{\mathbb{R}^d}(p, \mathrm{aff}(\sigma_p))$. Poorly shaped simplices are problematic for our purposes; we need to avoid them. Such simplices 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 $\Upsilon(\sigma) \geq \Upsilon_0$. The constant Υ_0 that bounds the thickness of the simplices plays an important role in our analysis.

Other parameters such as the volume [23], or the radius of the largest contained ball centred at the barycentre [21], can be used to characterize these poorly shaped (close to degenerate) simplices. We find a direct bound on the altitudes to be more convenient, in part due to the following consequence of Lemma 2.2:

Lemma 2.5 Let $\sigma = [p_0, \ldots, p_j]$ be a non-degenerate j-simplex in $\mathbb{R}^d$, and let P be the $d \times j$ matrix whose columns are $p_i - p_0$. Then

$$
s_j(P) \geq \sqrt{j}\Upsilon(\sigma)\Delta(\sigma). \tag{2}
$$

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

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

$\square$

The proof of Lemma 2.5 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 [23, 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 2.5 in the proof of Whitney's Lemma allows us to simplify the proof and sharpen the result:

Lemma 2.6 (Whitney) Suppose σ is a j-simplex whose vertices all lie within a distance ζ from a hyperplane, $H \subset \mathbb{R}^d$. Then

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

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

For any unit vector $u \in U$, $\sin \angle(\mathrm{aff}(\sigma), H) = \sin \angle(U, W) \leq \|u - \pi u\|$. Since the vectors $v_i = (p_i - p_0)$, $i \in \{1, \ldots, j\}$ form a basis for U, we may write $u = Pa$, where P is the $d \times j$ 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 $d_{\mathbb{R}^d}(p_i, H) \leq \zeta$ for all $0 \leq i \leq j$, W is at a distance less than ζ from H, and $\|v_i - \pi v_i\| \leq 2\zeta$. It follows then from Lemma 2.1 that
$$\|X\| \leq 2\sqrt{j}\zeta.$$
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 2.5. $\qquad\square$

Lemma 2.6, together with Lemma 2.4, implies the following useful observation:

Lemma 2.7 Suppose $\hat{c} \in \mathbb{M}$, and σ is a j-simplex such that $j \leq k = \dim \mathbb{M}$, and that each vertex $p \in \sigma$ has a unique closest point $\hat{p} \in \mathbb{M}$, with $\|p - \hat{p}\| < \epsilon$, and $\|\hat{p} - \hat{c}\| < r \leq \mathrm{rch}(\mathbb{M})$. If $\epsilon \leq \frac{C_0 r^2}{2\mathrm{rch}(\mathbb{M})}$, for some constant C_0, then

$$\sin \angle(\mathrm{aff}(\sigma), T_{\hat{c}}\mathbb{M}) \leq \frac{(1 + C_0)r^2}{\Upsilon(\sigma)\Delta(\sigma)\mathrm{rch}(\mathbb{M})}.$$

Proof   By Lemma 2.4, $d_{\mathbb{R}^d}(\hat{p}, T_{\hat{c}}\mathbb{M}) \leq \frac{r^2}{2\mathrm{rch}(\mathbb{M})}$. By the triangle inequality and the bound on ϵ, $d_{\mathbb{R}^d}(p, T_{\hat{c}}\mathbb{M}) \leq \frac{(1+C_0)r^2}{2\mathrm{rch}(\mathbb{M})}$. The result follows from Whitney's Lemma 2.6. $\qquad\square$

Thickness will also be used in a slightly different context. Specifically, as discussed in Section 3.2.1, a thickness assumption allows us to conclude that if c is the centre of a circumscribing ball for σ, then if the vertices of $\tilde{\sigma}$ lie close to those of σ, there will be a point $\tilde{c}$ close to c that is the centre of a circumscribing ball for $\tilde{\sigma}$. In order to exploit this observation, we will require a lower bound on the thickness of a perturbed simplex.

Lemma 2.8 Let $\sigma = [p_0, \ldots, p_j]$ and $\tilde{\sigma} = [\tilde{p}_0, \ldots, \tilde{p}_j]$ be j-simplices such that $\|\tilde{p}_i - p_i\| \leq \rho$ for all $i \in \{0, \ldots, j\}$. For any positive $K \leq 1$, if

$$\rho \leq \frac{(1 - K)\Upsilon(\sigma)^2 L(\sigma)}{14}, \tag{3}$$

then

$$D(\tilde{p}_i, \tilde{\sigma}) \geq KD(p_i, \sigma),$$

for all $i \in \{0, \ldots, j\}$. It follows that

$$\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma}) \geq K\Upsilon(\sigma)\Delta(\sigma),$$

and

$$\Upsilon(\tilde{\sigma}) \geq \left(1 - \frac{2\rho}{\Delta(\sigma)}\right) K\Upsilon(\sigma) \geq \frac{6}{7}K\Upsilon(\sigma).$$

Proof   Let $p, q \in \sigma$ with $\tilde{p}, \tilde{q}$ the corresponding vertices of $\tilde{\sigma}$. Let $v = p - q$ and $\tilde{v} = \tilde{p} - \tilde{q}$. Define $\theta = \angle(v, \mathrm{aff}(\sigma_p))$ and $\tilde{\theta} = \angle(\tilde{v}, \mathrm{aff}(\tilde{\sigma}_{\tilde{p}}))$. Since $\Upsilon(\sigma) \leq \Upsilon(\sigma_p)$, Whitney's Lemma 2.6 lets us bound $\angle(\mathrm{aff}(\sigma_p), \mathrm{aff}(\tilde{\sigma}_{\tilde{p}}))$ by the angle α defined by

$$\sin \alpha = \frac{2\rho}{\Upsilon(\sigma)\Delta(\sigma)}.$$

Also, by an elementary geometric argument,

$$\sin \gamma = \frac{2\rho}{\|v\|}$$

defines γ as an upper bound on the angle between the lines generated by v and $\tilde{v}$.

Thus we have

$$D(\tilde{p}, \tilde{\sigma}) = \|\tilde{v}\| \sin \tilde{\theta} \geq (\|v\| - 2\rho) \sin(\theta - \alpha - \gamma).$$

Using the addition formula for sine together with the facts that for $x, y \in [0, \frac{\pi}{2}]$, $(1 - x) \leq \cos x$; $2 \sin x \geq x$; and $\sin x + \sin y \geq \sin(x + y)$, we get

$$D(\tilde{p}, \tilde{\sigma}) \geq (\|v\| - 2\rho)\left[\left(1 - 2\left(\frac{2\rho}{\Upsilon(\sigma)\Delta(\sigma)} + \frac{2\rho}{\|v\|}\right)\right)\frac{D(p, \sigma)}{\|v\|} - \left(\frac{2\rho}{\Upsilon(\sigma)\Delta(\sigma)} + \frac{2\rho}{\|v\|}\right)\right].$$

7

For convenience, define $\mu = \frac{2\rho}{L(\sigma)} \geq \frac{2\rho}{\|v\|} \geq \frac{2\rho}{\Delta(\sigma)}$. Then

$$
\begin{aligned}
D(\tilde{p}, \tilde{\sigma}) &\geq \|v\| \, (1-\mu) \left[\left(1 - 2 \left(1 + \frac{1}{\Upsilon(\sigma)} \right) \mu \right) \frac{D(p, \sigma)}{\|v\|} - \left(1 + \frac{1}{\Upsilon(\sigma)} \right) \mu \right] \\
&\geq (1-\mu) \left[\left(1 - \frac{4\mu}{\Upsilon(\sigma)} \right) D(p, \sigma) - \frac{2\mu \, \|v\|}{\Upsilon(\sigma)} \right] \\
&\geq (1-\mu) \left[\left(1 - \frac{4\mu}{\Upsilon(\sigma)} \right) D(p, \sigma) - \frac{2\mu \, \|v\|}{\Upsilon(\sigma)^2 \Delta(\sigma)} D(p, \sigma) \right] \\
&\geq (1-\mu) \left(1 - \frac{4\mu}{\Upsilon(\sigma)} - \frac{2\mu}{\Upsilon(\sigma)^2} \right) D(p, \sigma) \\
&\geq \left(1 - \frac{7\mu}{\Upsilon(\sigma)^2} \right) D(p, \sigma) \\
&\geq K D(p, \sigma) \qquad \text{when } \mu \leq \frac{(1-K)\Upsilon(\sigma)^2}{7}.
\end{aligned}
$$

The condition on μ is satisfied when ρ satisfies Inequality (3).

