Local criteria for triangulation of manifolds

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh,
and Mathijs Wintraecken

December 11, 2017

Contents

1 Introduction

A *triangulation* of a manifold M is a homeomorphism $H\colon |\mathcal{A}| \to M$, where $\mathcal{A}$ is a simplicial complex, and $|\mathcal{A}|$ is its underlying topological space. If such a homeomorphism exists, we say that $\mathcal{A}$ *triangulates* M.

The purpose of this paper is to present criteria which ensure that a candidate map H is indeed a homeomorphism. This work is motivated by earlier investigations into the problem of algorithmically constructing a complex that triangulates a given manifold [BG14, BDG17]. It complements and is closely related to recent work that investigates a particular natural example of such a map [DVW15].

In the motivating algorithmic setting, we are given a compact manifold M, and a manifold simplicial complex $\mathcal{A}$ is constructed by working locally in Euclidean coordinate charts. Here we lay out criteria, based on local properties that arise naturally in the construction of $\mathcal{A}$, that guarantee that H is a homeomorphism. These criteria, which are summarized in Theorem 17, are based on metric properties of H within "compatible" coordinate charts (Definition 4). The Euclidean metric in the local coordinate chart is central to the analysis, but no explicit metric on $|\mathcal{A}|$ or M is involved, and no explicit assumption of differentiability is required of H or M. However, our only examples that meet the required local criteria are in the differentiable setting. We do not know whether or not our criteria for homeomorphism implicitly imply that M admits a differentiable structure. They do imply that $\mathcal{A}$ is piecewise linear (admits an atlas with piecewise linear transition functions).

Relation to other work

The first demonstrations that differentiable manifolds *can* always be triangulated were constructive. Cairns [Cai34] used coordinate charts to cover the manifold with embeddings of patches of Euclidean triangulations. He showed that if the complexes were sufficiently refined the embedding maps could be perturbed such that they remain embeddings and the images of simplices coincide where patches overlap. A global homeomorphic complex is obtained by identifying simplices with the same image. The technique was later refined and extended [Whi40, Mun68], but it is not easily adapted to provide triangulation guarantees for complexes constructed by other algorithms.

An alternative approach was developed by Whitney [Whi57] using his result that a manifold can be embedded into Euclidean space. A complex is constructed via a process involving the intersection of the manifold with a fine Cartesian grid in the ambient space, and it is shown that the *closest-point projection map*, which takes a point in the complex to its unique closest point in the manifold, is a homeomorphism.

More recently, Edelsbrunner and Shah [ES97] defined the restricted Delaunay complex of a subset M of Euclidean space as the nerve of the Voronoi diagram on M when the ambient Euclidean metric is used. They showed that if M is a compact manifold, then the restricted Delaunay complex is homeomorphic to M when the Voronoi diagram satisfies the *closed ball property (cbp)*: Voronoi faces are closed topological balls of the appropriate dimension.

Using the cbp, Amenta and Bern [AB99] demonstrated a specific sampling density that is sufficient to guarantee that the restricted Delaunay complex triangulates the surface. However, since the complex constructed by their reconstruction algorithm cannot be guaranteed to be exactly the restricted Delaunay complex, a new argument establishing homeomorphism was developed, together with a simplified version of the algorithm [ACDL02].

Although it was established in the context of restricted Delaunay triangulations, the cbp is an elegant topological result that applies in more general contexts. For example, it has been used to establish conditions for intrinsic Delaunay triangulations of surfaces [DZM08], and Cheng et al. [CDR05] have indicated how it can be applied for establishing weighted restricted Delaunay triangulations of smooth submanifolds of arbitrary dimension in Euclidean space.

However, the cbp is only applicable to Delaunay-like complexes that can be realized as the nerve of some kind of Voronoi diagram on the manifold. Thus, for example, it does not necessarily apply to the tangential Delaunay complex constructed by Boissonnat and Ghosh [BG14]. Secondly, even when a Delaunay-like complex is being constructed, it can be algorithmically difficult to directly verify the properties of the associated Voronoi structure; sampling criteria and conditions on the complex under construction are desired, but may not be easy to obtain from the cbp. A third deficiency of the cbp is that, although it can establish that a complex $\mathcal{A}$ triangulates the manifold M, it does not provide a specific triangulation $H : |\mathcal{A}| \to M$. Such a correspondence allows us to compare *geometric* properties of $|\mathcal{A}|$ and M.

In [BG14] Whitney's argument was adapted to demonstrate that the closest-point projection maps the tangential Delaunay complex homeomorphically onto the original manifold. The argument is intricate, and like Whitney's, is tailored to the specific complex under consideration. In contrast, the result of [ACDL02], especially in the formulation presented by Boissonnat and Oudot [BO05], guarantees a triangulation of a surface by any complex which satisfies a few easily verifiable properties. However, the argument relies heavily on the the codimension being 1.

If a set of vertices is contained within a sufficiently small neighbourhood on a Riemannian manifold, barycentric coordinates can be defined. So there is a natural map from a Euclidean simplex of the appropriate dimension to the manifold, assuming a correspondence between the vertices of the simplex and those on the manifold. Thus when a complex $\mathcal{A}$ is appropriately defined with vertices on a Riemannian manifold M, there is a natural *barycentric coordinate map* $|\mathcal{A}| \to M$. In [DVW15], conditions are presented which guarantee that this map is a triangulation. Although this map is widely applicable, the intrinsic criteria can be inconvenient, for example, in the setting of Euclidean submanifold reconstruction, and furthermore the closest-point projection map may be preferred for triangulation in that setting.

The argument in [DVW15] is based on a general result [DVW15, Proposition 16] for establishing that a given map is a triangulation of a differentiable manifold. However, the criteria include a bound on the differential of the map, which is not easy to obtain. The analysis required to show that the closest-point projection map meets this bound is formidable, and this motivated the current alternate approach. We have relaxed this constraint to a much more easily verifiable bound on the metric distortion of the map when viewed within a coordinate chart.

The sampling criteria for submanifolds imposed by our main result applied to the closest-point projection map (Theorem 39) are the most relaxed that we are aware of. The result could be applied to improve the sampling guarantees of previous works, e.g., [CDR05, BG14].

In outline, the argument we develop here is the same as that of [ACDL02], but extends the result to apply to abstract manifolds of arbitrary dimension and submanifolds of $\mathbb{R}^N$ of arbitrary codimension. We first show that the map H is a local homeomorphism, and thus a covering map, provided certain criteria are met. Then injectivity is ensured when we can demonstrate that each component of M contains a point y such that $H^{-1}(y)$ is a single point. A core technical lemma from Whitney [Whi57, Appendix II Lemma 15a] still lies at the heart of our argument.

Outline

The demonstration is developed abstractly in Section 2, without explicitly defining the map H. We assume that it has already been established that the restriction of H to any Euclidean simplex in $|\mathcal{A}|$ is an embedding. This is a nontrivial step that needs to be resolved from the specific properties of a particular choice of H. The criteria for local homeomorphism apply in a common coordinate chart (for $|\mathcal{A}|$ and M), and relate the size and quality of the simplices with the metric distortion of H, viewed in the coordinate domain. The requirement that leads to injectivity is also expressed in a local coordinate chart; it essentially demands that the images of vertices behave in a natural and expected way.

In Section 4, Theorem 17 is applied to the specific case where $M \subset \mathbb{R}^N$, and H is the projection to the closest-point on M. This is a refinement of the argument presented in [BG14], also correcting an error.

In Appendix C the argument presented in [DVW15] is reviewed. This is very similar to the argument presented in Section 2, but it exploits a detailed analysis of the differential of the map. Although it depends on differentiability, and more analysis, an advantage of this approach is that the size of the simplices is restricted by a bound that is linear with respect to the simplex quality, instead of quadratic. The main reason for reviewing this argument is that as a result of an error in [DVW15], the criteria in the announced theorems there do not in fact ensure injectivity of the map. The appendix illuminates the problem and corrected statements of those results are presented in Section C.1.

2 The homeomorphism criteria

We assume that $\mathcal{A}$ and M are both compact manifolds of dimension m, without boundary, and we have a map $H\colon |\mathcal{A}| \to M$ that we wish to demonstrate is a homeomorphism. We first show that H is a covering map, i.e., every $y \in M$ admits an open neighbourhood U_y such that $H^{-1}(y)$ is a disjoint union of open sets each of which is mapped homeomorphically onto U_y by H. In our setting it is sufficient to establish that H is a local homeomorphism whose image touches all components of M: Brouwer's invariance of domain then ensures that H is surjective, and, since $|\mathcal{A}|$ is compact, has the covering map property.

Notation 1 (simplices and stars) In this section, a simplex $\boldsymbol{\sigma}$ will always be a *full simplex*; a closed Euclidean simplex, specified by a set of vertices together with all the points with nonnegative barycentric coordinates. The relative interior of $\boldsymbol{\sigma}$ is denoted by $\mathrm{relint}(\boldsymbol{\sigma})$. If $\boldsymbol{\sigma}$ is a simplex of $\mathcal{A}$, the subcomplex consisting of all simplices that have $\boldsymbol{\sigma}$ as a face, together with the faces of these simplices, is called the *star* of $\boldsymbol{\sigma}$, denoted by $\underline{\mathrm{St}}(\boldsymbol{\sigma})$; the star of a vertex p is $\underline{\mathrm{St}}(p)$.

We also sometimes use the *open star* of a simplex $\boldsymbol{\sigma} \in \mathcal{C}$. This is the union of the relative interiors of the simplices in $\mathcal{C}$ that have $\boldsymbol{\sigma}$ as a face: $\mathrm{st}(\boldsymbol{\sigma}) = \bigcup_{\boldsymbol{\tau} \supseteq \boldsymbol{\sigma}} \mathrm{relint}(\boldsymbol{\tau})$. It is an open set in $|\mathcal{C}|$, and it is open in $\mathbb{R}^m$ if $\boldsymbol{\sigma} \notin \partial\mathcal{C}$.

Notation 2 (topology) If $A \subseteq \mathbb{R}^n$, then the topological closure, interior, and boundary of A are denoted respectively by $\overline{A}$, $\mathrm{int}(A)$, and $\partial A = \overline{A} \setminus \mathrm{int}(A)$.

Notation 3 (linear algebra) The Euclidean norm of $v \in \mathbb{R}^m$ is denoted by $|v|$, and $\|A\| = \sup_{|x|=1} |Ax|$ denotes the operator norm of the linear operator A.

We will work in local coordinate charts. To any given map $G \colon |\mathcal{C}| \to \mathbb{R}^m$, where $\mathcal{C}$ is a simplicial complex, we associate a piecewise linear map $\widehat{G}$ that agrees with G on the vertices of $\mathcal{C}$, and maps $x \in \boldsymbol{\sigma} \in \mathcal{C}$ to the point with the same barycentric coordinates with respect to the images of the vertices. The map $\widehat{G}$ is called the *secant map* of G with respect to $\mathcal{C}$.

The following definition provides the framework within which we will work.

Definition 4 (compatible atlases) We say that $|\mathcal{A}|$ and M have *compatible atlases* for $H \colon |\mathcal{A}| \to M$ if:

(1) There is a coordinate atlas $\{(U_p, \phi_p)\}_{p \in \mathcal{P}}$ for M, where the index set $\mathcal{P}$ is the set of vertices of $\mathcal{A}$ and the sets U_p are connected.

(2) For each $p \in \mathcal{P}$, there is a subcomplex $\widetilde{\mathcal{C}}_p$ of $\mathcal{A}$ that contains $\underline{\mathrm{St}}(p)$ and $H(|\widetilde{\mathcal{C}}_p|) \subset U_p$. Also, the secant map of $\Phi_p := \phi_p \circ H|_{|\widetilde{\mathcal{C}}_p|}$ defines a piecewise linear embedding of $|\widetilde{\mathcal{C}}_p|$ into $\mathbb{R}^m$. We denote this secant map

by $\widehat{\Phi}_p$. By definition, $\widehat{\Phi}_p$ preserves the barycentric coordinates within each simplex, and thus the collection $\{(\widetilde{\mathcal{C}}_p, \widehat{\Phi}_p)\}_{p\in\mathcal{P}}$ provides a piecewise linear atlas for $\mathcal{A}$.

Remark 5 The requirement in Definition 4 that the local patches U_p be connected implies that on each connected component M' of M, there is a $p \in \mathcal{P}$ such that $H(p) \in M'$.

We let $\mathcal{C}_p = \widehat{\Phi}_p(\widetilde{\mathcal{C}}_p)$, and we will work within the compatible local coordinate charts. Thus we are studying a map of the form

$$F_p \colon |\mathcal{C}_p| \subset \mathbb{R}^m \to \mathbb{R}^m,$$

where $\mathcal{C}_p$ is an m-manifold complex with boundary embedded in $\mathbb{R}^m$, and

$$F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1}, \tag{1}$$

as shown in the following diagram:

$$\begin{tikzcd}
|\mathcal{A}| \arrow[rrr, "H"] & & & M \\
& |\widetilde{\mathcal{C}}_p| \arrow[ul, hook'] \arrow[r, "H|_{|\widetilde{\mathcal{C}}_p|}"] \arrow[d, "\widehat{\Phi}_p"'] & U_p \arrow[ur, hook] \arrow[d, "\phi_p"] & \\
& |\mathcal{C}_p| \arrow[r, "\underbrace{F_p}_{\phi_p\circ H\circ\widehat{\Phi}_p^{-1}}"'] \arrow[dl, hook] & \phi_p(U_p) \arrow[dr, hook] & \\
\mathbb{R}^m & & & \mathbb{R}^m
\end{tikzcd}$$

We will focus on the map F_p, which can be considered as a local realisation of $H|_{|\widetilde{\mathcal{C}}_p|}$. By construction, F_p leaves the vertices of $\mathcal{C}_p$ fixed: if $q \in \mathbb{R}^m$ is a vertex of $\mathcal{C}_p$, then $F_p(q) = q$, since $\widehat{\Phi}_p$ conincides with $\phi_p \circ H$ on vertices.

Remark 6 The setting described here conforms to the paradigm laid out in the tangential complex work [BG14]. There one locally constructs a (weighted) Delaunay triangulation in the tangent space T_pM. This gives us the local patch $\mathcal{C}_p$ (the star of p in T_pM). The vertices of the constructed complex actually lie on M, and we recognise $\mathcal{C}_p$ as the orthogonal projection $\widehat{\Phi}_p$ of the corresponding complex $\widetilde{\mathcal{C}}_p$ with vertices on M.

In this context, the homeomorphism H that we are trying to establish is given by the closest-point projection map onto M, restricted to $|\mathcal{A}|$. Now we are going to work in local coordinate charts, given at each vertex $p \in \mathcal{P}$ by the orthogonal projection ϕ_p of some neighbourhood of p, $U_p \subset M$ into T_pM. We recognise that $\widehat{\Phi}_p$ really does coincide with the secant map of $\phi_p \circ H|_{|\widetilde{\mathcal{C}}_p|}$.

2.1 Local homeomorphism

Our goal is to ensure that there is some open $V_p \subset |\mathcal{C}_p|$ such that $F_p|_{V_p}$ is an embedding and that the sets $\widetilde{V}_p = \widehat{\Phi}_p^{-1}(V_p)$ are sufficiently large to cover $|\mathcal{A}|$. This will imply that H is a local homeomorphism. Indeed, if V_p is embedded by F_p, then $\widetilde{V}_p$ is embedded by $H|_{\widetilde{V}_p} = \phi_p^{-1} \circ F_p \circ \widehat{\Phi}_p|_{\widetilde{V}_p}$, since ϕ_p and $\widehat{\Phi}_p$ are both embeddings. Since $|\mathcal{A}|$ is compact, Brouwer's invariance of domain, together with Remark 5, implies that H is surjective, and a covering map. It will only remain to ensure that H is also injective.

We assume that we are given (i.e., we can establish by context-dependent means) a couple of properties of F_p. We assume that it is *simplexwise positive*, which means that it is continuous, and its restriction to any m-simplex is an orientation preserving topological embedding. As discussed in Appendix A, we say that F_p preserves the orientation of an m-simplex $\boldsymbol{\sigma} \subset \mathbb{R}^m$ if $F_p|_{\boldsymbol{\sigma}}$ has degree 1 at any point in the image of the interior of $\boldsymbol{\sigma}$, i.e., $\deg(F_p, \mathrm{int}(\boldsymbol{\sigma}), y) = 1$ for $y \in F_p(\mathrm{int}(\boldsymbol{\sigma}))$. (The other assumption we make is that F_p, when restricted to an m-simplex does not distort distances very much, as discussed below.)