The bound on $\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma})$ follows immediately from the bounds on the $D(\tilde{p}, \tilde{\sigma})$, and the bound on $\Upsilon(\tilde{\sigma})$ itself follows from the observation that

$$
\frac{\Delta(\sigma)}{\Delta(\tilde{\sigma})} \geq \frac{\Delta(\sigma)}{\Delta(\sigma) + 2\rho} \geq \left(1 - \frac{2\rho}{\Delta(\sigma)} \right) \geq \left(1 - \frac{\Upsilon(\sigma)^2}{7} \right) \geq \frac{6}{7},
$$

when ρ satisfies Inequality (3). $\qquad \square$

2.3 Delaunay and witness complexes

We recall the definitions of the witness complex and the restricted Delaunay triangulation. Let $L \subset \mathbb{R}^d$ be a finite set and let $X \subset \mathbb{R}^d$ be an arbitrary set.

Definition 2.9 A point $x \in X$ is a *witness* for $\sigma \subset L$ if

$$
\|p - x\| \leq \|q - x\| \quad \forall p \in \sigma \text{ and } q \in L \setminus \sigma.
$$

A *Delaunay centre* for σ is a point x that satisfies

$$
\|p - x\| \leq \|q - x\| \quad \forall p \in \sigma \text{ and } q \in L. \tag{4}
$$

If a simplex σ has a witness, we say that it is *witnessed*.

Note that a Delaunay centre is also a witness; the relaxed qualification on q in Equation (4) simply serves to demand that $\|p - x\| = \|p' - x\|$ for all $p, p' \in \sigma$. The formulation in terms of an inequality is convenient in the subsequent developments. In de Silva's original terminology [13], a Delaunay centre was referred to as a strong witness. In the context of the additional terminology we will subsequently introduce, we have found it convenient increase the difference between these names.

These definitions give rise to the simplicial complexes that are the focus of interest here:

Definition 2.10 The *witness complex* of $L \subset \mathbb{R}^d$ with respect to X is the abstract simplicial complex, $\mathrm{Wit}(L, X)$, on L defined by

$$\sigma \in \mathrm{Wit}(L, X) \iff \text{ every subsimplex } \tau \leq \sigma \text{ has a witness in } X.$$

The *Delaunay complex of L restricted to X* is the abstract simplicial complex, $\mathrm{Del}(L, X)$, on L defined by

$$\sigma \in \mathrm{Del}(L, X) \iff \text{ every subsimplex } \tau \leq \sigma \text{ has a Delaunay centre in } X.$$

Delaunay [16] showed that if $X = \mathbb{R}^d$ and L is in *general position*, i.e. $\mathrm{Del}(L, \mathbb{R}^d)$ has no $(d+1)$-simplices, then $\mathrm{Del}(L, \mathbb{R}^d)$ is a triangulation of $\mathbb{R}^d$. In particular, the natural mapping of $\mathrm{Del}(L, \mathbb{R}^d)$ into $\mathbb{R}^d$ is an embedding. In our context this implies that if L is in general position in $\mathbb{R}^d$, then for any $X \subseteq \mathbb{R}^d$, $\mathrm{Del}(L, X)$ has a natural geometric realization defined by the inclusion of L in $\mathbb{R}^d$. de Silva [13] showed that if $X = \mathbb{R}^d$, then $\mathrm{Del}(L, X) = \mathrm{Wit}(L, X)$; the result does not require a general position assumption, but it implies that $\mathrm{Wit}(L, X)$ is also naturally realized as an embedded simplicial complex when the general position assumption on L is met.

We will also use the relaxed notion of the witness complex, which was also introduced by de Silva [13].

Definition 2.11 Assume $\rho \geq 0$. A point $x \in X$ is a *ρ-witness* for $\sigma \subset L$ if

$$\|p - x\| \leq \|q - x\| + \rho \quad \forall p \in \sigma, \text{ and } q \in L \setminus \sigma.$$

A *ρ-Delaunay centre* for σ is a point x that satisfies

$$\|p - x\| \leq \|q - x\| + \rho \quad \forall p \in \sigma, \text{ and } q \in L.$$

Using ρ-witnesses we obtain a superset of the witness complex

Definition 2.12 The *ρ-witness complex*, $\mathrm{Wit}^\rho(L, X)$, is the abstract simplicial complex on L defined by

$$\sigma \in \mathrm{Wit}^\rho(L, X) \iff \text{ every subsimplex } \tau \leq \sigma \text{ has a } \rho\text{-witness in } X.$$

The *ρ-Delaunay complex of L restricted to X* is the abstract simplicial complex, $\mathrm{Del}^\rho(L, X)$, on L defined by

$$\sigma \in \mathrm{Del}^\rho(L, X) \iff \text{ every subsimplex } \tau \leq \sigma \text{ has a } \rho\text{-Delaunay centre in } X.$$

de Silva [13] also showed that $\mathrm{Wit}^\rho(L, \mathbb{R}^d) = \mathrm{Del}^\rho(L, \mathbb{R}^d)$. Our motivation for introducing $\mathrm{Wit}^\rho(L, W)$ is that we want a discrete set of witnesses W representing an embedded manifold $\mathbb{M}$. We would like to be able to relax the Delaunay condition without changing the Delaunay complex. For this we introduce the protection idea:

Definition 2.13 A point $x \in X$ is a *δ-protected witness* for $\sigma \subset L$ if

$$\|p - x\| < \|q - x\| - \delta \quad \forall p \in \sigma, \text{ and } q \in L \setminus \sigma.$$

If, in addition, x is a ρ-Delaunay centre for σ, then we say x is a *(δ, ρ)-Delaunay centre* for σ.

Protection is the opposite of relaxation, but a (δ, ρ)-Delaunay centre combines both notions.

A (δ, ρ)-Delaunay centre for σ is the centre of a closed ball $\overline{B}$ which contains all the vertices of σ and these lie within a distance ρ from $\partial \overline{B}$, and furthermore there are no points in $L \setminus \sigma$ within a distance of δ from $\partial \overline{B}$. Observe in particular that a (δ, ρ)-Delaunay centre for σ is a witness for σ, but it is not a Delaunay centre for σ. We will require $\delta > \rho$.

3   Equating the restricted Delaunay and witness complexes

Given a smooth compact k-manifold, $\mathbb{M} \subset \mathbb{R}^d$, we establish conditions on the finite set $W \subset \mathbb{R}^d$ of *witnesses*, and the set $L \subset W$ of *landmarks*, which will guarantee the equivalence of $\mathrm{Del}(L, \mathbb{M})$ and $\mathrm{Wit}(L, W)$. These conditions are specified in terms of $\mathrm{rch}(\mathbb{M})$ and three additional parameters by which the sampling radius ϵ of W with respect to $\mathbb{M}$, and the sampling radius λ of L with respect to W are controlled. These parameters are δ, which specifies the protection of the simplices, ρ, which is a relaxation parameter, closely tied to ϵ, and Υ_0, which is a constant that places a lower bound on the thickness of the simplices.