The local homeomorphism demonstration is based on Lemma 7 below, which is a particular case of an observation made by Whitney [Whi57, Appendix II Lemma 15a]. Whitney demonstrated a more general result from elementary first principles. The proof we give here is the same as Whitney's, except that we exploit elementary degree theory, as discussed in Appendix A, in order to avoid the differentiability assumptions Whitney made.

In the statement of the lemma, $\mathcal{C}^{m-1}$ refers to the $(m-1)$-*skeleton* of $\mathcal{C}$: the subcomplex consisting of simplices of dimension less than or equal to $m-1$. When $|\mathcal{C}|$ is a manifold with boundary, as in the lemma, then $\partial\mathcal{C}$ is the subcomplex containing all $(m-1)$-simplices that are the face of a single m-simplex, together with the faces of these simplices.

Lemma 7 (simplexwise positive embedding) *Assume $\mathcal{C}$ is an oriented m-manifold finite simplicial complex with boundary embedded in $\mathbb{R}^m$. Let*

$F\colon |\mathcal{C}| \to \mathbb{R}^m$ be simplexwise positive in $\mathcal{C}$. Suppose $V \subset |\mathcal{C}|$ is a connected open set such that $F(V) \cap F(|\partial\mathcal{C}|) = \emptyset$. If there is a $y \in F(V) \setminus F(|\mathcal{C}^{m-1}|)$ such that $F^{-1}(y)$ is a single point, then the restriction of F to V is a topological embedding.

Proof Notice that the topological boundary of $|\mathcal{C}| \subset \mathbb{R}^m$ is equal to the carrier of the boundary complex (see, e.g., [BDG13, Lemmas 3.6, 3.7]):

$$\partial|\mathcal{C}| = |\partial\mathcal{C}|.$$

Let $\Omega = |\mathcal{C}| \setminus |\partial\mathcal{C}|$. Since F is simplexwise positive, and $F(V)$ lies within a connected component of $\mathbb{R}^m \setminus F(\partial\Omega)$, the fact that $F^{-1}(y)$ is a single point implies that $F^{-1}(w)$ is a single point for any $w \in F(V) \setminus F(|C^{m-1}|)$ (Lemma 49). We need to show that F is also injective on $V \cap |\mathcal{C}^{m-1}|$.

We now show that $F(\mathrm{st}(\boldsymbol{\sigma}))$ is open for any $\boldsymbol{\sigma} \in \mathcal{C}^{m-1} \setminus \partial\mathcal{C}$, where $\mathrm{st}(\boldsymbol{\sigma})$ is the open star of $\boldsymbol{\sigma}$, defined in Notation 1. Suppose $x \in \mathrm{relint}(\boldsymbol{\tau})$ for some $\boldsymbol{\tau} \in \mathcal{C} \setminus \partial\mathcal{C}$. Since F is injective when restricted to any simplex, we can find a sufficiently small open (in $\mathbb{R}^m$) neighbourhood U of $F(x)$ such that $U \cap F(\partial\mathrm{st}(\boldsymbol{\tau})) = \emptyset$. Since the closure of the open star is equal to the carrier of our usual star:

$$\overline{\mathrm{st}(\boldsymbol{\tau})} = |\underline{\mathrm{St}}(\boldsymbol{\tau})|,$$

Lemma 49 implies that every point in $U \setminus F(|\underline{\mathrm{St}}(\boldsymbol{\tau})^{m-1}|)$ has the same number of points in its preimage. By the injectivity of F restricted to m-simplices, this number must be greater than zero for points near $F(x)$. It follows that $U \subseteq F(\mathrm{st}(\boldsymbol{\tau}))$.

If $x \in \mathrm{st}(\boldsymbol{\sigma})$, then $x \in \mathrm{relint}(\boldsymbol{\tau})$ for some $\boldsymbol{\tau} \in \mathcal{C} \setminus \partial\mathcal{C}$ that has $\boldsymbol{\sigma}$ as a face. Since $\mathrm{st}(\boldsymbol{\tau}) \subseteq \mathrm{st}(\boldsymbol{\sigma})$, we have $U \subseteq F(\mathrm{st}(\boldsymbol{\sigma}))$, and we conclude that $F(\mathrm{st}(\boldsymbol{\sigma}))$ is open.

Now, to see that F is injective on $|\mathcal{C}^{m-1}| \cap V$, suppose to the contrary that $w, z \in |\mathcal{C}^{m-1}| \cap V$ are two distinct points such that $F(w) = F(z)$. Since F is injective on each simplex, there are distinct simplices $\boldsymbol{\sigma}, \boldsymbol{\tau}$ such that $w \in \mathrm{relint}(\boldsymbol{\sigma})$ and $z \in \mathrm{relint}(\boldsymbol{\tau})$. So there is an open neighbourhood U of $F(w) = F(z)$ that is contained in $F(\mathrm{st}(\boldsymbol{\sigma})) \cap F(\mathrm{st}(\boldsymbol{\tau}))$.

We must have $\mathrm{st}(\boldsymbol{\sigma}) \cap \mathrm{st}(\boldsymbol{\tau}) = \emptyset$, because if $x \in \mathrm{st}(\boldsymbol{\sigma}) \cap \mathrm{st}(\boldsymbol{\tau})$, then $x \in \mathrm{relint}(\boldsymbol{\mu})$ for some $\boldsymbol{\mu}$ that has both $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$ as faces. But this means that both w and z belong to $\boldsymbol{\mu}$, contradicting the injectivity of $F|_{\boldsymbol{\mu}}$. It follows that points in the nonempty set $U \setminus |\mathcal{C}^{m-1}|$ have at least two points in their preimage, a contradiction. Thus $F|_V$ is injective, and therefore, by Brouwer's invariance of domain, it is an embedding. $\square$

Our strategy for employing Lemma 7 is to demand that the restriction of F_p to any m-simplex has low metric distortion, and use this fact to ensure that the image of $V_p \subset |\mathcal{C}_p|$ is not intersected by the image of the boundary of $|\mathcal{C}_p|$, i.e., we will establish that $F_p(V_p) \cap F_p(|\partial \mathcal{C}_p|) = \emptyset$. We need to also establish that there is a point y in $F_p(V_p) \setminus F_p(|\mathcal{C}_p^{m-1}|)$ such that $F^{-1}(y)$ is a single point. The metric distortion bound will help us here as well.

Definition 8 ($\boldsymbol{\xi}$-distortion map) A map $F \colon U \subset \mathbb{R}^m \to \mathbb{R}^m$ is a ξ-*distortion map* if for all $x, y \in U$ we have

$$\big||F(x) - F(y)| - |x - y|\big| \leq \xi \, |x - y| \, . \tag{2}$$

We are interested in ξ-distortion maps with small ξ. Equation (2) can be equivalently written

$$(1 - \xi)\, |x - y| \leq |F(x) - F(y)| \leq (1 + \xi)\, |x - y| \, ,$$

and it is clear that when $\xi < 1$, a ξ-distortion map is a bi-Lipschitz map. For our purposes the metric distortion constant ξ is more convenient than a bi-Lipschitz constant. It is easy to show that if F is a ξ-distortion map, with $\xi < 1$, then F is a homeomorphism onto its image, and F^{-1} is a $\frac{\xi}{1-\xi}$-distortion map (see Lemma 19(1)).

Assuming that $F_p|_{\boldsymbol{\sigma}}$ is a ξ-distortion map for each m-simplex $\boldsymbol{\sigma} \in \mathcal{C}_p$, we can bound how much it displaces points. Specifically, for any point $x \in |\mathcal{C}_p|$, we will bound $|x - F(x)|$. We exploit the fact that the $m+1$ vertices of $\boldsymbol{\sigma}$ remain fixed, and use *trilateration*, i.e., we use the estimates of the distances to the fixed vertices to estimate the location of $F(x)$. Here, the quality of the simplex comes into play.

Notation 9 (simplex quality) The *thickness* of $\boldsymbol{\sigma}$, denoted $t(\boldsymbol{\sigma})$ (or just t if there is no risk of confusion) is given by $\frac{a}{mL}$, where $a = a(\boldsymbol{\sigma})$ is the smallest altitude of $\boldsymbol{\sigma}$, and $L = L(\boldsymbol{\sigma})$ is the length of the longest edge. We set $t(\boldsymbol{\sigma}) = 1$ if $\boldsymbol{\sigma}$ has dimension 0.

Lemma 10 (trilateration) *Suppose $\boldsymbol{\sigma} \subset \mathbb{R}^m$ is an m-simplex, and $F \colon \boldsymbol{\sigma} \to \mathbb{R}^m$ is a ξ-distortion map that leaves the vertices of $\boldsymbol{\sigma}$ fixed. If $\xi \leq 1$, then for any $x \in \boldsymbol{\sigma}$,*

$$|x - F(x)| \leq \frac{3\xi L}{t},$$

where L is the length of the longest edge of $\boldsymbol{\sigma}$, and t is its thickness.

Proof Let $\{p_0, \dots, p_m\}$ be the vertices of $\boldsymbol{\sigma}$. For $x \in \boldsymbol{\sigma}$, let $\tilde{x} = F(x)$.

We choose p_0 as the origin, and observe that

$$p_i{}^{\mathsf{T}} x = \frac{1}{2}\left(|x|^2 + |p_i|^2 - |x - p_i|^2\right), \tag{3}$$

which we write in matrix form as $P^{\mathsf{T}} x = b$, where P is the $m \times m$ matrix whose i-th column is p_i, and b is the vector whose i-th component is given by the right-hand side of (3). Similarly, we have $P^{\mathsf{T}} \tilde{x} = \tilde{b}$ with the obvious definition of $\tilde{b}$. Then

$$\tilde{x} - x = (P^{\mathsf{T}})^{-1}(\tilde{b} - b).$$

Since $F(p_0) = p_0 = 0$, we have $||\tilde{x}| - |x|| \leq \xi\, |x|$, and so

$$||\tilde{x}|^2 - |x|^2| \leq \xi(2 + \xi)\, |x|^2 \leq 3\xi L^2.$$

Similarly, $||x - p_i|^2 - |\tilde{x} - p_i|^2| < 3\xi L^2$. Thus $|\tilde{b}_i - b_i| \leq 3\xi L^2$, and $|\tilde{b} - b| \leq 3\sqrt{m}\xi L^2$.

By [BDG13, Lemma 2.4] we have $\left\|(P^{\mathsf{T}})^{-1}\right\| \leq (\sqrt{m} t L)^{-1}$, and the stated bound follows. $\square$

2.1.1 Using $\underline{\mathrm{St}}(p)$ as $\widetilde{\mathcal{C}}_p$

For the local complex $\widetilde{\mathcal{C}}_p \subset \mathcal{A}$ introduced in Definition 4, we now make a specific choice: $\widetilde{\mathcal{C}}_p = \underline{\mathrm{St}}(p)$. This is the smallest complex allowed by the definition. For convenience, we define $\hat{p} = \widehat{\Phi}_p(p)$, so that $\mathcal{C}_p = \widehat{\Phi}_p(\widetilde{\mathcal{C}}_p) = \underline{\mathrm{St}}(\hat{p})$.

We assume that F_p is a ξ-distortion map on each simplex. The idea is to show that F_p is an embedding on a set V_p that includes all the points $x \in |\underline{\mathrm{St}}(\hat{p})|$ such that the barycentric coordinate of x associated with $\hat{p}$ in an m-simplex that contains x is at least $\frac{1}{m+1}$. (That is the homothetic copy of $|\underline{\mathrm{St}}(\hat{p})|$, "shrunk" by a factor of $1 - \frac{1}{m+1}$, using $\hat{p}$ as the origin.) To be more specific we define V_p to be the open set consisting of the points in $|\underline{\mathrm{St}}(\hat{p})|$ whose barycentric coordinate with respect to $\hat{p}$ is strictly larger than $\frac{1}{m+1} - \delta$, where $\delta > 0$ is arbitrarily small. Since the barycentric coordinates in each m-simplex sum to 1, and the piecewise linear maps $\widehat{\Phi}_p$ preserve barycentric coordinates, this ensures that the sets $\widehat{\Phi}_p^{-1}(V_p)$ cover $|\mathcal{A}|$.

In order to employ the simplexwise positive embedding lemma (Lemma 7), we need to establish that there is a point in $V_p \setminus |\mathcal{C}_p^{m-1}|$ that is not mapped to the image of any other point in $|\mathcal{C}_p|$. We choose the barycentre of a simplex for this purpose. We say that a simplicial complex is a *pure* m-dimensional simplicial complex if every simplex is the face of an m-simplex.

Lemma 11 (a point covered once) *Suppose $\mathcal{C}$ is a pure m-dimensional finite simplicial complex embedded in $\mathbb{R}^m$, and that for each $\boldsymbol{\sigma} \in \mathcal{C}$ we have $t(\boldsymbol{\sigma}) \geq t_0$. If $F\colon |\mathcal{C}| \to \mathbb{R}^m$ leaves the vertices of $\mathcal{C}$ fixed, and its restriction to any m-simplex in $\mathcal{C}$ is a ξ-distortion map with*

$$\xi \leq \frac{1}{6}\frac{m}{m+1}t_0^2, \tag{4}$$

then $F^{-1}(F(b)) = \{b\}$, where b is the barycentre of an m-simplex in $\mathcal{C}$ with the largest diameter.

Proof Let $\boldsymbol{\sigma} \in \mathcal{C}$ be an m-simplex with the largest diameter, i.e., $L(\boldsymbol{\sigma}) \geq L(\boldsymbol{\tau})$ for all $\boldsymbol{\tau} \in \mathcal{C}_p$. Let b be the barycentre of $\boldsymbol{\sigma}$. The distance from b to the boundary of $\boldsymbol{\sigma}$ is

$$\frac{a(\boldsymbol{\sigma})}{m+1} = \frac{mt(\boldsymbol{\sigma})L(\boldsymbol{\sigma})}{m+1}.$$

Using Lemma 10, we will be sure that $F^{-1}(F(b)) = \{b\}$ provided that for any $\boldsymbol{\tau} \in \mathcal{C}_p$ we have

$$\frac{3\xi L(\boldsymbol{\sigma})}{t(\boldsymbol{\sigma})} + \frac{3\xi L(\boldsymbol{\tau})}{t(\boldsymbol{\tau})} < \frac{mt(\boldsymbol{\sigma})L(\boldsymbol{\sigma})}{m+1},$$

which is satisfied when the constraint (4) is met. $\square$

Now we also need to ensure that $F_p(V_p) \cap F_p(|\partial\mathcal{C}_p|) = \emptyset$. Here we will explicitly use the assumption that $\mathcal{C}_p$ is $\underline{\mathrm{St}}(\hat{p})$. We say that $\underline{\mathrm{St}}(\hat{p})$ is a *full star* if its carrier is an m-manifold with boundary and $\hat{p}$ does not belong to $\partial\underline{\mathrm{St}}(\hat{p})$.

Lemma 12 (barycentric boundary separation) *Suppose $\underline{\mathrm{St}}(\hat{p})$ is a full m-dimensional star embedded in $\mathbb{R}^m$. Let $a_0 = \min_{\boldsymbol{\sigma}\in\underline{\mathrm{St}}(\hat{p})} a_{\hat{p}}(\boldsymbol{\sigma})$ be the smallest altitude of $\hat{p}$ in the m-simplices in $\underline{\mathrm{St}}(\hat{p})$. Suppose $x \in \boldsymbol{\sigma} \in \underline{\mathrm{St}}(\hat{p})$, where $\boldsymbol{\sigma}$ is an m-simplex, and $\lambda_{\boldsymbol{\sigma},\hat{p}}(x)$, the barycentric coordinate of x with respect to $\hat{p}$ in $\boldsymbol{\sigma}$, satisfies $\lambda_{\boldsymbol{\sigma},\hat{p}}(x) \geq \alpha$. Then $d_{\mathbb{R}^m}(x, |\partial\underline{\mathrm{St}}(\hat{p})|) \geq \alpha a_0$.*

If t_0 is a lower bound on the thicknesses of the simplices in $\underline{\mathrm{St}}(\hat{p})$, and s_0 is a lower bound on their diameters, then $d_{\mathbb{R}^m}(x, |\partial\underline{\mathrm{St}}(\hat{p})|) \geq \alpha m t_0 s_0$.