The sampling density of L reflects the resolution of the final simplicial representation of $\mathbb{M}$. It is governed by λ which is constrained only by Υ_0 and $\mathrm{rch}(\mathbb{M})$. The sampling density of W is constrained by the choice of λ.

Hypotheses 3.1 The sampling conditions demanded of L and W are as follows:

$\mathbf{H}_L$ $L \subset W$ is a λ-sparse λ-sample of W with

$$\lambda \leq \frac{\Upsilon_0 \mathrm{rch}(\mathbb{M})}{32},$$

$$\rho \leq \left(\frac{\Upsilon_0^2}{130}\right) \lambda$$

$$\delta \geq \left(\frac{644}{\Upsilon_0^2}\right) \rho,$$

and every simplex $\sigma \in \mathrm{Wit}^\rho(L, W)$ is Υ_0-thick and has a (δ, ρ)-Delaunay centre in W.

$\mathbf{H}_W$ $W \subset \mathbb{R}^d$ is an ϵ-sample of $\mathbb{M}$ with

$$\epsilon \leq \min\{\frac{\rho}{2}, \frac{2\lambda^2}{\mathrm{rch}(\mathbb{M})}\}.$$

Looking past the parameters and equations, the essential demand of $\mathbf{H}_L$ is that every $\sigma \in \mathrm{Wit}^\rho(L, W)$ is thick and has a (δ, ρ)-Delaunay centre $x \in \mathbb{R}^d$. As mentioned above, de Silva [13] has demonstrated that every simplex in $\mathrm{Wit}^\rho(L, W)$ already has a ρ-Delaunay centre. Also a thickness or equivalent quality constraint on the simplices is standard and is known to be a requirement [12, 22] for the restricted Delaunay triangulation to be an accurate representation of $\mathbb{M}$. Thus the principal novelty in our criteria is the protection demand, and it yields our main structural result:

Theorem 3.2 If $\mathbf{H}_W$ and $\mathbf{H}_L$ are satisfied, then

$$\mathrm{Wit}(L, W) = \mathrm{Del}(L, \mathbb{M}).$$

The protection requirement is not an unreasonable demand: As we demonstrate below, our criteria imply that each Delaunay simplex has a Delaunay centre that has some amount of protection. It is not difficult to verify that if there is a simplex in $\mathrm{Del}(L, \mathbb{R}^d)$ that does not have a protected Delaunay centre, then there is a $(d+1)$-simplex in $\mathrm{Del}(L, \mathbb{R}^d)$. In other words, a Delaunay simplex without any protected Delaunay centre only occurs in a point set that is not in "general position". Although de Silva [13] makes no such demand, it is customary, when dealing with Delaunay triangulations,

to assume that the vertex set is in general position. Thus this standard assumption is already demanding that there is some $\delta > 0$ such that the Delaunay simplices each have a δ-protected centre. The difference is that the general position assumption allows δ to be arbitrarily small, but we demand that δ be bounded below by ρ, which is in turn bounded below by ϵ. However, we may let ϵ be arbitrarily small.

Thus if ϵ is small enough we should expect that if we generate a λ-sparse λ-sample L of W, then (once non-thick simplices are exuded) it will probably meet our criteria. In this sense our sampling criteria are not unreasonable, yet when they are met we guarantee that the witness complex coincides with the restricted Delaunay complex.

3.1 Immediate consequences of protection: $\mathbf{Del}(L, \mathbb{M}) \subseteq \mathbf{Wit}(L, W)$

Any point sufficiently close to a Delaunay centre is a ρ-Delaunay centre.

Lemma 3.3 If $c \in \mathbb{R}^d$ is a Delaunay centre for σ, then for $r > 0$, any $x \in \overline{B}_{\mathbb{R}^d}(c; r)$ is a $2r$-Delaunay centre for σ.

Proof   By definition, for all $p \in \sigma$ and $q \in L$, we have

$$\|p - c\| \leq \|q - c\|.$$

For $x \in \overline{B}_{\mathbb{R}^d}(c; r)$, the triangle inequality gives

$$\|p - x\| - r \leq \|p - c\|,$$

and

$$\|q - c\| \leq \|q - x\| + r,$$

and the result follows from the definition of a $2r$-Delaunay centre. $\square$

As a direct consequence of Lemma 3.3 we get

Proposition 3.4 If $\mathbf{H}_W$ and $\mathbf{H}_L$ are satisfied, then

$$\mathrm{Del}(L, \mathbb{M}) \subseteq \mathrm{Wit}(L, W).$$

Proof   Since a (δ, ρ)-Delaunay centre is a witness, $\mathbf{H}_L$ gives us

$$\mathrm{Wit}^\rho(L, W) \subseteq \mathrm{Wit}(L, W).$$

The result follows because Lemma 3.3 guarantees $\mathrm{Del}(L, \mathbb{M}) \subseteq \mathrm{Wit}^\rho(L, W)$ when $\mathbf{H}_W$ is satisfied.
$\square$

3.2 (δ, ρ)-Delaunay centres and thick simplices: $\mathbf{Wit}(L, W) \subseteq \mathbf{Del}(L, \mathbb{M})$

The importance of protection stems from the fact that it confers a kind of stability on the witness complex. Centred at any protected witness for σ is a small ball within which every point is a witness for σ:

Lemma 3.5 If $z \in \mathbb{R}^d$ is a δ-protected witness for σ, and $r \leq \delta/2$, then any $x \in \overline{B}_{\mathbb{R}^d}(z;r)$ is a $(\delta - 2r)$-protected witness for σ.

Proof   By definition of δ-protection,

$$\|p - z\| < \|q - z\| - \delta \qquad \forall p \in \sigma, \quad q \in L \setminus \sigma. \tag{5}$$

By the triangle inequality,

$$\|p - x\| - r \leq \|p - z\|,$$

and

$$\|q - z\| \leq \|q - x\| + r.$$

Substituting these inequalities into Equation (5), we obtain,

$$\|p - x\| < \|q - x\| - (\delta - 2r) \qquad \forall p \in \sigma, \quad q \in L \setminus \sigma,$$

yielding the protection assertion on x. $\square$

Thus if σ has a sufficiently protected witness $z \in \mathbb{M}$, and W is sufficiently dense, then we are guaranteed that σ will be witnessed by a point in W.

In fact, if the δ-protected witness z is a Delaunay centre, then Lemma 3.5 combined with Lemma 3.3 shows that any $x \in \overline{B}_{\mathbb{R}^d}(z;r)$ is a $((\delta - 2r), 2r)$-Delaunay centre for σ. In the following we will demonstrate a converse observation: A (δ, ρ)-Delaunay centre has a protected (and exact) Delaunay centre nearby. As a consequence, we get

Proposition 3.6 If $\mathbf{H}_W$ and $\mathbf{H}_L$ are satisfied, then

$$\mathrm{Wit}(L, W) \subseteq \mathrm{Del}(L, \mathbb{M}).$$

For the demonstration we require a technical lemma concerning circumscribing balls of simplices whose vertices are subjected to a small perturbation.