Proof Since we are interested in the distance to the boundary, consider a point $y \in |\partial\underline{\mathrm{St}}(\hat{p})|$ such that the segment $[x, y]$ lies in $|\underline{\mathrm{St}}(\hat{p})|$. The segment passes through a sequence of m-simplices, $\boldsymbol{\sigma}_0 = \boldsymbol{\sigma}, \boldsymbol{\sigma}_1, \ldots, \boldsymbol{\sigma}_n$, that partition it into subsegments $[x_i, y_i] \subset \boldsymbol{\sigma}_i$ with $x_0 = x,\ y_n = y$ and $x_i = y_{i-1}$ for all $i \in \{1, \dots, n\}$.

Observe that $\lambda_{\boldsymbol{\sigma}_i,\hat{p}}(x_i) = \lambda_{\boldsymbol{\sigma}_{i-1},\hat{p}}(y_{i-1})$, and that

$$|x_i - y_i| \geq a_{\hat{p}}(\boldsymbol{\sigma}_i)|\lambda_{\boldsymbol{\sigma}_i,\hat{p}}(x_i) - \lambda_{\boldsymbol{\sigma}_i,\hat{p}}(y_i)|.$$

Thus

$$\begin{aligned}
|x - y| &= \sum_{i=0}^{n} |x_i - y_i| \\
&\geq \sum_{i=0}^{n} a_{\hat{p}}(\boldsymbol{\sigma}_i)|\lambda_{\boldsymbol{\sigma}_i,\hat{p}}(x_i) - \lambda_{\boldsymbol{\sigma}_i,\hat{p}}(y_i)| \\
&\geq a_0 \sum_{i=0}^{n} (\lambda_{\boldsymbol{\sigma}_i,\hat{p}}(x_i) - \lambda_{\boldsymbol{\sigma}_i,\hat{p}}(y_i)) \\
&= a_0(\lambda_{\boldsymbol{\sigma},\hat{p}}(x) - \lambda_{\boldsymbol{\sigma}_n,\hat{p}}(y)) = a_0\lambda_{\boldsymbol{\sigma},\hat{p}}(x) \\
&\geq a_0\alpha.
\end{aligned}$$

From the definition of thickness we find that $a_0 \geq t_0 m s_0$, yielding the second statement of the lemma. ◻

Lemma 12 allows us to quantify the distortion bound that we need to ensure that the boundary of $\underline{\mathrm{St}}(\hat{p})$ does not get mapped by F_p into the image of the open set V_p. The argument is the same as for Lemma 11, but there we were only concerned with the barycentre of the largest simplex, so the relative sizes of the simplices were not relevant as they are here (compare the bounds (4) and (5)).

Lemma 13 (boundary separation for $\boldsymbol{V_p}$) *Suppose $\underline{\mathrm{St}}(\hat{p})$ is a full star embedded in $\mathbb{R}^m$, and every m-simplex $\boldsymbol{\sigma}$ in $\underline{\mathrm{St}}(\hat{p})$ satisfies $s_0 \leq L(\boldsymbol{\sigma}) \leq L_0$, and $t(\boldsymbol{\sigma}) \geq t_0$. If the restriction of F_p to any m-simplex in $\underline{\mathrm{St}}(\hat{p})$ is a ξ-distortion map, with*

$$\xi < \frac{1}{6}\frac{m}{m+1}\frac{s_0}{L_0}t_0^2, \tag{5}$$

then $F_p(V_p) \cap F_p(|\partial\underline{\mathrm{St}}(\hat{p})|) = \emptyset$, where V_p is the set of points with barycentric coordinate with respect to $\hat{p}$ in a containing m-simplex strictly greater than $\frac{1}{m+1} - \delta$, with $\delta > 0$ an arbitrary, suffiently small parameter.

Proof If $x \in \boldsymbol{\sigma} \in \underline{\mathrm{St}}(\hat{p})$ has barycentric coordinate with respect to $\hat{p}$ larger than $\frac{1}{m+1} - \delta$, and $y \in \boldsymbol{\tau} \in \partial\underline{\mathrm{St}}(\hat{p})$, then Lemmas 10 and 12 ensure that $F_p(x) \neq F_p(y)$ provided

$$\frac{3\xi L(\boldsymbol{\sigma})}{t(\boldsymbol{\sigma})} + \frac{3\xi L(\boldsymbol{\tau})}{t(\boldsymbol{\tau})} \leq \left(\frac{1}{m+1} - \delta\right) m s_0 t_0,$$

which is satisfied by (5) when $\delta > 0$ satisfies

$$\delta \leq \frac{1}{m+1} - \frac{6L_0\xi}{m s_0 t_0^2}.$$

$\square$

When inequality (5) (and therefore also inequality (4)) is satisfied, we can employ the embedding lemma (Lemma 7) to guarantee that V_p is embedded:

Lemma 14 (local homeomorphism) *Suppose $\mathcal{A}$ is a compact m-manifold complex (without boundary), with vertex set $\mathcal{P}$, and M is an m-manifold. A map $H \colon |\mathcal{A}| \to M$ is a covering map if the following criteria are satisfied:*

(1) **compatible atlases** *There are compatible atlases for H, with $\widetilde{\mathcal{C}}_p = \underline{\mathrm{St}}(p)$ for each $p \in \mathcal{P}$* (Definition 4)*.*

(2) **simplex quality** *For each $p \in \mathcal{P}$, every simplex $\boldsymbol{\sigma} \in \underline{\mathrm{St}}(\hat{p}) = \widehat{\Phi}_p(\underline{\mathrm{St}}(p))$ satisfies $s_0 \leq L(\boldsymbol{\sigma}) \leq L_0$ and $t(\boldsymbol{\sigma}) \geq t_0$* (Notation 9)*.*

(3) **distortion control** *For each $p \in \mathcal{P}$, the map*

$$F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1} \colon |\underline{\mathrm{St}}(\hat{p})| \to \mathbb{R}^m,$$

when restricted to any m-simplex in $\underline{\mathrm{St}}(\hat{p})$, is an orientation-preserving ξ-distortion map with

$$\xi < \frac{m s_0 t_0^2}{6(m+1)L_0}$$

(Definitions 45 and 8)*.*

$\square$

2.2 Injectivity

Having established that H is a covering map, to ensure that H is injective it suffices to demonstrate that on each component of M there is a point with only a single point in its preimage. Injectivity follows since the number of points in the preimage is locally constant for covering maps.

Since each simplex is embedded by H, it is sufficient to show that for each vertex $q \in \mathcal{P}$, if $H(q) \in H(\boldsymbol{\sigma})$, then q is a vertex of $\boldsymbol{\sigma}$. This ensures that $H^{-1}(H(q)) = \{q\}$, and by Remark 5 each component of M must contain the image of a vertex.

In practice, we typically don't obtain this condition directly. The complex $\mathcal{A}$ is *constructed* by means of the local patches $\mathcal{C}_p$, and it is with respect to these patches that the vertices behave well.

Definition 15 (vertex sanity) If $H\colon |\mathcal{A}| \to M$ has compatible atlases (Definition 4), then H exibits *vertex sanity* if: for all vertices $p, q \in \mathcal{P}$, if $\phi_p \circ H(q) \in |\underline{\mathrm{St}}(\hat{p})| = \widehat{\Phi}_p(|\underline{\mathrm{St}}(p)|)$, then q is a vertex of $\underline{\mathrm{St}}(p)$.

Together with the distortion bounds that imposed on F_p, Definition 15 ensures that the image of a vertex cannot lie in the image of a simplex to which it does not belong:

Lemma 16 (injectivity) *If $H\colon |\mathcal{A}| \to M$ satisfies the hypotheses of Lemma 14 as well as Definition 15, then H is injective, and therefore a homeomorphism.*

Proof Towards a contradiction, suppose that $H(q) \in H(\boldsymbol{\sigma})$ and that q is not a vertex of the m-simplex $\boldsymbol{\sigma}$. This means there is some $x \in \boldsymbol{\sigma}$ such that $H(x) = H(q)$. Let p be a vertex of $\boldsymbol{\sigma}$. The vertex sanity hypothesis (Definition 15) implies that $\phi_p \circ H(q)$ must be either outside of $|\underline{\mathrm{St}}(\hat{p})|$, or belong to its boundary. Thus Lemmas 12 and 10, and the bound on ξ from Lemma 14(3) imply that the barycentric coordinate of x with respect to p must be smaller than $\frac{1}{m+1}$: Let $\hat{x} = \widehat{\Phi}_p(x)$, and $\hat{\boldsymbol{\sigma}} = \widehat{\Phi}_p(\boldsymbol{\sigma})$. Lemma 10 says that

$$|F_p(\hat{x}) - \hat{x}| \leq \frac{3\xi L_0}{t_0} < \frac{m s_0 t_0}{2(m+1)} \leq \frac{a_0}{2(m+1)},$$

where a_0 is a lower bound on the altitudes of $\hat{p}$, as in Lemma 12. Since $F_p(\hat{x}) = \phi_p \circ H(x)$ is at least as far away from $\hat{x}$ as $\partial\underline{\mathrm{St}}(\hat{p})$, Lemma 12 implies that the barycentric coordinate of $\hat{x} \in \hat{\boldsymbol{\sigma}}$ with respect to $\hat{p}$ must be no larger than $\frac{1}{2(m+1)}$. Since $\widehat{\Phi}_p$ preserves barycentric coordinates, and the argument works for any vertex p of $\boldsymbol{\sigma}$, we conclude that all the barycentric coordinates of x in $\boldsymbol{\sigma}$ are strictly less than $\frac{1}{m+1}$. We have reached a contradiction with the fact that the barycentric coordinates of x must sum to 1. $\square$

2.3 Main result

To recap, Lemmas 14 and 16 yield the following triangulation result. In the bound on ξ from Lemma 14(3), we replace the factor $\frac{m}{m+1}$ with $\frac{1}{2}$, the lower bound attained when $m = 1$.

Theorem 17 (triangulation) *Suppose $\mathcal{A}$ is a compact m-manifold complex (without boundary), with vertex set $\mathcal{P}$, and M is an m-manifold. A map $H\colon |\mathcal{A}| \to M$ is a homeomorphism if the following criteria are satisfied:*

(1) **compatible atlases** *There are compatible atlases*

$$\{(\widetilde{\mathcal{C}}_p, \widehat{\Phi}_p)\}_{p\in\mathcal{P}}, \quad \widetilde{\mathcal{C}}_p \subset \mathcal{A}, \quad \textit{and} \quad \{(U_p, \phi_p)\}_{p\in\mathcal{P}}, \quad U_p \subset M,$$

for H, with $\widetilde{\mathcal{C}}_p = \underline{\mathrm{St}}(p)$ for each $p \in \mathcal{P}$, the vertex set of $\mathcal{A}$ (Definition 4)*.*

(2) **simplex quality** *For each $p \in \mathcal{P}$, every simplex $\boldsymbol{\sigma} \in \underline{\mathrm{St}}(\hat{p}) = \widehat{\Phi}_p(\underline{\mathrm{St}}(p))$ satisfies $s_0 \leq L(\boldsymbol{\sigma}) \leq L_0$ and $t(\boldsymbol{\sigma}) \geq t_0$* (Notation 9)*.*

(3) **distortion control** *For each $p \in \mathcal{P}$, the map*

$$F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1}\colon |\underline{\mathrm{St}}(\hat{p})| \to \mathbb{R}^m,$$

when restricted to any m-simplex in $\underline{\mathrm{St}}(\hat{p})$, is an orientation-preserving ξ-distortion map with

$$\xi < \frac{s_0 t_0^2}{12 L_0}$$

(Definitions 45 and 8).

(4) **vertex sanity** *For all vertices $p, q \in \mathcal{P}$, if $\phi_p \circ H(q) \in |\underline{\mathrm{St}}(\hat{p})|$, then q is a vertex of $\underline{\mathrm{St}}(p)$.*

$\square$

Remark 18 The constants L_0, s_0, and t_0 that constrain the simplices in the local complex $\underline{\mathrm{St}}(\hat{p})$, and the metric distortion of F_p in Theorem 17 can be considered to be local, i.e., they may depend on $p \in \mathcal{P}$.

3 Metric and differentiable distortion maps

In this short section we gather some useful lemmas on general distortion maps, and on differentiable distortion maps.

3.1 Metric distortion maps

The definition of a distortion map (Definition 8) makes sense in a more general context: a map $F : (X, d_X) \to (Y, d_Y)$ between metric spaces is a ξ-*distortion map* if

$$|d_Y(F(x), F(y)) - d_X(x, y)| \leq \xi d_X(x, y) \quad \text{for all } x, y \in X. \tag{6}$$

Lemma 19 (inverse and composition of distortion maps) (1) *If $F : (X, d_X) \to (Y, d_Y)$ is a ξ-distortion map with $\xi < 1$, then F^{-1} is a $\frac{\xi}{1-\xi}$-distortion map.*

(2) *Suppose $F_i : (X_i, d_{X_i}) \to (X_{i+1}, d_{X_{i+1}})$, $1 \leq i \leq k$, are respectively ξ_i-distortion maps. Then*

$$F_k \circ F_{k-1} \circ \cdots \circ F_1 : X_1 \to X_{k+1}$$

is a $\left(\sum_{W \in \wp(\{k\})} \prod_{i \in W} \xi_i\right)$-distortion map, where $\wp(\{k\})$ is the set of nonempty subsets of $\{1, \ldots, k\}$.

In particular, $F_2 \circ F_1$ is a $(\xi_1 + \xi_2 + \xi_1\xi_2)$-distortion map.

Proof (1) Let $u = F(x)$ and $v = F(y)$. Then (6) becomes

$$|d_X(F^{-1}(u), F^{-1}(v)) - d_Y(u, v)| \leq \xi d_X(F^{-1}(u), F^{-1}(v)). \tag{7}$$

But it follows from (7) that $d_X(F^{-1}(u), F^{-1}(v)) \leq \frac{1}{1-\xi} d_Y(u, v)$, and plugging this back into (7) yields the result.

(2) Using the observation that

$$d_{X_2}(F_1(x), F_1(y)) \leq (1 + \xi_1) d_{X_1}(x, y),$$

we find

$$\begin{aligned}
&\left|d_{X_3}(F_2 \circ F_1(x), F_2 \circ F_1(x)) - d_{X_1}(x,y)\right| \\
&\qquad \le \left|d_{X_3}(F_2 \circ F_1(x), F_2 \circ F_1(x)) - d_{X_2}(F_1(x), F_1(y))\right| \\
&\qquad\qquad\qquad\qquad\qquad\qquad + \left|d_{X_2}(F_1(x), F_1(y)) - d_{X_1}(x,y)\right| \\
&\qquad \le \xi_2 d_{X_2}(F_1(x), F_1(y)) + \xi_1 d_{X_1}(x,y) \\
&\qquad \le \xi_2(1+\xi_1) d_{X_1}(x,y) + \xi_1 d_{X_1}(x,y) \\
&\qquad = (\xi_1 + \xi_2 + \xi_2\xi_1) d_{X_1}(x,y).
\end{aligned}$$

This establishes the bound for the composition of two maps, but it also serves as the inductive step for the general bound. If G is the η-distortion map defined by $G = F_{k-1} \circ \cdots \circ F_1$, then $F_k \circ G$ is a $(\xi_k+\eta+\xi_k\eta)$-distortion map. Now notice that

$$\sum_{W \in \wp(\{k\})} \prod_{i \in W} \xi_i = \xi_k + \left(\sum_{W \in \wp(\{k-1\})} \prod_{i \in W} \xi_i \right) + \xi_k \left(\sum_{W \in \wp(\{k-1\})} \prod_{i \in W} \xi_i \right).$$

□

3.2 Differentiable distortion maps

Although the homeomorphism demonstration that yields Theorem 17 makes no explicit requirement of differentiability, it is convenient to exploit differentiability when it is available. We collect here some observations relating metric distortion and bounds on the differential of a map.