3.2.1 Perturbations and circumballs

Lemma 3.7 Suppose $B = B_{\mathbb{R}^d}(c;r)$, with $r < \lambda$, is a circumscribing ball for a j-simplex $\sigma = [p_0, \ldots, p_j]$. Suppose also that $\tilde{\sigma} = [\tilde{p}_0, \ldots, \tilde{p}_j]$ is such that $\tilde{p}_0 = p_0$ and $\|\tilde{p}_i - p_i\| \leq \rho$ for all $i \in [1, \ldots, j]$. If

$$\rho \leq \frac{\Upsilon(\sigma)^2 \Delta(\sigma)}{5},$$

then there is a circumscribing ball $\tilde{B} = B_{\mathbb{R}^d}(\tilde{c}; \tilde{r})$ for $\tilde{\sigma}$ with

$$|\tilde{r} - r| \leq \|\tilde{c} - c\| < \left(\frac{12\lambda}{\Upsilon(\sigma)^2 \Delta(\sigma)} \right) \rho.$$

Proof   Recall that $N(\sigma)$ is the $(d - j)$-dimensional flat of points equidistant to all the vertices of σ. So $c \in N(\sigma)$. We wish to find $\tilde{c}$, the projection of c into $N(\tilde{\sigma})$.

The proof becomes an exercise in linear algebra once we recognize that $[\tilde{c}, c]$ is parallel to $\mathrm{aff}(\tilde{\sigma})$, and therefore the vector $(\tilde{c} - c)$ may be represented as a linear combination of vectors representing

the edges of $\tilde{\sigma}$. This is Equation (6). We obtain the bound on $\|\tilde{c} - c\|$ by bounding the magnitude of the vector of these coefficients. We obtain the Expression (8) for this vector using the fact that $\tilde{c}$ itself is described by a system of linear equations (7) equating the squared distances from $\tilde{c}$ to the vertices of $\tilde{\sigma}$. The Expression (8) involves the matrix $(\tilde{P}^\mathsf{T}\tilde{P})^{-1}$, where the columns of $\tilde{P}$ are the edge vectors for $\tilde{\sigma}$. A bound on the norm of this matrix is found with the aid of Lemma 2.5, which introduces the dependence on the thickness of σ.

Choose a coordinate such that $p_0 = \tilde{p}_0$ is the origin. Define the matrices $\tilde{P} = [\tilde{p}_1 \cdots \tilde{p}_j]$, and $P = [p_1 \cdots p_j]$. Since $\tilde{c} - c$ is orthogonal to $N(\tilde{\sigma})$, it lies in the column space of $\tilde{P}$:

$$\tilde{c} - c = \tilde{P}b, \tag{6}$$

for some unknown vector of coefficients, b.

We also have that $\tilde{c} \in N(\tilde{\sigma})$ is defined by the j linear equations

$$\tilde{c}^2 = (\tilde{c} - \tilde{p}_i)^2,$$

for $i = 1$ to j, where we are using the shorthand $v^2 = v \cdot v$ for $v \in \mathbb{R}^d$. Expanding this out, we get

$$\tilde{p}_i \cdot \tilde{c} = \frac{1}{2}\tilde{p}_i^2,$$

which we can summarize in matrix notation as

$$\tilde{P}^\mathsf{T}\tilde{c} = \tilde{w}, \tag{7}$$

where $\tilde{w}$ is the $j \times 1$ vector whose coefficients are $\frac{1}{2}\tilde{p}_i^2$. Similarly, we have $P^\mathsf{T}c = w$, with the obvious definition of w.

Expanding Equation (7) with Equation (6), and subtracting $P^\mathsf{T}c = w$, we find

$$(\tilde{P} - P)^\mathsf{T}c + \tilde{P}^\mathsf{T}\tilde{P}b = \tilde{w} - w.$$

Defining $\check{P} = \tilde{P} - P$ and $\check{w} = \tilde{w} - w$, we obtain an expression for b:

$$b = (\tilde{P}^\mathsf{T}\tilde{P})^{-1}\left(\check{w} - \check{P}c\right). \tag{8}$$

We desire a bound on $\|b\|$. To start, we have $\|c\| < \lambda$, by hypothesis. Lemma 2.1 gives $\|\check{P}\| \le \rho\sqrt{j}$, and so,

$$\|\check{P}c\| < \rho\sqrt{j}\lambda$$

The i^{th} component of $\check{w}$ is $\frac{1}{2}(\tilde{p}_i^2 - p_i^2)$. Defining $\check{p}_i = \tilde{p}_i - p_i$, we get

$$\frac{1}{2}(\tilde{p}_i^2 - p_i^2) = p_i \cdot \check{p}_i + \frac{1}{2}\check{p}_i^2$$
$$\le \rho\|p_i\| + \frac{1}{2}\rho^2$$
$$\le \rho\Delta(\sigma) + \frac{1}{2}\rho^2$$
$$\le \rho\frac{3}{2}\Delta(\sigma) \qquad \text{when } \rho \le \Delta(\sigma).$$

Since there are j components in $\breve{w}$, we have

$$\|\breve{w}\| \le \rho \frac{3}{2}\sqrt{j}\Delta(\sigma).$$

Finally, to obtain a bound on

$$(\tilde{P}^\mathsf{T}\tilde{P})^{-1} = (P^\mathsf{T}P + (P^\mathsf{T}\breve{P} + \breve{P}^\mathsf{T}P + \breve{P}^\mathsf{T}\breve{P}))^{-1},$$

recall that for a $j \times j$ matrix A, $\|A^{-1}\| = s_j(A)^{-1}$. We will use the fact that if A and B are $j \times j$ matrices, then

$$s_j(A + B) \ge s_j(A) - s_1(B).$$

Indeed, for any unit vector v, $\|(A + B)v\| \ge \|Av\| - \|Bv\| \ge s_j(A) - s_1(B)$. Thus, since

$$\left\|P^\mathsf{T}\breve{P}\right\| = \left\|\breve{P}^\mathsf{T}P\right\| \le \left\|\breve{P}\right\|\|P\| \le \rho j \Delta(\sigma),$$

and

$$\left\|\breve{P}^\mathsf{T}\breve{P}\right\| \le \rho^2 j,$$

we get

$$\left\|(\tilde{P}^\mathsf{T}P)^{-1}\right\| \le (s_j(P)^2 - (\rho 2j\Delta(\sigma) + \rho^2 j))^{-1}.$$

$$\le \left(s_j(P)^2 - \rho\frac{5}{2}j\Delta(\sigma)\right)^{-1} \qquad \text{when } \rho \le \frac{\Delta(\sigma)}{2}$$

$$= s_j(P)^{-2}\left(1 - \rho\left(\frac{5j\Delta(\sigma)}{2s_j(P)^2}\right)\right)^{-1}$$

$$\le 2s_j(P)^{-2} \qquad \text{when } \rho \le \frac{s_j(P)^2}{5j\Delta(\sigma)}.$$

Putting it all together we have

$$\|b\| \le \left\|(\tilde{P}^\mathsf{T}P)^{-1}\right\|(\|\breve{w}\| + \left\|\breve{P}c\right\|)$$

$$< 2s_j(P)^{-2}\left(\frac{3}{2}\sqrt{j}\Delta(\sigma) + \sqrt{j}\lambda\right)\rho$$

$$= \rho\left(\frac{\sqrt{j}(3\Delta(\sigma) + 2\lambda)}{s_j(P)^2}\right).$$

We now can bound the distance between the centre c of the circumscribing ball B for σ, and the nearest centre for a circumscribing ball for $\tilde{\sigma}$. Since $\left\|\tilde{P}\right\| = \|P + \breve{P}\| \le \sqrt{j}(\Delta(\sigma) + \rho)$, we get

$$\|\tilde{c} - c\| \le \left\|\tilde{P}\right\|\|b\|$$

$$< \rho\sqrt{j}(\Delta(\sigma) + \rho)\left(\frac{\sqrt{j}(3\Delta(\sigma) + 2\lambda)}{s_j(P)^2}\right)$$

$$= \rho\left(\frac{j(3\Delta(\sigma)^2 + 2\Delta(\sigma)\lambda + 3\Delta(\sigma)\rho + 2\lambda\rho)}{s_j(P)^2}\right)$$

$$< \rho\left(\frac{j12\Delta(\sigma)\lambda}{s_j(P)^2}\right) \qquad \text{using } \Delta(\sigma) < 2\lambda \text{ and } \rho \le \frac{\Delta(\sigma)}{2}.$$

Recalling that $s_j(P)^2 \geq j\Upsilon(\sigma)^2\Delta(\sigma)^2$ from Lemma 2.5, and since

$$\frac{\Upsilon(\sigma)^2\Delta(\sigma)}{5} \leq \min\left\{\frac{s_j(P)^2}{5j\Delta(\sigma)}, \frac{\Delta(\sigma)}{2}\right\},$$

we get

$$\|\tilde{c} - c\| < \left(\frac{12\lambda}{\Upsilon(\sigma)^2\Delta(\sigma)}\right)\rho$$

$$\text{when} \quad \rho \leq \frac{\Upsilon(\sigma)^2\Delta(\sigma)}{5}.$$

Since $p_0 = \tilde{p}_0$, we obtain $|\tilde{r} - r| \leq \|\tilde{c} - c\|$ by the triangle inequality. □

3.2.2 Stability from protection

Lemma 3.7 provides the main calculation needed for the following observation:

Lemma 3.8 Suppose that $x \in \mathbb{R}^d$ is a (δ, ρ)-Delaunay centre for a Υ_0-thick σ. Suppose also that $L(\sigma) \geq \lambda$, and that there is a vertex $p \in \sigma$ with $d_{\mathbb{R}^d}(x, p) < \lambda$. If

$$\rho \leq \frac{\Upsilon_0^2}{27}\lambda,$$

and

$$\delta \geq \frac{128}{\Upsilon_0^2}\rho,$$

then σ has a $(\delta - (\frac{128}{\Upsilon_0^2})\rho)$-protected Delaunay centre $\tilde{c} \in B_{\mathbb{R}^d}(x; (\frac{64}{\Upsilon_0^2})\rho)$.

Proof Without loss of generality, let p_0 be a vertex of σ that minimizes the distance to x, and let $B = B_{\mathbb{R}^d}(x; r)$, where $r = d_{\mathbb{R}^d}(x, p_0) < \lambda$. Consider the simplex $\tilde{\sigma} = [\tilde{p}_0, \ldots, \tilde{p}_j]$, where $\tilde{p}_i$ is the projection of p_i into ∂B. By construction $\tilde{p}_0 = p_0$, and $\|\tilde{p}_i - p_i\| \leq \rho$ since x is a ρ-Delaunay centre.

We wish to apply Lemma 3.7, except that the roles of σ and $\tilde{\sigma}$ are reversed. Employing Lemma 2.8, we demand
$$\rho \leq \frac{6K^2\Upsilon(\sigma)^2\lambda}{35} \leq \frac{\Upsilon(\tilde{\sigma})^2\Delta(\tilde{\sigma})}{5},$$

which is satisfied, as is Inequality (3) of Lemma 2.8, by $\rho \leq \frac{\Upsilon_0^2}{27}\lambda$, when we choose K to be the larger of the roots of the equation $\frac{6}{35}K^2 = \frac{1}{14}(1 - K)$. In this case, Lemma 3.7 guarantees that σ has a circumscribing ball with centre $\tilde{c}$ such that

$$\|\tilde{c} - x\| < \frac{12\lambda}{\Upsilon(\tilde{\sigma})^2\Delta(\tilde{\sigma})}\rho \leq \frac{12}{\frac{6}{7}K^2\Upsilon_0^2}\rho \leq \frac{64}{\Upsilon_0^2}\rho,$$

and Lemma 3.5 confirms that it is in fact a $(\delta - (\frac{128}{\Upsilon_0^2})\rho)$-protected Delaunay centre for σ. □

15

Lemma 3.8 ensures that a sufficiently thick simplex with a (δ, ρ)-Delaunay centre w is a Delaunay simplex with a protected Delaunay centre near w. The following lemma indicates that protection also confers stability on the restricted Delaunay triangulation. The main observation is that if a thick simplex σ has a sufficiently protected Delaunay centre $\tilde{c}$ near the manifold, then σ must belong to the restricted Delaunay triangulation. The protection assumption provides a way to circumvent the notorious fragility of the restricted Delaunay triangulation: we no longer need to know the manifold exactly in order to ensure that $\sigma \in \mathrm{Del}(L, \mathbb{M})$. The heart of the proof of this lemma lies in the observation that $N(\sigma)$ intersects $\mathbb{M}$ near $\tilde{c}$. The demonstration of this fact is relegated to Appendix A.

Lemma 3.9 Let $\sigma = [p_0, \ldots, p_j]$ be a Υ_0-thick j-simplex with $j \leq k = \dim \mathbb{M}$, and $\Delta(\sigma) \geq K_1\lambda$. Suppose that $\tilde{c} \in \mathbb{R}^d$ is a 8μ-protected Delaunay centre for σ, and that $\tilde{c}$ has a unique closest point $\hat{c} \in \mathbb{M}$ with $\|\tilde{c} - \hat{c}\| = \mu \leq \frac{\mathrm{rch}(\mathbb{M})}{20}$. Suppose also that each vertex p of σ has a unique closest point $\hat{p} \in \mathbb{M}$ such that $\|\hat{p} - \hat{c}\| < K_0\lambda \leq \mathrm{rch}(\mathbb{M})$, and $\|p - \hat{p}\| < \epsilon \leq \frac{C_0 K_0^2 \lambda^2}{2\mathrm{rch}(\mathbb{M})}$.

If
$$\lambda \leq \frac{K_1 \Upsilon_0 \mathrm{rch}(\mathbb{M})}{4(1 + C_0)K_0^2}, \tag{9}$$

then σ has a Delaunay centre $c \in \mathbb{M}$ with $\|\tilde{c} - c\| \leq 4\mu$.

Proof   Let $\alpha = \angle(\mathrm{aff}(\sigma), T_{\hat{c}}\mathbb{M})$. Applying Lemma 2.7 with $r = K_0\lambda$, we find

$$\sin \alpha \leq \frac{(1 + C_0)K_0^2 \lambda}{K_1 \Upsilon_0 \mathrm{rch}(\mathbb{M})}.$$

By Lemma 2.3, $\alpha = \angle(N_{\hat{c}}\mathbb{M}, N(\sigma))$, and the assumption on λ implies $\sin \alpha \leq \frac{1}{4}$. Thus we may apply Lemma A.1 to conclude that there is a $c \in \mathbb{M}$ which is the centre of a circumscribing ball for σ, and $\|c - \tilde{c}\| \leq 4\mu$. Since $\tilde{c}$ is 8μ-protected, Lemma 3.5, ensures that c will be a Delaunay centre for σ. $\qquad\square$

The proof of Proposition 3.6 now becomes an exercise in verifying that the conditions of Lemma 3.9 are met.

Proof of Proposition 3.6   Let $w \in W$ be a (δ, ρ)-Delaunay centre for $\sigma \in \mathrm{Wit}(L, W)$. We will use Lemma 3.8 to show that σ has a protected Delaunay centre $\tilde{c}$ near w.