Recall that if $F\colon U \subseteq \mathbb{R}^m \to \mathbb{R}^m$ is differentiable, then the *differential* of F at x is the linear map defined by

$$dF_x(v) = \frac{d}{dt} F \circ \alpha(t)\Big|_{t=0},$$

where $\alpha\colon I \subset \mathbb{R} \to \mathbb{R}^m$ is any curve such that $\alpha(0) = x$ and $\alpha'(0) = v$, where α' is the derivative with respect to t.

Bounds on the differential of F are closely related to the metric distortion of F. If $A\colon U \subseteq \mathbb{R}^m \to \mathbb{R}^m$ is a linear map, we let $\|A\|$ denote the operator norm: $\|A\| = \sup_{|v|=1} |Av|$. Associated with A are m nonnegative numbers called the *singular values* of A, denoted $s_i(A)$, $1 \le i \le m$, ordered such that $s_i(A) \ge s_j(A)$ if $i \le j$. We only mention the singular values because they provide a notational convenience. The largest singular value is defined by $s_1(A) = \|A\|$, and the smallest singular value is $s_m(A) = \inf_{|v|=1} |Av|$.

Lemma 20 *If $F\colon U \subseteq \mathbb{R}^m \to \mathbb{R}^m$ is a differentiable ξ-distortion map, then*

$$|s_i(dF_x) - 1| \leq \xi \quad \text{for all } x \in U \text{ and } 1 \leq i \leq m.$$

Proof For $x \in U$, and $v \in T_x\mathbb{R}^m = \mathbb{R}^m$ with $|v| = 1$, let $\alpha(t) = x + tv$. Since F is a ξ-distortion map, we have

$$(1-\xi)\,|(x+tv)-x| \leq |F(x+tv) - F(x)| \leq (1+\xi)\,|(x+tv)-x|\,,$$

so for all $t \neq 0$,

$$1 - \xi \leq \frac{|F(x+tv) - F(x)|}{|t|} \leq 1 + \xi.$$

Since F is differentiable,

$$\lim_{t\to 0} \frac{|F(x+tv) - F(x)|}{|t|} = \left|\lim_{t\to 0} \frac{F(x+tv) - F(x)}{t}\right| = |dF_x(v)|\,,$$

so

$$1 - \xi \leq |dF_x(v)| \leq 1 + \xi,$$

which yields the claimed result. $\square$

So a bound on the metric distortion of a differentiable map directly yields the same bound on the amount the singular values of the differential can differ from 1. We are interested in a converse assertion: we want to bound the metric distortion of F, given a bound on (the singular values of) its differential. This can only be done with caveats.

Lemma 21 *Suppose $\boldsymbol{\omega}$ is a convex set, $\boldsymbol{\omega} \subseteq U \subseteq \mathbb{R}^m$, and $F\colon U \to \mathbb{R}^m$ is a differentiable map such that $F(\boldsymbol{\omega}) \subseteq V \subseteq F(U)$ for some convex set V. If F is injective, and*

$$|s_i(dF_x) - 1| \leq \xi < 1, \quad \text{for all } x \in U \text{ and } 1 \leq i \leq m,$$

then $F|_{\boldsymbol{\omega}}$ is a ξ-distortion map.

Proof Since $\boldsymbol{\omega}$ is convex, the length of the image of the line segment between x and y provides an upper bound on the distance between $F(x)$ and $F(y)$, and this length can be bounded because of the bound on dF:

$$|F(y) - F(x)| \leq \int_0^1 \left|dF_{x+t(y-x)}(y-x)\right|\, dt \leq (1+\xi)\,|y-x|\,. \tag{8}$$

To get a lower bound we use the fact that injectivity and the bound on the singular values imply that F is invertible. For any point $z = F(x)$ we have

$$\left\| dF_z^{-1} \right\| = s_m(dF_x)^{-1} \leq (1 - \xi)^{-1}.$$

Since, for $x, y \in \boldsymbol{\omega}$ the segment $[F(x), F(y)]$ is contained in $F(U)$, we can use the same argument as in (8), but using F^{-1} instead of F, so

$$|x - y| \leq (1 - \xi)^{-1} |F(x) - F(y)| .$$

Therefore, combining with (8) we have

$$(1 - \xi) |x - y| \leq |F(x) - F(y)| \leq (1 + \xi) |x - y| .$$

□

4 Submanifolds of Euclidean space

As a specific application of Theorem 17, we consider a smooth (or at least C^2) compact m-dimensional submanifold of Euclidean space: $M \subset \mathbb{R}^N$. A simplicial complex $\mathcal{A}$ is built whose vertices are a finite set $\mathcal{P}$ sampled from the manifold: $\mathcal{P} \subset M$. The motivating model for this setting is the tangential Delaunay complex [BG14]. In that case $\mathcal{A}$ is constructed as a subcomplex of a weighted Delaunay triangulation of the ambient space $\mathbb{R}^N$, so it is necessarily embedded. However, in general we do not need to assume *a priori* that $\mathcal{A}$ is embedded in $\mathbb{R}^N$. Instead, we assume only that the embedding of the vertex set $\mathcal{P} \hookrightarrow \mathbb{R}^N$ defines an *immersion* $\iota \colon |\mathcal{A}| \to \mathbb{R}^N$. By this we mean that for any vertex $p \in \mathcal{P}$ we have that the restriction of ι to $|\underline{\mathrm{St}}(p)|$ is an embedding.

At each point $x \in M$, the tangent space $T_xM \subset T_x\mathbb{R}^N$ is naturally viewed as an affine flat in $\mathbb{R}^N$, with the vector-space structure defined by taking the distinguished point x as the origin. The maps involved in Theorem 17 will be defined by projection maps. The coordinate charts are defined using the orthogonal projection $\mathrm{pr}_{T_pM} \colon \mathbb{R}^N \to T_pM$. As discussed in Section 4.3, for a sufficiently small neighbourhood $U_p \subset M$, we obtain an embedding

$$\phi_p = \mathrm{pr}_{T_pM}|_{U_p} \colon U_p \subset M \to T_pM \cong \mathbb{R}^m,$$

which will define our coordinate maps for M.

For the map $H \colon |\mathcal{A}| \to M$, we will employ the closest point projection map defined in Section 4.1 and discussed further in Section 4.4. There is an open neighbourhood $U_M \subset \mathbb{R}^N$ of M on which each point has a unique closest point on M, so the closest-point projection map $\mathrm{pr}_M \colon U_M \to M$ is well-defined. We define $H = \mathrm{pr}_M \circ \iota$.

As demanded by Definition 4, for each $p \in \mathcal{P}$ the coordinate map $\widehat{\Phi}_p$ for $\mathcal{A}$ is the secant map of $\phi_p \circ H$ restricted to $\widetilde{\mathcal{C}}_p = |\underline{\mathrm{St}}(p)|$, and since pr_{T_pM} is already a linear map, and pr_M is the identity on the vertices, this means $\widehat{\Phi}_p = \mathrm{pr}_{T_pM} \circ \iota|_{|\underline{\mathrm{St}}(p)|}$.

In Sections 4.1 and 4.2 we review some of the geometric concepts and standard results that we will use in the rest of the section. In order to bound the metric distortion of the maps $F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1}$ via Lemma 19, we are free to choose any convenient metric on M. As is common in computational geometry, we employ here the metric of the ambient space $\mathbb{R}^N$, rather than the intrinsic metric of geodesic distances.

4.1 Submanifold geometry

Since $M \subset \mathbb{R}^N$ is compact, for any $x \in \mathbb{R}^N$ there is a point $z \in M$ that realizes the distance to M, i.e.,

$$\delta_M(x) := d_{\mathbb{R}^N}(x, M) := \inf_{y \in M} d_{\mathbb{R}^N}(x, y) = d_{\mathbb{R}^N}(x, z).$$

The *medial axis* of M is the set of points $\mathrm{ax}(M) \subset \mathbb{R}^N$ that have more than one such closest point on M. In other words, if $x \in \mathrm{ax}(M)$, then an open ball $B_{\mathbb{R}^N}(x, r)$, centred at x and of radius $r = \delta_M(x)$, will be tangent to M at two or more distinct points. The *cut locus* of M is the closure of the medial axis, and is denoted $\overline{\mathrm{ax}}(M)$. The *reach* of M is defined by $\mathrm{rch}(M) := d_{\mathbb{R}^N}(M, \overline{\mathrm{ax}}(M))$. We observe below that for compact C^2 submanifolds, $\mathrm{rch}(M) > 0$. Thus, by definition, every point x in the open neighbourhood $U_M := \mathbb{R}^N \setminus \overline{\mathrm{ax}}(M)$ of M has a unique closest point $\check{x} \in M$. The *closest-point projection* map

$$\mathrm{pr}_M \colon U_M \to M,$$

takes x to this closest point: $\mathrm{pr}_M(x) = \check{x}$.

To each point $x \in M$, we associate a *normal space*

$$N_xM = \{n \in T_x\mathbb{R}^N \mid n \cdot v = 0 \ \forall v \in T_xM\}$$

of vectors orthogonal to T_xM. Thus $T_x\mathbb{R}^N = N_xM \oplus T_xM$. As with the tangent space, the normal space at x is naturally identified with an affine flat in $\mathbb{R}^N$. It has dimension $k = N - m$ and is orthogonal to T_xM.

The tubular neighbourhood theorem is a well known result in differential topology. The statement presented here is adapted from [Fed59, Theorem 4.8(13)], and the regularity assertions in the second paragraph are demonstrated in [Foo84].

Theorem 22 (tubular neighbourhood) *There is a natural structure on U_M given by partitioning it into subsets of points that all project via pr_M onto the same point of M. This allows us to identify U_M as a portion of the* normal bundle *of M,*

$$NM := \{(x, n) \in \mathbb{R}^N \times \mathbb{R}^N \mid x \in M,\ n \in N_x M\}.$$

The map $\psi : U_M \to NM$ given by $x \mapsto (x, x - \mathrm{pr}_M(x))$ is a diffeomorphism onto its image, with inverse $\psi(U_M) \subset NM \to \mathbb{R}^N$ given by $(x, n) \mapsto x + n$.

If M is a C^j submanifold, then ψ is a C^{j-1} diffeomorphism onto its image [Foo84], and the function δ_M has is C^j on $U_M \setminus M$.

If $x \in U_M$, with $\check{x} = \mathrm{pr}_M(x)$, then the ball $B_{\mathbb{R}^N}(x, r)$ of radius $r = |x - \check{x}|$ is tangent to M at $\check{x}$, and $(x - \check{x}) \in N_{\check{x}} M$. For any point $y \in M$, the *local reach* [AEM07] of M at y is

$$\mathrm{rch}(y, M) := \sup\{r \in \mathbb{R} \mid (y + ru) \in U_M\ \forall u \in N_y M \text{ with } |u| = 1\}.$$

By the tubular neighbourhood theorem, $\mathrm{rch}(y, M)$ can be equivalently defined as

$$\mathrm{rch}(y, M) = \sup\{r \in \mathbb{R} \mid \mathrm{pr}_M(y + ru) = y\ \forall u \in N_y M \text{ with } |u| = 1\}.$$

With this formulation, it is easy to see the following standard observation:

Lemma 23 *For any $y \in M$, any open ball that is tangent to M at y and with radius $r \leq \mathrm{rch}(y, M)$, does not intersect M.*

This property is useful for bounding the extrinsic curvatures of M at y (see, e.g., [BLW17a, Lemma 3.3]). However, it can be awkward to work with the local reach since it is not known to be continuous, even when M is smooth. The *local feature size* at $y \in M$ is the distance from y to the medial axis; equivalently, it is the supremum of the radii of balls centred at y and contained in U_M [Fed59, p. 432]:

$$\mathrm{lfs}(y) := \sup\{r \mid B_{\mathbb{R}^N}(y, r) \subset U_M\}.$$

Although it was introduced by Federer, the local feature size was later rediscovered, and given its name, by Amenta and Bern [AB99]. Since the local reach at y is the distance to the medial axis measured only in directions orthogonal to M at y, we have $\mathrm{rch}(y, M) \geq \mathrm{lfs}(y)$. It is a short exercise to show that the local feature size is 1-Lipschitz:

$$|\mathrm{lfs}(y) - \mathrm{lfs}(x)| \leq |x - y|, \quad \text{for all } x, y \in M.$$

The *reach* of M is the infimum of the local feature size, or equivalently, the infimum of the local reach:

$$\mathrm{rch}(M) := d_{\mathbb{R}^N}(M, \overline{\mathrm{ax}}(M)) = \inf_{y \in M} \mathrm{lfs}(y) = \inf_{y \in M} \mathrm{rch}(y, M).$$

The last equality comes from the observation that since M is compact, there is some y^* for which $\mathrm{lfs}(y^*) = \mathrm{rch}(M)$, and we must have $\mathrm{rch}(y^*, M) = \mathrm{lfs}(y^*)$ because $z - y^*$ must lie in $N_{y^*}M$ for all $z \in \overline{\mathrm{ax}}(M) \cap \overline{B}_{\mathbb{R}^N}(y^*, \mathrm{lfs}(y^*))$.

Observe also, that since, by the tubular neighbourhood theorem, $\mathrm{lfs}(x) > 0$ at any point $x \in M$, the continuity of the local feature size implies that $\mathrm{rch}(M) > 0$ for any compact C^2 submanifold.

Remark 24 (locally bounding the local reach) We will often need a lower bound R_{rch} on $\mathrm{rch}(x, M)$ in a neighbourhood of a point $p \in M$, and moreover, the *size* of this neighbourhood will depend on the size of the bound. For example, we will have the following awkward self-referential definition: $U_p = B_{\mathbb{R}^N}(p, r) \cap M$, where $r \leq R_{\mathrm{rch}}$ and $\mathrm{rch}(x, M) \geq R_{\mathrm{rch}}$ for all $x \in U_p$.

This can be easily resolved if we choose $R_{\mathrm{rch}} = \mathrm{rch}(M)$, the global bound on the local reach. However, the local reach could vary by orders of magnitude over the manifold, making it inefficient to use a global bound to govern the size of the simplices in the constructed simplicial approximation.

The Lipschitz property of lfs makes it well-suited to bound the local reach in a small neighbourhood of $p \in M$. For example,

$$\mathrm{rch}(x, M) \geq \mathrm{lfs}(x) \geq (1 - \epsilon)\,\mathrm{lfs}(p), \quad \text{for all } x \text{ with } |p - x| \leq \epsilon\,\mathrm{lfs}(p).$$

Thus we can choose $r = \epsilon\,\mathrm{lfs}(p)$ and $R_{\mathrm{rch}} = (1 - \epsilon)\,\mathrm{lfs}(p)$, and the criterion $r \leq R_{\mathrm{rch}}$ is satisfied provided $\epsilon \leq \frac{1}{2}$.

4.2 Affine flats and angles

The angle between two vectors $u, v \in \mathbb{R}^N \setminus \{0\}$ is denoted $\angle(u, v)$ (this angle is $\leq \pi$). If $K \subseteq \mathbb{R}^n$ is a linear subspace, $\mathrm{pr}_K(u)$ is the orthogonal projection of into K. We define $\angle(u, K)$ to be $\pi/2$ if $\mathrm{pr}_K(u) = 0$, and otherwise $\angle(u, K) = \angle(u, \mathrm{pr}_K(u))$ (thus $\angle(u, K) \leq \pi/2$). If K and L are two linear subspaces, then

$$\angle(K, L) = \sup_{u \in K \setminus \{0\}} \angle(u, L).$$

The definition is only interesting when $\dim K \leq \dim L$. Observe that if $\dim K = \dim L$, then $\angle(K, L) = \angle(L, K)$. If K and L are affine flats in $\mathbb{R}^N$, then $\angle(K, L)$ is the angle between the corresponding parallel vector subspaces. If $\boldsymbol{\sigma}$ is a simplex with $\dim \boldsymbol{\sigma} \leq \dim L$, then $\angle(\boldsymbol{\sigma}, L) := \angle(\mathrm{aff}(\boldsymbol{\sigma}), L)$, where $\mathrm{aff}(\boldsymbol{\sigma})$ is the affine hull of $\boldsymbol{\sigma}$. We denote the orthogonal complement of a linear subspace $K \subseteq \mathbb{R}^N$ by $K^\perp$. A short exercise yields the following observations:

Lemma 25 (1) *If K, L are subspaces of $\mathbb{R}^N$, then*

$$\angle(L^\perp, K^\perp) = \angle(K, L).$$

(2) *If $Q \subset \mathbb{R}^N$ is a subspace of codimension 1, then*

$$\angle(Q^\perp, K) = \pi/2 - \angle(K, Q).$$

The following standard observation (cf. [Fed59, Theorem 4.8(7)]) follows easily from Lemma 23.