Since L is a λ-sample of W, there is a vertex p of σ such that $d_{\mathbb{R}^d}(w, p) < \lambda$. Since by Hypothesis $\mathbf{H}_L$, L is λ-sparse, and satisfies the other conditions of Lemma 3.8, there is a $(\delta - (\frac{128}{\Upsilon_0^2})\rho)$-protected Delaunay centre with $d_{\mathbb{R}^d}(\tilde{c}, w) < (\frac{64}{\Upsilon_0^2})\rho$.

Let $\hat{c} \in \mathbb{M}$ be the closest point to $\tilde{c}$. This point is unique since $\tilde{c}$ is at a distance less than $\mathrm{rch}(\mathbb{M})$ from $\mathbb{M}$. Indeed, since W is an ϵ-sample of $\mathbb{M}$, it follows that

$$\begin{aligned}
\mu = \|\tilde{c} - \hat{c}\| = d_{\mathbb{R}^d}(\tilde{c}, \mathbb{M}) &\leq d_{\mathbb{R}^d}(\tilde{c}, w) + d_{\mathbb{R}^d}(w, \mathbb{M}) \\
&< \frac{64}{\Upsilon_0^2}\rho + \epsilon \\
&\leq \left(\frac{64}{\Upsilon_0^2} + \frac{1}{2}\right)\rho,
\end{aligned} \tag{10}$$

16

which Hypothesis $\mathbf{H}_L$ ensures is less than the $\frac{\mathrm{rch}(\mathbb{M})}{20}$ required by Lemma 3.9.

As usual, let p be a vertex of σ and $\hat{p}$ its closest point on $\mathbb{M}$. In order to establish a bound on $\|\hat{p} - \hat{c}\|$, we note that $\|p - \hat{p}\| < \epsilon$, since $L \subset W$, and that $\|p - w\| \leq \lambda + \rho$ since w is a (δ, ρ)-Delaunay centre. Thus

$$\|\hat{p} - \hat{c}\| \leq \|\hat{p} - p\| + \|p - w\| + \|w - \tilde{c}\| + \|\tilde{c} - \hat{c}\|$$
$$< \epsilon + (\lambda + \rho) + \frac{64}{\Upsilon_0^2}\rho + \left(\frac{64}{\Upsilon_0^2} + \frac{1}{2}\right)\rho$$
$$\leq \lambda + \left(\frac{128}{\Upsilon_0^2} + 2\right)\rho$$
$$\leq 2\lambda \qquad \text{when } \rho \leq \left(\frac{128}{\Upsilon_0^2} + 2\right)^{-1}\lambda.$$

This bound on ρ with respect to λ is satisfied when

$$\rho \leq \frac{\Upsilon_0^2}{130}\lambda,$$

as imposed by Hypothesis $\mathbf{H}_L$.

Thus, once we have established that $\tilde{c}$ is 8μ protected, the Hypotheses $\mathbf{H}_W$ and $\mathbf{H}_L$, which demand

$$\epsilon \leq \frac{2\lambda^2}{\mathrm{rch}(\mathbb{M})},$$

and

$$\lambda \leq \frac{\Upsilon_0\mathrm{rch}(\mathbb{M})}{32},$$

ensure that we may apply Lemma 3.9 with $K_0 = 2$, $K_1 = 1$, and $C_0 = 1$, and thus $\sigma \in \mathrm{Del}(L, \mathbb{M})$.

Recalling Inequality (10), the condition on the protection is satisfied when

$$\left(\delta - \frac{128}{\Upsilon_0^2}\rho\right) \geq 8\left(\frac{64}{\Upsilon_0^2} + \frac{1}{2}\right)\rho,$$

or

$$\delta \geq \left(\frac{640}{\Upsilon_0^2} + 4\right)\rho,$$

which is satisfied when

$$\delta \geq \frac{644}{\Upsilon_0^2}\rho,$$

as required by Hypothesis $\mathbf{H}_L$. $\qquad\square$

A   Almost normal flats intersect $\mathbb{M}$

This appendix is devoted to proving the following technical lemma, which asserts that, for $j \leq k = \dim\mathbb{M}$, if a $(d - j)$-flat, N, passes through a point $\tilde{c}$ that is close to $\mathbb{M}$, and the normal space at the point on $\mathbb{M}$ closest to $\tilde{c}$ makes a small angle with N, then N must intersect $\mathbb{M}$ in that vicinity. The technical difficulty stems from the fact that the codimension may be greater than one.

Lemma A.1 Let $\tilde{c} \in \mathbb{R}^d$ be such that it has a unique closest point $\hat{c}$ on $\mathbb{M}$ and $\|\tilde{c} - \hat{c}\| \leq \mu \leq \frac{\mathrm{rch}(\mathbb{M})}{20}$. Let $j \leq k = \dim \mathbb{M}$, and let N be a $(d-j)$-dimensional affine flat passing through $\tilde{c}$ such that $\angle(N_{\hat{c}}\mathbb{M}, N) \leq \alpha$ with $\sin \alpha \leq \frac{1}{4}$. Then there exists an $x \in N \cap \mathbb{M}$ such that $\|\tilde{c} - x\| \leq 4\mu$.

The idea of the proof is to consider the k-dimensional affine space $\tilde{T}_{\hat{c}}\mathbb{M}$ that passes through $\hat{c}$ and is orthogonal to N. We show that the orthogonal projection onto $\tilde{T}_{\hat{c}}\mathbb{M}$ induces, in some neighbourhood V of $\hat{c}$, a diffeomorphism between $\mathbb{M} \cap V$, and $\tilde{T}_{\hat{c}}\mathbb{M} \cap V$ (Lemma A.6). We use $T_{\hat{c}}\mathbb{M}$ as an intermediary in this calculation (Lemma A.5). Then, since N intersects $T_{\hat{c}}\mathbb{M}$ near $\hat{c}$ (Lemma A.4), we can argue that it must also intersect $\mathbb{M}$ because the established diffeomorphisms make a correspondence between points along segments parallel to N.

The final bounds are established in Lemma A.7, from which Lemma A.1 follows by a direct calculation, together with the following observations: If $\dim N = \dim N_{\hat{c}}\mathbb{M}$, then $\angle(N_{\hat{c}}\mathbb{M}, N) = \angle(N, N_{\hat{c}}\mathbb{M})$, and if $\dim N \geq \dim N_{\hat{c}}\mathbb{M}$, then there is an affine subspace $\tilde{N} \subset N$, such that $\dim \tilde{N} = \dim N_{\hat{c}}\mathbb{M}$, and $\angle(N_{\hat{c}}\mathbb{M}, \tilde{N}) = \angle(N_{\hat{c}}\mathbb{M}, N)$. Indeed, we may take $\tilde{N}$ to be the orthogonal projection of $N_{\hat{c}}\mathbb{M}$ into N.

We will use the following results [17, 18].

Lemma A.2   1. For any point $q \in \mathbb{M}$ such that $\|p - q\| = t\,\mathrm{rch}(\mathbb{M})$ for some $0 < t < 1$, $\sin \angle(pq, T_p\mathbb{M}) \leq t/2$.

2. Let q be a point in $T_p\mathbb{M}$ such that $\|p - q\| = t\,\mathrm{rch}(\mathbb{M})$ for some $0 < t \leq 1/4$. Let q' be the point on $\mathbb{M}$ closest to q. Then $\|q - q'\| \leq 2t\|p - q\|$.

The following lemma, which bounds the angle between nearby tangent spaces, is a particular case of a more general result [6, Lemma 5.5], which we demonstrate here in order to obtain an appropriate explicit constant.

Proposition A.3 (Tangent variation) Let p and q be points in $\mathbb{M}$ such that $\|p - q\| \leq \frac{\mathrm{rch}(\mathbb{M})}{10}$. Then,
$$\sin \angle(T_p\mathbb{M}, T_q\mathbb{M}) \leq \frac{5\,\|p - q\|}{\mathrm{rch}(\mathbb{M})}$$

Proof   Let $t = \frac{\|p - q\|}{\mathrm{rch}(p)}$. We will show that for any unit vector u in T_p there exists a unit vector v in T_q such that $\sin \angle(u, v) = 5t$.

For a unit vector u in $T_p\mathbb{M}$, let $p_u \in T_p\mathbb{M}$ be defined as
$$p_u = t\,\mathrm{rch}(\mathbb{M}) \cdot u$$

Let v denote the unit vector in $T_q\mathbb{M}$ which makes the smallest angle with the unit vector u.

Let p_u' denote the point closest to p_u on $\mathbb{M}$. Then, from Lemma A.2 (2), we have
$$\|q - p_u'\| \leq \|q - p\| + \|p - p_u\| + \|p_u - p_u'\| \leq 2t\,(1 + t)\,\mathrm{rch}(\mathbb{M}) \tag{11}$$

Using Lemma A.2 (1), we have
$$\mathrm{dist}(p, T_q\mathbb{M}) \leq \|p - q\| \sin \angle(pq, T_q\mathbb{M}) \leq t^2 \mathrm{rch}(\mathbb{M})/2 \tag{12}$$

Using Lemma A.2 (2), we have

$$\begin{aligned} \operatorname{dist}(p_u, T_q\mathbb{M}) &\leq \operatorname{dist}(p'_u, T_q\mathbb{M}) + \|p_u - p'_u\| \\ &\leq 2t^2(1+t)^2\operatorname{rch}(\mathbb{M}) + 2t^2\operatorname{rch}(\mathbb{M}) \end{aligned} \tag{13}$$

Let $\eta = \operatorname{dist}(p, T_q\mathbb{M}) + \operatorname{dist}(p_u, T_q\mathbb{M})$. From Eq. (12) and (13), we have

$$\begin{aligned} \eta &\leq 2.5t^2\operatorname{rch}(\mathbb{M}) + 2t^2(1+t)^2\operatorname{rch}(\mathbb{M}) \\ &= 4.5t^2\operatorname{rch}(\mathbb{M}) + 2t^2(t^2 + 2t)\operatorname{rch}(\mathbb{M}) \end{aligned}$$

Therefore,

$$\sin \angle(u, v) \leq \frac{\eta}{\|p - p_u\|} \leq 4.5t + 2t(t^2 + 2t) \leq 5t$$

$\square$

We now bound distances to the intersection of N and $T_{\hat{c}}\mathbb{M}$.

Lemma A.4 Let $\tilde{c}$, $\hat{c}$ be points in $\mathbb{R}^d$ such that the projection of $\tilde{c}$ onto $\mathbb{M}$ is $\hat{c}$ and $\|\tilde{c} - \hat{c}\| \leq \mu$. Let N be a $d - k$ dimensional affine flat passing through $\tilde{c}$ such that $\angle(N, N_{\hat{c}}\mathcal{M}) \leq \alpha$. For all $x \in N \cap T_{\hat{c}}\mathbb{M}$, we have

1. $\|\tilde{c} - x\| \leq \frac{\mu}{\cos \alpha}$

2. $\|\hat{c} - x\| \leq \left(1 + \frac{1}{\cos \alpha}\right)\mu$

Proof For a point $x \in N \cap T_{\hat{c}}\mathbb{M}$, let u_x denote the unit vector from $\tilde{c}$ to x, and let $v_x \in N_{\hat{c}}\mathbb{M}$ be the unit vector that makes the smallest angle with u_x. Let H denote the hyperplane passing through $\hat{c}$ and orthogonal to v_x. Since $\|\tilde{c} - \hat{c}\| \leq \mu$, $\operatorname{dist}(\tilde{c}, H) \leq \mu$. Therefore,

$$\|\tilde{c} - x\| \leq \frac{\operatorname{dist}(\tilde{c}, H)}{\cos \alpha}$$

and

$$\|\hat{c} - x\| \leq \|\hat{c} - \tilde{c}\| + \|\tilde{c} - x\| \leq \left(1 + \frac{1}{\cos \alpha}\right)\mu$$

$\square$

The following lemma is a direct consequence of the definition of the angle between two affine spaces.

Lemma A.5 Let p be a point in $\mathbb{M}$ and let $\widetilde{T}_p\mathbb{M}$ denote a k-dimensional flat passing through p with $\angle(T_p\mathbb{M}, \widetilde{T}_p\mathbb{M}) \leq \alpha < \frac{\pi}{2}$. Let f_p^α denote the orthogonal projection of $T_p\mathbb{M}$ onto $\widetilde{T}_p\mathbb{M}$, then :

1. The map f_p^α is bijective.

2. For $r > 0$, $f_p^\alpha(B_p(r)) \supseteq \widetilde{B}_p(r\cos\alpha)$ where $B_p(r) = B_{\mathbb{R}^d}(p; r) \cap T_p\mathbb{M}$ and $\widetilde{B}_p(r) = B_{\mathbb{R}^d}(p; r) \cap \widetilde{T}_p\mathbb{M}$.

19

Lemma A.6 Let p be a point in $\mathbb{M}$, and let $\widetilde{T}_p\mathbb{M}$ be a k-dimensional affine flat passing through p with $\angle(T_p\mathbb{M}, \widetilde{T}_p\mathbb{M}) \le \alpha$. There exists an $r(\alpha)$ satisfying :

$$\frac{6\,r(\alpha)}{\mathrm{rch}(\mathbb{M})} + \sin\alpha \; < \; 1 \quad \text{and} \quad r(\alpha) \le \frac{\mathrm{rch}(\mathbb{M})}{10}$$

such that the orthogonal projection map, g_p^α, of $B_\mathbb{M}(p, r(\alpha)) = B_{\mathbb{R}^d}(p; r(\alpha)) \cap \mathbb{M}$ into $\widetilde{T}_p\mathbb{M}$ satisfy the following conditions:

(1) g_p^α is a diffeomorphism.

(2) $g_p^\alpha(B_\mathbb{M}(p, r(\alpha))) \supseteq \widetilde{B}_p(r(\alpha)\cos\alpha_1)$ where $\sin\alpha_1 = \frac{r(\alpha)}{2\,\mathrm{rch}(\mathbb{M})} + \sin\alpha$.

(3) Let $x \in g_p^\alpha(B_\mathbb{M}(p, r(\alpha)))$, then $\|x - (g_p^\alpha)^{-1}(x)\| \le \|p - x\|\tan\alpha_1$

Proof   1. Let $\pi_{\widetilde{T}_p\mathbb{M}}$ denote the orthogonal projection of $\mathbb{R}^d$ onto $\widetilde{T}_p\mathbb{M}$. The derivative of this map, $D\pi_{\widetilde{T}_p\mathbb{M}}$, has a kernel of dimension $(d - k)$ that is parallel to the orthogonal complement of $\widetilde{T}_p\mathbb{M}$ in $\mathbb{R}^d$.