Lemma 26 *Given any two points $x, y \in M \subset \mathbb{R}^N$, we have*

(1) $\sin \angle([x, y], T_x M) \leq \frac{|x-y|}{2\,\mathrm{rch}(x,M)}$.

(2) $d_{\mathbb{R}^N}(y, T_x M) \leq \frac{|x-y|^2}{2\,\mathrm{rch}(x,M)}$.

The following lemma is a local adaptation of results presented in [BLW17a]:

Lemma 27 (tangent space variation) *Suppose $B_{\mathbb{R}^N}(c, r) \subset U_M$, $x, y \in B = B_{\mathbb{R}^N}(c, r) \cap M$ and $\mathrm{rch}(z, M) \geq R_{\mathrm{rch}}$ for all z in B. If $r < R_{\mathrm{rch}}$, then*

$$\sin \angle(T_x M, T_y M) \leq \frac{|y - x|}{R_{\mathrm{rch}}},$$

and

$$\angle(T_x M, T_y M) \leq \frac{\pi\,|y - x|}{2 R_{\mathrm{rch}}}.$$

Proof See Appendix B for a sketch of how this result follows from the arguments presented in [BLW17a]. $\square$

Remark 28 Since the bounds in Lemma 27 are vacuous if $|y - x| \geq R_{\mathrm{rch}}$, in practice we require that either $c = x$ or $r < \frac{1}{2} R_{\mathrm{rch}}$.

We will need to bound the angle between a simplex with vertices on M and the nearby tangent spaces. To this end we employ a result established by Whitney [Whi57, p. 127] in the formulation presented in [BDG13, Lemma 2.1]:

Lemma 29 (Whitney's angle bound) *Suppose $\boldsymbol{\sigma}$ is a j-simplex whose vertices all lie within a distance η from a k-dimensional affine space, $K \subset \mathbb{R}^N$, with $k \geq j$. Then*

$$\sin \angle(\boldsymbol{\sigma}, K) \leq \frac{2\eta}{tL},$$

where t is the thickness of $\boldsymbol{\sigma}$ and L is the length of its longest edge.

Lemma 30 (simplices lie close to $\boldsymbol{M}$) *Let $\boldsymbol{\sigma} \subset U_M$ be a simplex with vertices on M, and R_{rch} a constant such that $\mathrm{rch}(\check{x}, M) \geq R_{\mathrm{rch}}$ for all $x \in \boldsymbol{\sigma}$, where $\check{x} = \mathrm{pr}_M(x)$. Then for all $x, y \in \boldsymbol{\sigma}$,*

$$d_{\mathbb{R}^N}(y, T_{\check{x}}M) < \frac{2L(\boldsymbol{\sigma})^2}{R_{\mathrm{rch}}}, \quad \textit{and in particular,} \quad \delta_M(x) < \frac{2L(\boldsymbol{\sigma})^2}{R_{\mathrm{rch}}}.$$

Proof Since $\check{x}$ is the closest point on M to x, and every vertex lies on M, we must have $|x - \check{x}| < L(\boldsymbol{\sigma})$, and therefore, for any vertex p of $\boldsymbol{\sigma}$,

$$|p - \check{x}| \leq |p - x| + |x - \check{x}| < 2L(\boldsymbol{\sigma}). \tag{9}$$

By Lemma 26(2) we have $d_{\mathbb{R}^N}(p, T_{\check{x}}M) < 2L^2/R_{\mathrm{rch}}$. This is true for all vertices p of $\boldsymbol{\sigma}$, and since the function $d_{\mathbb{R}^N}(\,\cdot\,, T_{\check{x}}M)$ is affine on $\boldsymbol{\sigma}$, it is also true for any $y \in \boldsymbol{\sigma}$. The second inequality follows by taking $y = x$, since

$$d_{\mathbb{R}^N}(x, T_{\check{x}}M) = |x - \check{x}|\,.$$

$\square$

Remark 31 If p is a vertex of $\boldsymbol{\sigma}$, then the constraint $\boldsymbol{\sigma} \subset U_M$ of Lemma 30 is ensured if $L(\boldsymbol{\sigma}) < \mathrm{lfs}(p)$, since lfs is the distance to the medial axis. In practice R_{rch} is defined either in terms of $\mathrm{rch}(M)$ or in terms of $\mathrm{lfs}(p)$.

For example, since $|p - \check{x}| \leq 2L(\boldsymbol{\sigma})$, we have that $\mathrm{lfs}(\check{x}) \geq \mathrm{lfs}(p) - 2L$. So the requirements of Lemma 30 are satisfied by demanding $L \leq \epsilon\, \mathrm{lfs}(p)$, with $\epsilon < \frac{1}{2}$, and setting $R_{\mathrm{rch}} = (1 - 2\epsilon)\, \mathrm{lfs}(p)$. Alternatively, we can simply demand $L < \mathrm{rch}(M)$, and use $R_{\mathrm{rch}} = \mathrm{rch}(M)$. Of course, other variations are possible.

Lemma 32 (simplex-tangent space angle bounds) *Suppose* $\boldsymbol{\sigma} \subset \mathbb{R}^N$ *is a simplex of dimension* $\leq m$ *with vertices on* M*. If* p *is a vertex of* $\boldsymbol{\sigma}$*, then*

(1)
$$\sin \angle(\boldsymbol{\sigma}, T_p M) \leq \frac{L}{t \, \mathrm{rch}(p, M)}.$$

(2) *In addition, suppose* $\boldsymbol{\sigma} \subset U_M$*, and there is a ball* $B_{\mathbb{R}^N}(c, r) \subset U_M$ *such that for any* $x \in \boldsymbol{\sigma}$*,* $\check{x} = \mathrm{pr}_M(x) \in B = B_{\mathbb{R}^N}(c, r) \cap M$ *and* $\mathrm{rch}(z, M) \geq R_{\mathrm{rch}}$ *for all* $z \in B$*. If* $r < R_{\mathrm{rch}}$*, then*

$$\sin \angle(\boldsymbol{\sigma}, T_{\check{x}} M) \leq \frac{3L}{t R_{\mathrm{rch}}}.$$

Proof (1) By Lemma 26(2), all the vertices of $\boldsymbol{\sigma}$ are within a distance $\eta = L^2/(2 \, \mathrm{rch}(p, M))$ from $T_p M$, and so Lemma 29 ensures

$$\sin \angle(\boldsymbol{\sigma}, T_p M) \leq \frac{2\eta}{tL} = \frac{L}{t \, \mathrm{rch}(p, M)}.$$

(2) By Lemma 27 and (9),

$$\sin \angle(T_p M, T_{\check{x}} M) \leq \frac{|p - \check{x}|}{R_{\mathrm{rch}}} \leq \frac{2L}{R_{\mathrm{rch}}},$$

and the result follows using part 1. $\square$

4.3 Distortion of orthogonal projection: $\widehat{\Phi}_p$ and ϕ_p

The coordinate maps ϕ_p and $\widehat{\Phi}_p$ are defined in terms of the orthogonal projection to $T_p M$. The size of the neighbourhoods used to define the coordinate charts are constrained by the requirement that these maps be embeddings, which we establish by ensuring that they are ξ-distortion maps with $\xi < 1$.

Lemma 33 (definition and distortion of $\boldsymbol{\phi_p}$) *Let* $U_p = B_{\mathbb{R}^N}(p, r) \cap M$*, where* $r = \rho R_{\mathrm{rch}}$*, with* $\rho < \frac{1}{2}$*, and* $\mathrm{rch}(x, M) \geq R_{\mathrm{rch}}$ *for all* $x \in U_p$ (see Remark 24)*. Define* $\phi_p := \mathrm{pr}_{T_p M}|_{U_p}$*. Then* ϕ_p *is a* ξ*-distortion map with*

$$\xi = 4\rho^2.$$

Proof For any distinct $x, y \in U_p$, we have, from Lemma 26(1), that

$$\sin \angle([x,y], T_x M) \leq \rho,$$

and, from Lemma 27, that

$$\sin \angle(T_p M, T_x M) \leq \rho.$$

Combining these bounds we have

$$\sin \angle([x,y], T_p M) \leq 2\rho.$$

Letting $\widetilde{x} = \phi_p(x)$, $\widetilde{y} = \phi_p(y)$, and $\theta = \angle([x,y], T_p M)$, we find

$$\begin{aligned} |x-y| - |\widetilde{x} - \widetilde{y}| &= (1 - \cos\theta)\, |x-y| \\ &\leq (1 - \sqrt{1 - 4\rho^2})\, |x-y| \\ &\leq 4\rho^2 \, |x-y| \, . \end{aligned}$$

The result follows, since $|\widetilde{x} - \widetilde{y}| \leq |x-y|$. $\square$

Remark 34 (differential of $\boldsymbol{\phi_p}$) It straight forward to verify from the definitions that for any $x \in U_p$,

$$d(\mathrm{pr}_{T_p M}|_M)_x = \mathrm{pr}_{T_p M}|_{T_x M}.$$

The domain of the map $\widehat{\Phi}_p$, i.e., an upper bound on the allowable size of the simplices in $\iota(|\underline{\mathrm{St}}(p)|)$, is governed by the following bound on the metric distortion of the projection from a simplex.

Lemma 35 (simplexwise distortion of $\widehat{\boldsymbol{\Phi}}_{\boldsymbol{p}}$) *Suppose $\boldsymbol{\sigma} \subset \mathbb{R}^N$ is a simplex of dimension $\leq m$ with vertices on M. If p is a vertex of $\boldsymbol{\sigma}$, and*

$$L(\boldsymbol{\sigma}) < t(\boldsymbol{\sigma}) \,\mathrm{rch}(p, M),$$

then the restriction of $\mathrm{pr}_{T_p M}$ to $\boldsymbol{\sigma}$ is a ξ-distortion map with

$$\xi = \left(\frac{L(\boldsymbol{\sigma})}{t(\boldsymbol{\sigma}) \,\mathrm{rch}(p, M)} \right)^2 .$$

Proof Let $x, y \in \boldsymbol{\sigma}$ and set $\hat{x} = \mathrm{pr}_{T_pM}(x),\ \hat{y} = \mathrm{pr}_{T_pM}(y)$. By Lemma 32(1),

$$\sin \angle(\boldsymbol{\sigma}, T_pM) \leq \frac{L}{t\,\mathrm{rch}(p, M)}.$$

So, putting $\theta = \angle([x, y], [\hat{x}, \hat{y}]) \leq \angle(\boldsymbol{\sigma}, T_pM)$ we find

$$\begin{aligned} |x - y| - |\hat{x} - \hat{y}| &= (1 - \cos\theta)\,|x - y| \\ &\leq \left(1 - \sqrt{1 - \left(\frac{L}{t\,\mathrm{rch}(p, M)}\right)^2}\right)|x - y| \\ &\leq \left(\frac{L}{t\,\mathrm{rch}(p, M)}\right)^2 |x - y|\,. \end{aligned}$$

The result follows since $|\hat{x} - \hat{y}| \leq |x - y|$. $\square$

Therefore, in order to use our framework, we need to ensure that for each simplex $\boldsymbol{\sigma} \in \mathcal{A}$, the simplex $\iota(\boldsymbol{\sigma}) \subset \mathbb{R}^N$ satisfies $L < tR_{\mathrm{rch}}$, where $R_{\mathrm{rch}} \leq \mathrm{rch}(p, M)$ for each vertex p of $\iota(\boldsymbol{\sigma})$ (we will in fact need a stronger bound than this). Then, in conformance with Section 2.1.1, we set $\widehat{\Phi}_p = \mathrm{pr}_{T_pM} \circ \iota|_{|\underline{\mathrm{St}}(p)|}$, and we *require* that it be an embedding. Although we have established that pr_{T_pM} is an embedding on each simplex, and we have assumed that $\iota|_{|\underline{\mathrm{St}}(p)|}$ is an embedding, these criteria do not imply that $\widehat{\Phi}_p$ is an embedding. We leave this as a requirement of the embedding theorem (requirement (a) of Theorem 39), i.e., something that needs to be established in context. For the case of the tangential Delaunay complex, the embedding follows naturally because $\phi_p(|\underline{\mathrm{St}}(p)|)$ is seen as a weighted Delaunay triangulation in T_pM [BG14].

4.4 Distortion of the closest-point projection map: H

Recall that $H\colon |\mathcal{A}| \to M$ is the map that we wish to show is a homeomorphism. In our current context H is based on the closest-point projection map $\mathrm{pr}_M\colon \mathbb{R}^N \to M$. As discussed at the beginning of this section, we define $H = \mathrm{pr}_M \circ \iota$, where $\iota\colon |\mathcal{A}| \to \mathbb{R}^N$ is the immersion of our simplicial complex into $\mathbb{R}^N$. Once H is shown to be a homeomorphism, it follows that ι is in fact an embedding, but we don't assume this a priori. The metric on $|\mathcal{A}|$ (i.e., the edge lengths of the Euclidean simplices) is *defined* by ι, so ι itself does not contribute to the metric distortion of H.

An upper bound for the metric distortion of pr_M was demonstrated by [Fed59, Theorem 4.8(8)]; the proof we present here is similar, but less general, since we require M to be a differentiable submanifold:

Lemma 36 (upper bound for pr_M distortion) *Let $x, y \in U_M$ and $R_{\mathrm{rch}} = \min\{\mathrm{rch}(\check{x}, M), \mathrm{rch}(\check{y}, M)\}$, where $\check{x} = \mathrm{pr}_M(x)$, and $\check{y} = \mathrm{pr}_M(y)$, as usual. If $a \geq \max\{\delta_M(x), \delta_M(y)\}$ for some $a < R_{\mathrm{rch}}$, then*

$$|\check{y} - \check{x}| \leq \left(1 - \frac{a}{R_{\mathrm{rch}}}\right)^{-1} |y - x| . \tag{10}$$

Proof Let $\widetilde{x}$, $\widetilde{y}$ be the orthogonal projection of x and y into the line ℓ generated by $\check{y} - \check{x}$. Then

$$|y - x| \geq |\widetilde{y} - \widetilde{x}| \geq |\check{y} - \check{x}| - |\check{y} - \widetilde{y}| - |\check{x} - \widetilde{x}| . \tag{11}$$

Let Q_x be the hyperplane through $\check{x}$ and orthogonal to $[x, \check{x}]$. Then, using Lemma 25(2), we have

$$|\check{x} - \widetilde{x}| = |x - \check{x}| \cos \angle([x, \check{x}], \ell) = |x - \check{x}| \sin \angle(\ell, Q_x).$$

Since $T_{\check{x}}M \subseteq Q_x$, we have $\angle(\ell, Q_x) \leq \angle(\ell, T_{\check{x}}M)$, and so by Lemma 26(1),

$$|\check{x} - \widetilde{x}| \leq \frac{|\check{y} - \check{x}|}{2\,\mathrm{rch}(\check{x}, M)} |x - \check{x}| \leq \frac{a}{2R_{\mathrm{rch}}} |\check{y} - \check{x}| .$$

Likewise,

$$|\check{y} - \widetilde{y}| \leq \frac{a}{2R_{\mathrm{rch}}} |\check{y} - \check{x}| .$$

Thus (11) yields

$$|y - x| \geq \left(1 - \frac{a}{R_{\mathrm{rch}}}\right) |\check{y} - \check{x}| ,$$

and hence the result. $\square$

Lemma 37 (simplexwise distortion of $\boldsymbol{H}$) *Suppose $\boldsymbol{\sigma} \subset U_M$ is a simplex of dimension $\leq m$ whose vertices lie in M, and there is a ball $B_{\mathbb{R}^N}(c, r)$ such that for all $z \in \boldsymbol{\sigma}$, $\check{z} \in B = B_{\mathbb{R}^N}(c, r) \cap M$, where $\check{z} = \mathrm{pr}_M(z)$. Let R_{rch} be a lower bound on $\mathrm{rch}(\widetilde{z}, M)$ for all $\widetilde{z} \in B$. If $r < R_{\mathrm{rch}}$, and $L(\boldsymbol{\sigma}) < t(\boldsymbol{\sigma}) R_{\mathrm{rch}}/3$, then the restriction of pr_M to $\boldsymbol{\sigma}$ is a ξ-distortion map with*

$$\xi = \frac{12L^2}{t^2 R_{\mathrm{rch}}^2}.$$

Proof By Lemma 30,

$$\delta_M(x) < a = \frac{2L^2}{R_{\text{rch}}} \quad \text{for any } x \in \boldsymbol{\sigma}. \tag{12}$$

Thus it follows from Lemma 36 that

$$|\breve{y} - \breve{x}| \leq \left(1 - \frac{2L^2}{R_{\text{rch}}^2}\right)^{-1} |y - x| \leq \left(1 + \frac{4L^2}{R_{\text{rch}}^2}\right) |y - x| \,, \tag{13}$$

for any $x, y \in \boldsymbol{\sigma}$.