We will first show that Dg_p^α is nonsingular for all $x \in B_\mathbb{M}(p, r(\alpha))$. From Proposition A.3 and the fact that $\angle(T_p\mathbb{M}, \widetilde{T}_p\mathbb{M}) \le \alpha$, we have

$$\begin{aligned}
\sin\angle(\widetilde{T}_p\mathbb{M}, T_x\mathbb{M}) &\le& \sin\angle(T_x\mathbb{M}, T_p\mathbb{M}) + \sin\angle(T_p\mathbb{M}, \widetilde{T}_p\mathbb{M}) \\
&\le& \frac{5\,r(\alpha)}{\mathrm{rch}(\mathbb{M})} + \sin\alpha \; < \; 1
\end{aligned}$$

Since g_p^α is the restriction of $\pi_{\widetilde{T}_p\mathbb{M}}$ to $B_\mathbb{M}(p, r(\alpha))$, the above inequality implies that Dg_p^α is nonsingular. Therefore, g_p^α is a local diffeomorphism.

Let $x, y \in B_\mathbb{M}(p, r(\alpha))$. From Lemma A.2 (1) and Proposition A.3, we have

$$\begin{aligned}
\sin\angle([x, y], \widetilde{T}_p\mathbb{M}) &\le& \sin\angle([x, y], T_x\mathbb{M}) + \sin\angle(T_x\mathbb{M}, T_p\mathbb{M}) + \sin\angle(\widetilde{T}_p, T_p\mathbb{M}) \\
&\le& \frac{\|x - y\|}{2\mathrm{rch}(\mathbb{M})} + \frac{5\|p - x\|}{\mathrm{rch}(\mathbb{M})} + \sin\alpha \\
&\le& \frac{6\,r(\alpha)}{\mathrm{rch}(\mathbb{M})} + \sin\alpha \; < \; 1
\end{aligned}$$

This implies $g_p^\alpha(x) \ne g_p^\alpha(y)$.

Since g_p^α is nonsingular and injective on $B_\mathbb{M}(p, r(\alpha))$, it is a diffeomorphism onto its image.

2. Notice that, for $x \in B_\mathbb{M}(p, r(\alpha))$, the angle α_1 is a bound on the angle between $[p, x]$ and $\widetilde{T}_p\mathbb{M}$. The inclusion $g_p^\alpha(B_\mathbb{M}(p, r(\alpha))) \supseteq \widetilde{B}_p(r(\alpha)\cos\alpha_1)$ follows since $[x, g_P^\alpha(x)]$ is orthogonal to $\widetilde{T}_p\mathbb{M}$.

3. Follows similarly. $\square$

Lemma A.7 Let $\tilde{c}$, $\hat{c}$ be points in $\mathbb{R}^d$ such that the projection of $\tilde{c}$ onto $\mathbb{M}$ is $\hat{c}$ and $\|\tilde{c} - \hat{c}\| \le \mu$. Let N be a $d - k$ dimensional affine flat passing through $\tilde{c}$ such that $\angle(N, N_{\hat{c}}\mathcal{M}) \le \alpha$. If

$$\mu \le \frac{r(\alpha)\cos\alpha\cos\alpha_1}{1 + \cos\alpha}$$

20

Then there exists an $x \in N \cap \mathbb{M}$ such that

$$\|\tilde{c} - x\| \le \left(\frac{1}{\cos \alpha} + \left(1 + \frac{1}{\cos \alpha} \right) (\sin \alpha + \sin \alpha_1) \right) \mu \,.$$

Proof   Let $\tilde{T}_{\hat{c}}\mathbb{M}$ denote the orthogonal complement of N in $\mathbb{R}^d$ passing through $\hat{c}$. By Lemma 2.3, $\angle(T_{\hat{c}}\mathbb{M}, \tilde{T}_{\hat{c}}\mathbb{M}) = \angle(N, N_{\hat{c}}\mathcal{M})$.

Let $\hat{x} \in N \cap T_{\hat{c}}\mathbb{M}$ and $\tilde{x} = f_{\hat{c}}^{\alpha}(\hat{x})$. Then from Lemma A.4, we have

$$\|\tilde{x} - \hat{c}\| \le \|\hat{x} - \hat{c}\| \le \left(1 + \frac{1}{\cos \alpha} \right) \mu$$

and

$$\|\hat{x} - \tilde{x}\| \le \|\hat{x} - \hat{c}\| \sin \alpha \le \left(1 + \frac{1}{\cos \alpha} \right) \sin \alpha \, \mu \,.$$

Using the fact that $\mu \le \frac{r(\alpha) \cos \alpha \cos \alpha_1}{1 + \cos \alpha}$, we have $\|\tilde{x} - \hat{c}\| \le \left(1 + \frac{1}{\cos \alpha} \right) \mu \le r(\alpha) \cos \alpha_1$. Therefore, from Lemma A.4, there exists an $x \in B_{\mathbb{M}}(p, r(\alpha))$ such that $g_p^{\alpha}(x) = \tilde{x}$ and

$$\|\tilde{x} - x\| \le \|\tilde{x} - \hat{c}\| \tan \alpha_1 \le \left(1 + \frac{1}{\cos \alpha} \right) \tan \alpha_1 \, \mu.$$

Therefore

$$
\begin{aligned}
\|\tilde{c} - x\| &\le \|\tilde{c} - \hat{x}\| + \|\hat{x} - \tilde{x}\| + \|\tilde{x} - x\| \\
&\le \frac{\mu}{\cos \alpha} + \left(1 + \frac{1}{\cos \alpha} \right) (\sin \alpha + \tan \alpha_1) \, \mu \\
&= \left(\frac{1}{\cos \alpha} + \left(1 + \frac{1}{\cos \alpha} \right) (\sin \alpha + \tan \alpha_1) \right) \mu.
\end{aligned}
$$

Note that the line segment $[\tilde{c}, x] \in N$.   □

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Comp. Geom.*, 22(4):481–504, 1999.

[2] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60:125–135, 1998.

[3] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. of the ACM-SIAM SODA*, pages 213–222, 2000.

[4] D. Attali, H. Edelsbrunner, and Y. Mileyko. Weak witnesses for Delaunay triangulations of submanifolds. In *Proc. ACM Sympos. on Solid and Physical Modeling*, pages 143–150, 2007.

[5] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. In *ACM SoCG*, pages 324–333, 2010.

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

[7] J.-D. Boissonnat, L. Guibas, and S. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. *Discrete and Comp. Geom.*, 42(1):37–70, 2009.

[8] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, September 2005.

[9] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *Int. J. of Computer Vision*, 76(1):1–12, 2008.

[10] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction. In J.D. Boissonnat and M. Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*, pages 231–273. Springer, 2006.

[11] F. Chazal and S. Y. Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proc. 24th ACM Sympos. Comput. Geom.*, pages 232–241, 2008.

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

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

[14] V. de Silva and G. Carlsson. Topological estimation using witness complexes. In *Proc. Sympos. Point-Based Graphics*, pages 157–166, 2004.

[15] Vin de Silva. A weak characterisation of the delaunay triangulation. *Geometriae Dedicata*, 135(1):39–64, 2008.

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

[17] Herbert Federer. Curvature measures. *Trans. of the American Mathematical Society*, pages 418–491, 1959.

[18] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. In *Proc. of the ACM Symp. on Computational Geometry*, pages 329–337, 2003.

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

[20] L. J. Guibas and S. Y. Oudot. Reconstruction using witness complexes. *Discrete and Comp. Geom.*, 40:325–356, 2008.

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

[22] S. Y. Oudot. On the topology of the restricted Delaunay triangulation and witness complex in higher dimensions. Technical report, Stanford University, November 2006. LANL arXiv:0803.1296v1 [cs.CG].

[23] H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957.