We now need to establish a lower bound on $|\breve{y} - \breve{x}|$. Let Q_x be the hyperplane through $\breve{x}$ and orthogonal to $[x, \breve{x}]$, and let $\widehat{y}$ and $\widehat{\breve{y}}$ be the orthogonal projection of y and $\breve{y}$ into Q_x. We have

$$|\breve{y} - \breve{x}| \geq \left|\widehat{\breve{y}} - \breve{x}\right| \geq |\widehat{y} - \breve{x}| - \left|\widehat{\breve{y}} - \widehat{y}\right|. \tag{14}$$

To get a lower bound on $|\widehat{y} - \breve{x}| = |y - x| \cos \angle([x, y], Q_x)$, notice that

$$\angle([x, y], Q_x) \leq \angle(\text{aff}(\boldsymbol{\sigma}), Q_x) \leq \angle(\text{aff}(\boldsymbol{\sigma}), T_{\breve{x}}M),$$

and by Lemma 32(2),

$$\sin \angle(\text{aff}(\boldsymbol{\sigma}), T_{\breve{x}}M) \leq \frac{3L}{tR_{\text{rch}}}.$$

Thus,

$$\cos \angle([x, y], Q_x) \geq \left(1 - \left(\frac{3L}{tR_{\text{rch}}}\right)^2\right)^{\frac{1}{2}} \geq \left(1 - \frac{9L^2}{t^2 R_{\text{rch}}^2}\right),$$

and so

$$|\widehat{y} - \breve{x}| \geq \left(1 - \frac{9L^2}{t^2 R_{\text{rch}}^2}\right) |y - x| \,. \tag{15}$$

To get an upper bound on the second term on the right side of (14), let Q_y be the hyperplane through $\breve{y}$ and orthogonal to $[\breve{y}, y]$. We have

$$\begin{aligned}
\left|\widehat{\breve{y}} - \widehat{y}\right| &= |\breve{y} - y| \cos \angle([\breve{y}, y], Q_x) \\
&= |\breve{y} - y| \sin \angle(Q_x, Q_y) \quad \text{by Lemma 25(2)} \\
&\leq |\breve{y} - y| \sin \angle(T_{\breve{x}}, T_{\breve{y}}) \\
&\leq |\breve{y} - y| \, \frac{|\breve{y} - \breve{x}|}{R_{\text{rch}}} \quad \text{by Lemma 27} \\
&\leq \frac{2L^2}{R_{\text{rch}}^2} \, |\breve{y} - \breve{x}| \quad \text{by (12).}
\end{aligned}$$

Now, using (13), we have

$$\left|\hat{\check{y}} - \hat{y}\right| \leq \frac{2L^2}{R_{\mathrm{rch}}^2}\left(1 + \frac{4L^2}{R_{\mathrm{rch}}^2}\right)|y - x| < \frac{3L^2}{R_{\mathrm{rch}}^2}|y - x|\,, \tag{16}$$

since the hypothesis $3L/R_{\mathrm{rch}} < t$ implies $4L^2/R_{\mathrm{rch}}^2 < \frac{1}{2}$.

Finally, plugging (15) and (16) back into (14), we get

$$|\check{y} - \check{x}| \geq \left(1 - \frac{12L^2}{t^2 R_{\mathrm{rch}}^2}\right)|y - x|\,.$$

Comparing this lower bound with (13), we arrive at the stated value for the metric distortion ξ. $\square$

4.5 Triangulation criteria for submanifolds

In order to employ Theorem 17 we first ensure that we meet the compatible atlases criteria (Definition 4). In Lemma 33 we defined

$$U_p = B_{\mathbb{R}^N}(p, r) \cap M, \quad \text{where } r = \rho R_{\mathrm{rch}} < \tfrac{1}{2} R_{\mathrm{rch}}.$$

We need to ensure that $H(|\underline{\mathrm{St}}(p)|) \subseteq U_p$. In our context, this means that we require $\mathrm{pr}_M(\iota(|\underline{\mathrm{St}}(p)|)) \subset U_p$, and Lemma 30 ensures that if $L(\boldsymbol{\sigma}) \leq L_0$ for all $\boldsymbol{\sigma} \in \iota(\underline{\mathrm{St}}(p))$, then it is sufficient to choose $r = L_0 + 2L_0^2/R_{\mathrm{rch}}$, or

$$\rho = \frac{L_0}{R_{\mathrm{rch}}}\left(1 + \frac{2L_0}{R_{\mathrm{rch}}}\right). \tag{17}$$

Our bound on L_0 itself will be much smaller than R_{rch}, which in turn will be expressed in terms of $\mathrm{lfs}(p)$ or $\mathrm{rch}(M)$, so we will have $r < \mathrm{lfs}(p)$, and thus $B_{\mathbb{R}^N}(p, r) \subset U_M$, ensuring also that each $\boldsymbol{\sigma} \in \iota(\underline{\mathrm{St}}(p))$ also lies in U_M.

We need to establish the metric distortion of $F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1}$ restricted to any m-simplex in $\widehat{\Phi}_p(\iota(\underline{\mathrm{St}}(p)))$, and ensure that it meets the distortion control criterion of Theorem 17.

Anticipating the bound we will need to meet the distortion-control criterion of Theorem 17, we impose the constraint

$$\frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2} \leq \frac{1}{16^2}, \tag{18}$$

where t_0 is a lower bound on the thickness: $t(\boldsymbol{\sigma}) \geq t_0$ for all $\boldsymbol{\sigma} \in \iota(\underline{\mathrm{St}}(p))$. We remark that L_0 and R_{rch} may be considered to be *local constants*, i.e., they may depend on the vertex $p \in \mathcal{P}$, however t_0 and the ratio L_0/R_{rch} will be global constants.

Using (18) together with Lemmas 35 and 19(1), we can bound the metric distortion of $\widehat{\Phi}_p^{-1}$ as

$$\xi_1 = \frac{L^2}{t^2 R_{\mathrm{rch}}^2}\left(1 - \frac{L^2}{t^2 R_{\mathrm{rch}}^2}\right)^{-1} \leq \left(\frac{16^2}{16^2 - 1}\right)\frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}.$$

Lemma 37 gives us the distortion of H:

$$\xi_2 = 12\frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}.$$

For ϕ_p, using Lemma 33 and (17) we get the distortion bound

$$\xi_3 = 4\rho^2 = \frac{4L_0^2}{R_{\mathrm{rch}}^2}\left(1 + \frac{2L_0}{R_{\mathrm{rch}}}\right)^2 \leq \frac{9}{2}\frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}.$$

Lemma 19(2) says that the distortion of F_p is no more than

$$\xi = \xi_1 + \xi_2 + \xi_3 + \xi_1\xi_2 + \xi_1\xi_3 + \xi_2\xi_3 + \xi_1\xi_2\xi_3.$$

Using (18) we find that F_p is a ξ-distortion map with

$$\xi = \frac{19L_0^2}{t_0^2 R_{\mathrm{rch}}^2}. \tag{19}$$

Observe that (18) implies that all the maps involved have distortion less than 1.

Now we need to ensure that this bound meets the distortion-bound requirement for Theorem 17. We have chosen to use here the properties of the Euclidean simplices in the ambient space $\mathbb{R}^N$, but in Theorem 17 we are considering simplices in the local coordinate space; for us these are the projected simplices, e.g., $\hat{\boldsymbol{\sigma}} = \widehat{\Phi}_p(\boldsymbol{\sigma}) = \mathrm{pr}_{T_pM}(\boldsymbol{\sigma})$. Using Lemma 35, the distortion properties of the affine map pr_{T_pM} imply

$$a(\hat{\boldsymbol{\sigma}}) \geq \left(1 - \frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}\right) a(\boldsymbol{\sigma}), \quad L(\hat{\boldsymbol{\sigma}}) \leq L(\boldsymbol{\sigma}),$$

and therefore

$$t(\hat{\boldsymbol{\sigma}}) \geq \left(1 - \frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}\right) t(\boldsymbol{\sigma}).$$

We can thus set

$$\hat{t}_0 = \left(1 - \frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}\right) t_0, \quad \hat{L}_0 = L_0, \quad \hat{s}_0 = \left(1 - \frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}\right) s_0,$$

where s_0 is a lower bound for the diameters of the simplices in $\iota(\underline{\mathrm{St}}(p))$.

Then, in order to meet the distortion control criterion of Theorem 17, we require

$$\frac{19L_0^2}{t_0^2 R_{\mathrm{rch}}^2} < \frac{\hat{s}_0 \hat{t}_0^2}{12\hat{L}_0} = \left(1 - \frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}\right)^3 \frac{s_0 t_0^2}{12L_0}. \tag{20}$$

It is convenient to define $\mu_0 = s_0/L_0$. Then, observing that

$$\left(1 - \frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}\right)^3 \geq \left(1 - 3\frac{L_0^2}{t_0^2 R_{\mathrm{rch}}^2}\right),$$

we see that (20) is satisfied if

$$L_0^2 \leq \frac{\mu_0 t_0^4 R_{\mathrm{rch}}^2}{16^2}. \tag{21}$$

Now we consider R_{rch}. For any $x \in U_p$, the Lipschitz continuity of lfs ensures that $\mathrm{lfs}(x) > \mathrm{lfs}(p) - \rho R_{\mathrm{rch}}$, where ρ is given by (17). Our constraint (21) on L_0 implies

$$\rho R_{\mathrm{rch}} \leq \frac{9}{8} L_0 \leq \frac{9}{2^7} R_{\mathrm{rch}}. \tag{22}$$

Thus we need $R_{\mathrm{rch}} \leq \mathrm{lfs}(p) - \frac{9}{128} R_{\mathrm{rch}}$, which is satisfied by

$$R_{\mathrm{rch}} = \frac{128}{137} \mathrm{lfs}(p). \tag{23}$$

Of course, we can also choose $R_{\mathrm{rch}} = \mathrm{rch}(M)$ independent of p. Plugging these values back into (21) gives us two alternatives for the bound on L_0:

$$L_0 \leq \frac{\mu_0^{\frac{1}{2}} t_0^2 \, \mathrm{lfs}(p)}{18} \quad \text{or} \quad L_0 \leq \frac{\mu_0^{\frac{1}{2}} t_0^2 \, \mathrm{rch}(M)}{16}. \tag{24}$$

We now only need to establish that F_p is simplexwise positive to arrive at our triangulation theorem for submanifolds.

Lemma 38 *F_p is simplexwise positive on* $\mathrm{pr}_{T_pM}(\iota(|\underline{\mathrm{St}}(p)|))$.

Proof Recall that $F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1}$, and observe that the restriction of F_p to an m-simplex $\hat{\boldsymbol{\sigma}} \in \mathrm{pr}_{T_pM}(\iota(\underline{\mathrm{St}}(p)))$ is differentiable. The bound (21) together with (19) implies F_p is a ξ-distortion map with $\xi < 1$, so Lemma 20 ensures that for any m-simplex $\boldsymbol{\sigma} \in \mathrm{pr}_{T_pM}(\iota(\underline{\mathrm{St}}(p)))$, the differential of F_p does not vanish on $\boldsymbol{\sigma}$.

We choose an orientation on U_p; thus we have an orientation on each tangent space T_xM, $x \in U_p$. The projection map $\mathrm{pr}_{T_pM}|_{T_xM} = d(\mathrm{pr}_{T_pM}|_M)_x = d(\phi_p)_x$ (see Remark 34) is then orientation preserving, because of continuity: it is certainly true when $x = p$, and Lemmas 33 and 20 imply that the differential $d(\phi_p)$ is nondegenerate on U_p. Thus ϕ_p is simplexwise positive.

For an m-simplex $\boldsymbol{\sigma} \in \iota(\underline{\mathrm{St}}(p))$ we *define* the orientation such that $\mathrm{pr}_{T_pM}|_{\boldsymbol{\sigma}}$ is positive. Equivalently, we define the orientation on $\iota(|\underline{\mathrm{St}}(p)|)$ to be such that $\widehat{\Phi}_p$ is simplexwise positive. This is not problematic, since we require that $\widehat{\Phi}_p$ be an embedding. Thus $\widehat{\Phi}_p$ is simplexwise positive by our definitions.

It remains to show that $\mathrm{pr}_M|_{\boldsymbol{\sigma}}$ is positive for each m-simplex $\boldsymbol{\sigma} \in \iota(\underline{\mathrm{St}}(p))$. First observe that for any $x \in \mathbb{R}^N$, the kernel of $d(\mathrm{pr}_M)_x$ is the subspace of $T_x\mathbb{R}^N$ corresponding to $N_{\check{x}}M$ (under the canonical identification $T_x\mathbb{R}^N \cong \mathbb{R}^N \cong T_{\check{x}}\mathbb{R}^N$, where $\check{x} = \mathrm{pr}_M(x)$). Also, observe that if $x \in M$, then $d(\mathrm{pr}_M|_{T_xM})_x = \mathrm{id}_{T_xM}$. This is easily seen by observing that $d(\mathrm{pr}_M|_{T_xM})_x = d(\mathrm{pr}_M)_x|_{T_xM}$ and using curves on M to apply the definition of the differential.

Thus at the central vertex $p \in \boldsymbol{\sigma}$, we can express $d(H|_{\boldsymbol{\sigma}})_p$ as

$$d(\mathrm{pr}_M|_{\boldsymbol{\sigma}})_p = \mathrm{id}_{T_pM} \circ \mathrm{pr}_{T_pM}|_{\boldsymbol{\sigma}}.$$

Since we have established that $\mathrm{pr}_{T_pM}|_{\boldsymbol{\sigma}}$ is positive, it follows that $d(\mathrm{pr}_M|_{\boldsymbol{\sigma}})_p$ is positive, and since $d(\mathrm{pr}_M|_{\boldsymbol{\sigma}})$ is nondegenerate on $\boldsymbol{\sigma}$ (because $d(F_p)$ is), it follows that $H|_{\iota(|\underline{\mathrm{St}}(p)|)}$ is simplexwise positive.

Thus F_p is orientation preserving on each simplex of $\underline{\mathrm{St}}(\hat{p})$. ∎

Theorem 39 (triangulation for submanifolds) *Let $M \subset \mathbb{R}^N$ be a compact C^2 manifold, and $\mathcal{P} \subset M$ a finite set of points such that for each connected component M_c of M, $M_c \cap \mathcal{P} \neq \emptyset$. Suppose that $\mathcal{A}$ is a simplicial complex whose vertices, $\mathcal{A}^0$, are identified with $\mathcal{P}$, by a bijection $\mathcal{A}^0 \to \mathcal{P}$ such that the resulting piecewise linear map $\iota\colon |\mathcal{A}| \to \mathbb{R}^N$ is an immersion, i.e., $\iota|_{|\underline{\mathrm{St}}(p)|}$ is an embedding for each vertex p.*

If:

(a) *For each vertex $p \in \mathcal{P}$, the projection* $\mathrm{pr}_{T_pM}|_{\iota(|\underline{\mathrm{St}}(p)|)}$ *is an embedding and p lies in the interior of* $\mathrm{pr}_{T_pM}(\iota(|\underline{\mathrm{St}}(p)|))$.

(b) *There are constants $0 < t_0 \leq 1$, $0 < \mu_0 \leq 1$, and $\epsilon_0 > 0$ such that for each simplex $\boldsymbol{\sigma} \in \iota(\mathcal{A})$, and each vertex $p \in \boldsymbol{\sigma}$,*

$$t(\boldsymbol{\sigma}) \geq t_0, \quad \mu_0 \epsilon_0 \,\mathrm{lfs}(p) \leq L(\boldsymbol{\sigma}) \leq \epsilon_0 \,\mathrm{lfs}(p), \quad \epsilon_0 \leq \frac{\mu_0^{\frac{1}{2}} t_0^2}{18}.$$

(c) *For any vertices $p, q \in \mathcal{P}$, if*

$$q \in U_p = B_{\mathbb{R}^N}(p, r) \cap M, \quad \textit{where } r = \frac{\mathrm{lfs}(p)}{15},$$

then $\mathrm{pr}_{T_pM}(q) \in \mathrm{pr}_{T_pM}(\iota(\underline{\mathrm{St}}(p)))$ *if and only if q is a vertex of* $\underline{\mathrm{St}}(p)$.

Then:

(1) *ι is an embedding, so the complex $\mathcal{A}$ may be identified with $\iota(\mathcal{A})$.*

(2) *The closest-point projection map* $\mathrm{pr}_M|_{|\mathcal{A}|}$ *is a homeomorphism $|\mathcal{A}| \to M$.*

(3) *For any $x \in \boldsymbol{\sigma} \in \mathcal{A}$,*

$$\delta_M(x) = |\check{x} - x| \leq \tfrac{7}{3}\epsilon_0^2 \,\mathrm{lfs}(\check{x}), \quad \sin \angle(\boldsymbol{\sigma}, T_{\check{x}}) \leq \frac{13\epsilon_0}{4t_0},$$

where $\check{x} = \mathrm{pr}_M(x)$.

Proof Observe that the local embedding condition (a) implies that $\mathcal{A}$ is a compact m-manifold without boundary. Condition (b) is a reformulation of the first inequality of (24). Condition (c) is the vertex sanity condition of Theorem 17; using (22) and (23) we get $r = 9\,\mathrm{lfs}(p)/137 < \mathrm{lfs}(p)/15$.

Thus, from our argument above, the criteria of Theorem 17 are satisfied, and $H = \mathrm{pr}_M \circ \iota$ is a homeomorphism. It follows that ι is injective, and since $|\mathcal{A}|$ is compact, ι must be an embedding. The second consequence, that $\mathrm{pr}_M|_{|\mathcal{A}|}$ is a homeomorphism, is now immediate.

For the third consequence, notice that, as argued to obtain (23), the Lipschitz continuity of lfs imples that

$$\mathrm{lfs}(p) \leq \frac{137}{128} \,\mathrm{lfs}(\check{x}).$$

So, using Lemma 30 and (23), we have

$$\delta_M(x) \le \frac{2L^2}{R_{\mathrm{rch}}} \le \frac{2\epsilon_0^2 \operatorname{lfs}(p)^2}{R_{\mathrm{rch}}} \le 2\left(\frac{137}{128}\right)^2 \epsilon_0^2 \operatorname{lfs}(\check{x}) \le \tfrac{7}{3}\epsilon_0^2 \operatorname{lfs}(\check{x}).$$

The second inequality follows from Lemma 32(2) and (23). $\square$

If we use the global bound $R_{\mathrm{rch}} = \operatorname{rch}(M)$, to bound the size of the simplices, then the third consequence of Theorem 39 can be tightened. In this context we obtain the following variation of Theorem 39, which is a corollary in the sense that it follows from essentially the same proof, even though it does not follow from the statement of Theorem 39.

Corollary 40 *If the conditions* (b) *and* (c) *in Theorem 39 are replaced by*

(b′) *There are constants* $0 < t_0 \le 1$, $0 < \mu_0 \le 1$, *and* $\epsilon_0 > 0$ *such that for each simplex* $\boldsymbol{\sigma} \in \iota(\mathcal{A})$, *and each vertex* $p \in \boldsymbol{\sigma}$,

$$t(\boldsymbol{\sigma}) \ge t_0, \quad \mu_0\epsilon_0 \operatorname{rch}(M) \le L(\boldsymbol{\sigma}) \le \epsilon_0 \operatorname{rch}(M), \quad \epsilon_0 \le \frac{\mu_0^{\frac{1}{2}} t_0^2}{16}.$$

(c′) *For any vertices* $p, q \in \mathcal{P}$, *if* $q \in U_p = B_{\mathbb{R}^N}(p, r) \cap M$, *where* $r = \operatorname{rch}(M)/14$, *then* $\operatorname{pr}_{T_pM}(q) \in \operatorname{pr}_{T_pM}(\iota(\underline{\mathrm{St}}(p)))$ *if and only if* q *is a vertex of* $\underline{\mathrm{St}}(p)$.

then the conclusions of Theorem 39 hold, and consequence (3) can be tightened to:

(3′) *For any* $x \in \boldsymbol{\sigma} \in \mathcal{A}$,

$$\delta_M(x) = |\check{x} - x| \le 2\epsilon_0^2 \operatorname{rch}(M), \quad \sin\angle(\boldsymbol{\sigma}, T_{\check{x}}) \le \frac{3\epsilon_0}{t_0},$$

where $\check{x} = \operatorname{pr}_M(x)$.

Proof This follows from the proof of Theorem 39, using $R_{\mathrm{rch}} = \operatorname{rch}(M)$ instead of (23), e.g., from (21) we obtain the second alternative in (24), and from (22) we get $r = 9\operatorname{rch}(M)/128 < \operatorname{rch}(M)/14$. $\square$

A Elementary degree theory

We recall here some basic ideas in degree theory. Our primary motivation is to facilitate the statement and proof of Lemma 7, recovering what we need of a result of Whitney [Whi57, Appendix II Lemma 15a], without using differentiability assumptions.

We are interested in the degree of continuous maps $F\colon \overline{\Omega} \to \mathbb{R}^m$, where $\Omega \subset \mathbb{R}^m$ is an open, bounded, and nonempty domain in $\mathbb{R}^m$, as can be found Chapter IV, Sections 1 and 2 of [OR09], for example. As with most modern treatments of degree theory, Outerelo and Ruiz start by defining the degree for a regular value of a differentiable map, where the idea is transparent. For $y \in \mathbb{R}^m \setminus F(\partial\Omega)$, the degree is defined as

$$\deg(F, \Omega, y) := \sum_{x \in F^{-1}(y)} \operatorname{sgn}_x(F),$$

where $\operatorname{sgn}_x(F)$ denotes the sign (± 1) of the Jacobian determinant (the determinant of the differential of F) at x. Thus the degree at y counts the number of points in the preimage, accounting for the local orientation of the map.

It is then shown that the degree is locally constant on the (open) set of regular values, and after showing that it is also invariant under homotopies that avoid conflicts between y and the image of the boundary, the degree is defined for an arbitrary point in $\mathbb{R}^m \setminus F(\partial\Omega)$, and it is constant on each connected component of $\mathbb{R}^m \setminus F(\partial\Omega)$.

Then the definition of degree is extended to continuous maps $F\colon \overline{\Omega} \to \mathbb{R}^m$ by means of the Weierstrass approximation theorem (here $\|\cdot\|_\infty$ denotes the supremum norm):

Lemma 41 ([OR09, Proposition and Definition IV.2.1]) *Let $F\colon \overline{\Omega} \to \mathbb{R}^m$ be a continuous map, and let $y \in \mathbb{R}^m \setminus F(\partial\Omega)$. Then there exists a smooth map $G : \overline{\Omega} \to \mathbb{R}^m$ such that $\|F - G\|_\infty < d_{\mathbb{R}^m}(y, F(\partial\Omega))$. For all such G, the degree $\deg(G, \Omega, y)$ is defined ($y \in \mathbb{R}^m \setminus G(\partial\Omega)$) and is the same, and we define the degree of F by*

$$\deg(F, \Omega, y) = \deg(G, \Omega, y).$$

Furthermore, G can be chosen such that y is a regular value of $G|_\Omega$, and then

$$\deg(F, \Omega, a) = \sum_{x \in G^{-1}(y)} \operatorname{sgn}_x(G).$$

The locally constant nature of the degree is the main property we wish to exploit:

Lemma 42 ([OR09, Proposition IV.2.3]) *Let $F: \overline{\Omega} \to \mathbb{R}^m$ be a continuous map. Then the degree $y \mapsto \deg(F, \Omega, y)$ is constant on every connected component of $\mathbb{R}^m \setminus F(\partial\Omega)$.*

Notice that if F is a topological embedding, then $\deg(F, \Omega, y) = \pm 1$ for any point $y \in F(\Omega)$. We will also have occasion to use the following:

Lemma 43 ([OR09, Corollary IV.2.5(3)]) *Given a continuous mapping $F: \overline{\Omega} \to \mathbb{R}^m$, two disjoint open subsets $\Omega_1, \Omega_2 \subset \Omega$, and a point $y \notin F(\overline{\Omega} \setminus (\Omega_1 \cap \Omega_2))$,*

$$\deg(F, \Omega, y) = \deg(F, \Omega_1, y) + \deg(F, \Omega_2, y).$$

Since no connectedness assumptions are made on the open sets in question, a straightforward inductive argument allows us to strengthen the statement of Lemma 43:

Lemma 44 *Suppose $\Omega_1, \ldots, \Omega_n$ are mutually disjoint open subsets of the open domain $\Omega \subset \mathbb{R}^m$, and $F: \overline{\Omega} \to \mathbb{R}^m$ is a continuous map. If $y \notin F(\overline{\Omega} \setminus \bigcup_{i=1}^n \Omega_i)$, then*

$$\deg(F, \Omega, y) = \sum_{i=1}^{n} \deg(F, \Omega_i, y).$$

A.1 Orientation and cogent maps

A simplex is *oriented* by choosing an orientation for its affine hull, or equivalently, by ordering its vertices; any even permutation of this order describes the same orientation. An m-simplex in $\mathbb{R}^m$ has a natural orientation induced from the canonical orientation of $\mathbb{R}^m$ defined by the standard basis. Our convention is that $\boldsymbol{\sigma} \subset \mathbb{R}^m$ is positively oriented if its vertices v_i, $0 \leq i \leq m$, are ordered such that the basis $\{v_1 - v_0, \ldots, v_m - v_0\}$ defines the same orientation as the canonical basis of $\mathbb{R}^m$.

In the case that concerns us, where $\mathcal{C}$ is a finite pure m-complex piecewise linearly embedded in $\mathbb{R}^m$ (so that we can naturally view $|\mathcal{C}| \subset \mathbb{R}^m$), we assume that the m-simplices carry the canonical orientation inherited from the ambient space.

Definition 45 (orientation preserving map) If $\boldsymbol{\sigma} \subset \mathbb{R}^m$ is an m-simplex, we say that a continuous topological embedding $F\colon \boldsymbol{\sigma} \to \mathbb{R}^m$ is *orientation preserving*, or *positive*, if $\deg(F, \operatorname{relint}(\boldsymbol{\sigma}), y) = 1$, for any point $y \in \operatorname{int}(F(\boldsymbol{\sigma}))$, otherwise F is *orientation reversing*.

Definition 46 (cogent maps) Suppose $\mathcal{C}$ is a pure m-complex. We call a map $F\colon |\mathcal{C}| \to \mathbb{R}^m$ *cogent* with respect to $\mathcal{C}$ if it is continuous and its restriction to each simplex is an embedding.

For cogent maps, the points in $\mathbb{R}^m \setminus F(|\mathcal{C}^{m-1}|)$ are analogous to the regular values of a differentiable map. If F is a cogent map and $x \in \operatorname{relint}(\boldsymbol{\sigma})$, where $\boldsymbol{\sigma}$ is an m-simplex in $\mathcal{C}$, we define

$$\operatorname{sgn}_x(F) := \deg(F, \operatorname{relint}(\boldsymbol{\sigma}), F(x)). \tag{25}$$

This makes the analogy with the degree of a differentiable map transparent:

Lemma 47 (degree of cogent maps) *Suppose $\mathcal{C}$ is a finite pure m-complex embedded in $\mathbb{R}^m$. Let $\Omega = |\mathcal{C}| \setminus \partial|\mathcal{C}|$. If a continuous map $F\colon |\mathcal{C}| \to \mathbb{R}^m$ is cogent with respect to $\mathcal{C}$, then for any $y \in \mathbb{R}^m \setminus F(|\mathcal{C}^{m-1}|)$* (recall that $\partial|\mathcal{C}| \subseteq |\mathcal{C}^{m-1}|$ [BDG13, Lemmas 3.6, 3.7]),

$$\deg(F, \Omega, y) = \sum_{x \in F^{-1}(y)} \operatorname{sgn}_x(F).$$

A very minor modification to the proof of Lemma 41 would allow us to sharpen the bound on $\|F - G\|_\infty$ in our context, so that

$$\|F - G\|_\infty < \min_{\boldsymbol{\sigma} \cap F^{-1}(y) \neq \emptyset} d_{\mathbb{R}^m}(y, F(\partial \boldsymbol{\sigma})),$$

and Lemma 47 follows immediately from the same proof. But rather than digging into the proof of that lemma, we can avoid getting our hands dirty and just exploit Lemma 44.

Proof of Lemma 47 The preimage of y is a finite set of points: $F^{-1}(y) = \{x_i\}$, $i \in \{1, \dots, n\}$. For each i, let $\boldsymbol{\sigma}_i$ be the m-simplex that contains x_i in its interior, and set $\Omega_i = \operatorname{relint}(\boldsymbol{\sigma}_i)$. Then the statement follows from Lemma 44 and the definition (25) of $\operatorname{sgn}_{x_i}(F)$. $\square$

Definition 48 (simplexwise positive) A map is *simplexwise positive* if it is cogent and its restriction to any m-simplex is orientation preserving.

Lemmas 42 and 47 yield:

Lemma 49 *If $\mathcal{C}$ is a finite pure m-complex embedded in $\mathbb{R}^m$, and $F\colon |\mathcal{C}| \to \mathbb{R}^m$ is simplexwise positive, then for any connected open subset W of $\mathbb{R}^m \setminus F(\partial|\mathcal{C}|)$, any two points of W not in $F(|\mathcal{C}^{m-1}|)$ are covered the same number of times (i.e., have the same number of points in their preimage under F).*

Remark 50 Whitney's Lemma [Whi57, Appendix II Lemma 15] is a combination of Lemmas 49 and 7, but applies in more generality. However, he assumes that the restriction of F to each m-simplex is a smooth map. The generality can be fully recovered without invoking this differentiability assumption.

The definition of the degree of a map $\overline{\Omega} \to \mathbb{R}^m$ can be naturally extended to the case where $\overline{\Omega}$ is an oriented abstract manifold with boundary. Using this, the assumption that $\mathcal{C}$ is embedded in $\mathbb{R}^m$ can be dropped. Also, the assumption that $\mathcal{C}$ be finite can be relaxed, at least if we assume that F is proper (so that the number of points in the preimage of a point is finite — Whitney seems to assume this).

Whitney also only assumed that $\mathcal{C}$ was a *pseudomanifold* with boundary: a pure m-complex such that any $(m-1)$-simplex is a face of either 2 or 1 m-simplices (those that are the face of only 1 m-simplex define the boundary complex).

Brouwer's original exposition of degree theory [Bro12] used simplicial approximations, and piecewise linear maps, rather than differentiable maps as the foundation. Even though the simplicial aspect is attractive and natural in our setting, there would be no economy in using this approach for our purposes. However, Brouwer's exposition was based on the notion of pseudomanifolds [OR09, p. 23], so that approach would also allow us to recover the pseudomanifold aspect of Whitney's lemma.

B On tangent space variation

The purpose of this appendix is to sketch the demonstration of Lemma 27, which is adapted from arguments presented in [BLW17b]. The argument relies on this convexity result:

Lemma 51 (convexity) *Suppose $B_{\mathbb{R}^N}(c, r) \subset U_M$ is such that $\mathrm{rch}(x, M) \geq R_{\mathrm{rch}}$ for all $x \in B = B_{\mathbb{R}^N}(c, r) \cap M$. If $r < R_{\mathrm{rch}}$, then B is geodesically convex* *in the sense that for any $x, y \in B$, any minimizing geodesic between x and y is contained in B.*

Proof This follows from the same argument that produced [BLW17b, Theorem 3.6], but using R_{rch} instead of $\mathrm{rch}(M)$. Here is an overview of the adjustments that must be done to the argument:

[BLW17b, Corollary 2.3] holds with $\mathrm{rch}(M)$ replaced by $\mathrm{rch}(p, M)$.

[BLW17b, Lemma 3.1] holds with $\mathrm{rch}(M)$ replaced with R_{rch}, where R_{rch} is a lower bound on $\mathrm{rch}(\gamma(t), M)$, for all relevant t.

[BLW17b, Lemma 3.2] also holds more generally. In fact the essential argument has nothing to do with reach or manifolds, it says this:

If $\alpha\colon I \to \mathbb{R}^N$ is parameterized by arc length, and has curvature bounded by $1/R$, i.e., $|\alpha''(t)| \leq 1/R$ for all $t \in I$, then for any $a, b \in I$,

$$\angle(\alpha'(a), \alpha'(b)) \leq \frac{\mathrm{len}(\alpha([a, b]))}{R}. \tag{26}$$

This comes directly from using the bound on the curvature to give a bound on the length of the "indicatrix of tangents", i.e., the curve traced out on the sphere by the unit tangent vectors.

This allows us to bound the distance between the endpoints of $\alpha(I)$ for sufficiently small I. Let $I = [0, \ell]$, and $\alpha(0) = a$ and $\alpha(\ell) = b$. Then, letting $v = \alpha'(\ell/2)$ and integrating $\langle b - a, v\rangle = \int_0^\ell \langle \alpha'(s), v\rangle \, ds$ in two parts, we get

$$|b - a| \geq 2R \sin\left(\frac{\ell}{2R}\right), \quad \text{assuming } \ell \leq \pi R. \tag{27}$$

(Actually, it seems the argument is okay as long as $\ell \leq 2\pi R$, but after ℓ becomes larger than πR the bound becomes smaller, finally vanishing when we've come full circle.)

[BLW17b, Lemma 3.3]: Again this is really just an argument about curvature controlled curves. The stated bound holds if γ is any space curve, and $\mathrm{rch}(M)$ is replaced with R, where $|\gamma''| \leq R$, and $\mathrm{len}(\gamma) < \pi R$.

[BLW17b, Lemma 3.4]: We can replace this statement with: If p and q are connected by a minimizing geodesic γ, and $|p - q| < 2R_{\mathrm{rch}}$, where R_{rch} is a lower bound on the local reach along γ, then $d_M(p, q) < \pi R_{\mathrm{rch}}$.

The lens-shaped region described in [BLW17b, Corollary 3.5] is of course constructed using the short arc of a circle of radius R_{rch}, where R_{rch} is a lower bound on the local reach along the geodesic.

Now the convexity argument for [BLW17b, Corollary 3.5] goes through when we replace the ball of radius less than reach with the ball $B_{\mathbb{R}^N}(c,r) \subset U_M$ such that $\mathrm{rch}(x,M) \geq R_{\mathrm{rch}}$ for all $x \in B = B_{\mathbb{R}^N}(c,r) \cap M$, and $r < R_{\mathrm{rch}}$. Indeed, note that the condition $B_{\mathbb{R}^N}(c,r) \subset U_M$ ensures that $B_{\mathbb{R}^N}(c,r)$ can only intersect a single connected component of M (since the medial axis separates topological components). $\square$

Using the kind of argument that leads to Equation (26) we find: If $\mathrm{rch}(x,M) \leq R_{\mathrm{rch}}$ for all x lying on a minimizing geodesic between a and b on M, then

$$\angle(T_aM, T_bM) \leq \frac{d_M(a,b)}{R_{\mathrm{rch}}}. \tag{28}$$

This is found by using the curvature bound provided by R_{rch} to bound the angle between any vector $u \in T_aM$ and $v \in T_bM$, obtained by parallel transport of u along the geodesic to b. This bound does not require a bound on the distance between a and b, but it becomes vacuous if $d_M(a,b) \geq \pi R_{\mathrm{rch}}/2$.

Using (27), we get

$$\sin\left(\tfrac{1}{2}\angle(T_aM, T_bM)\right) \leq \frac{|a-b|}{2R_{\mathrm{rch}}}, \tag{29}$$

when R_{rch} is a lower bound on the local reach along the geodesic, and the length of a minimizing geodesic doesn't exceed πR_{rch}. As discussed in the "proof" of Lemma 51, the argument of [BLW17b, Lemma 3.4] implies that if a and b are connected by a minimizing geodesic γ, and $|a-b| < 2R_{\mathrm{rch}}$, where R_{rch} is a lower bound on the local reach along γ, then $d_M(a,b) < \pi R_{\mathrm{rch}}$.

Lemma 52 *Suppose $a, b \in M$ are connected by a minimizing geodesic γ, and $\mathrm{rch}(x,M) \leq R$ for all x along γ. Then*

$$\sin\angle(T_aM, T_bM) \leq \frac{|a-b|}{R}.$$

If $|a-b| \leq 2R$, then

$$\angle(T_aM, T_bM) \leq \frac{\pi\,|a-b|}{2R}.$$

Proof The first inequality follows from (29) and the observation (from the angle sum formula) that $\sin(2\theta) \leq 2\sin\theta$.

The second claim follows from the observation above that $|a-b| \leq 2R_{\mathrm{rch}}$ implies $d_M(a,b) \leq \pi R_{\mathrm{rch}}$, and so (28) shows that $\theta = \frac{1}{2}\angle(T_aM, T_bM) \leq \pi/2$. We then use the observation that in this case $\frac{2\theta}{\pi} \leq \sin\theta$. $\square$

Combining Lemmas 51 and 52, we obtain Lemma 27.

C Exploiting strong differential bounds

The triangulation criteria of Theorem 17 presented in Section 2 are based on the triangulation demonstration presented in [DVW15, Proposition 16]. The main motivation for presenting the new argument in Section 2 is that the methods of [DVW15] require an intricate analysis of the differential of the map F_p, which makes the application of the result, considerably more difficult than meeting the purely metric criteria of Theorem 17.

The motivation for reviewing the previous method here is that the demonstration of the triangulation result [DVW15, Proposition 16] was incorrect, and in fact the statement of the proposition does not ensure the injectivity of the map H. We correct the problem here by employing the vertex sanity assumption (Definition 15) introduced in Section 2.2, and provide an erratum in Section C.1. Although the criteria for this method are more difficult to establish, once they are established, a stronger result is obtained, as mentioned in Remark 58, so these results may still be of interest.

For the method of Theorem 17, if F_p is differentiable on each m-simplex, as is the case in the setting of Section 4, then Lemma 21 says that to show that F_p is a ξ-distortion map, it is sufficient to show that for any vector w tangent to a point u in the domain of F_p,

$$(1-\xi)\,|w| \leq |d(F_p)_u w| \leq (1+\xi)\,|w|\,. \tag{30}$$

In the analogous result, [DVW15, Proposition 16], a stronger bound on the differential is demanded: we require that

$$\|d(F_p)_u - \mathrm{id}\| \leq \xi \tag{31}$$

for all u in the domain. The bound (31) is strictly stronger than (30). It is not difficult to establish that (31) implies that F_p is a ξ-distortion map [DVW15, Lemma 11] on each simplex. However, whereas (30) only constrains how much dF_p can change the magnitude of a vector, (31) also constrains how much the direction can change. For this kind of bound on the differential, there is no need to exploit the trilateration lemma (Lemma 10):

Lemma 53 ([DVW15, Lemma 12]) *Suppose $\boldsymbol{\omega} \subseteq \mathbb{R}^m$ is a convex set and $F\colon \boldsymbol{\omega} \to \mathbb{R}^m$ is a smooth map with a fixed point $p \in \boldsymbol{\omega}$. If*

$$\|dF_x - \mathrm{Id}\| \leq \xi \quad \textit{for all } x \in \boldsymbol{\omega},$$

then

$$|F(x) - x| \leq \xi\,|x - p| \quad \textit{for all } x \in \boldsymbol{\omega}.$$

The added control obtained by the strong bound (31) on the differential enables us to ensure that F_p embeds the boundary of $\underline{\mathrm{St}}(p)$. This in turn implies that F_p embeds all of $\underline{\mathrm{St}}(p)$. This follows from the following corrollary to Whitney's lemma [Whi57, Lemma AII.15a] (which is demonstrated as Lemmas 49 and 7 in this work). We include the proof here since the proof in [DVW15] erroneously states that Whitney's proof implies that a simplexwise positive map is a local homeomorphism; this is not true, but such a map is an open map.

Definition 54 (smooth on $\mathcal{C}$) Given a simplicial complex $\mathcal{C}$, we say that a map $F\colon |\mathcal{C}| \to \mathbb{R}^m$ is *smooth on* $\mathcal{C}$ if for each $\boldsymbol{\sigma} \in \mathcal{C}$ the restriction $F\big|_{\boldsymbol{\sigma}}$ is smooth.

Lemma 55 ([DVW15, Lemma 13]) *Let $\mathcal{C}$ be a (finite) simplicial complex embedded in $\mathbb{R}^m$ such that* $\mathrm{int}(|\mathcal{C}|)$ *is nonempty and connected, and $\partial|\mathcal{C}|$ is a compact, connected $(m-1)$-manifold. Suppose $F\colon |\mathcal{C}| \to \mathbb{R}^m$ is smooth on $\mathcal{C}$ and simplexwise positive. If the restriction of F to $\partial|\mathcal{C}|$ is an embedding, then F is a topological embedding.*

Proof The assumptions on $\mathrm{int}(|\mathcal{C}|)$ and $\partial|\mathcal{C}|$ imply that $\mathcal{C}$ is a pure m-complex, and that each $(m-1)$-simplex is either a boundary simplex, or the face of exactly two m-simplices. (This is a nontrivial exercise, and requires the demand that $\partial|\mathcal{C}|$ be connected, which was absent in the original statement of the lemma.)

By the same argument as in the proof of Lemma 7, $F(\mathrm{int}(|\mathcal{C}|))$ is open. By the Jordan-Brouwer separation theorem [OR09, §IV.7], $\mathbb{R}^n \setminus F(\partial|\mathcal{C}|)$ consists of two open components, one of which is bounded. Since $F(|\mathcal{C}|)$ is compact, Lemma 49 implies that $F(\mathrm{int}(|\mathcal{C}|))$ must coincide with the bounded component, and in particular $F(\mathrm{int}(|\mathcal{C}|)) \cap F(\partial|\mathcal{C}|) = \emptyset$, so $F(\mathrm{int}(|\mathcal{C}|))$ is a single connected component.

We need to show that F is injective. First we observe that the set of points in $F(\mathrm{int}(|\mathcal{C}|))$ that have exactly one point in the preimage is nonempty.

It suffices to look in a neigbhourhood of a point $y \in F(\partial|\mathcal{C}|)$. Choose $y = F(x)$, where x is in the relative interior of $\boldsymbol{\sigma}^{m-1} \subset \partial|\mathcal{C}|$. Then there is a neighbourhood V of y such that V does not intersect the image of any other simplex of dimension less than or equal to $n-1$. Let $\boldsymbol{\sigma}^m$ be the unique m-simplex that has σ^{m-1} as a face. Then $F^{-1}(V \cap F(|\mathcal{C}|) \subset \boldsymbol{\sigma}^m$, and it follows that every point in $V \cap \mathrm{int}(|\mathcal{C}|)$ has a unique point in its image.

Now the injectivity of F follows from Lemma 7. □

We then obtain bounds on the differential of F_p sufficient to ensure that it embeds $\underline{\mathrm{St}}(p)$:

Lemma 56 ([DVW15, Lemma 14]) *Suppose $\mathcal{C} = \underline{\mathrm{St}}(\hat{p})$ is a t_0-thick, pure m-complex embedded in $\mathbb{R}^m$ such that all of the m-simplices are incident to a single vertex, $\hat{p}$, and $\hat{p} \in \mathrm{int}(|\mathcal{C}|)$* (i.e., $\underline{\mathrm{St}}(\hat{p})$ is a full star). *If $F \colon |\mathcal{C}| \to \mathbb{R}^m$ is smooth on $\mathcal{C}$, and satisfies*

$$\|dF - \mathrm{Id}\| < mt_0 \tag{32}$$

on each m-simplex of $\mathcal{C}$, then F is an embedding.

If the requirements of Lemma 56 are met, then H is a local homemorphism, and we are left with ensuring that it is injective, in order to guarantee that it is a homeomorphism. To this end we employ the vertex sanity assumption (Definition 15), and we arrive at the following triangulation result, which can replace the flawed [DVW15, Proposition 16]:

Proposition 57 (triangulation with differential control) *Suppose $\mathcal{A}$ is a compact m-manifold complex (without boundary), with vertex set $\mathcal{P}$, and M is an m-manifold. A map $H \colon |\mathcal{A}| \to M$ is a homeomorphism if the following criteria are satisfied:*

(1) **compatible atlases** *There are compatible atlases*

$$\{(\widetilde{\mathcal{C}}_p, \widehat{\Phi}_p)\}_{p\in\mathcal{P}}, \quad \widetilde{\mathcal{C}}_p \subset \mathcal{A}, \quad \textit{and} \quad \{(U_p, \phi_p)\}_{p\in\mathcal{P}}, \quad U_p \subset M,$$

for H, with $\widetilde{\mathcal{C}}_p = \underline{\mathrm{St}}(p)$ for each $p \in \mathcal{P}$, the vertex set of $\mathcal{A}$ (Definition 4).

(2) **simplex quality** *For each $p \in \mathcal{P}$, every simplex $\boldsymbol{\sigma} \in \underline{\mathrm{St}}(\hat{p}) = \widehat{\Phi}_p(\underline{\mathrm{St}}(p))$ satisfies $s_0 \leq L(\boldsymbol{\sigma}) \leq L_0$ and $t(\boldsymbol{\sigma}) \geq t_0$* (Notation 9).

(3) **distortion control** *For each $p \in \mathcal{P}$, the map*

$$F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1} \colon |\underline{\mathrm{St}}(\hat{p})| \to \mathbb{R}^m,$$

is smooth on $\underline{\mathrm{St}}(\hat{p})$, and simplexwise positive, and for any m-simplex $\boldsymbol{\sigma} \in \underline{\mathrm{St}}(\hat{p})$ and any $u \in \boldsymbol{\sigma}$, we have

$$\|d(F_p)_u - \mathrm{id}\| < \xi = \frac{s_0 t_0}{2L_0}$$

(Definitions 54 and 48).

(4) **vertex sanity** *For all vertices $p, q \in \mathcal{P}$, if $\phi_p \circ H(q) \in |\underline{\mathrm{St}}(\hat{p})|$, then q is a vertex of $\underline{\mathrm{St}}(p)$.*

Proof Since the requirements of Lemma 56 are met, we only need to show that H is injective. The argument the same as for Lemma 16, with minor modifications.

Towards a contradiction, suppose that $H(q) \in H(\boldsymbol{\sigma})$ and that q is not a vertex of the m-simplex $\boldsymbol{\sigma}$. This means there is some $x \in \boldsymbol{\sigma}$ such that $H(x) = H(q)$. Let p be a vertex of $\boldsymbol{\sigma}$. The vertex sanity hypothesis (Definition 15) implies that $\phi_p \circ H(q)$ must be either outside of $|\underline{\mathrm{St}}(\hat{p})|$, or belong to its boundary. Thus Lemmas 12 and 53, and the bound on ξ imply that the barycentric coordinate of x with respect to p must be smaller than $\frac{1}{m+1}$: Let $\hat{x} = \widehat{\Phi}_p(x)$, and $\hat{\boldsymbol{\sigma}} = \widehat{\Phi}_p(\boldsymbol{\sigma})$. Lemma 53 says that

$$|F_p(\hat{x}) - \hat{x}| < \xi L_0 = \frac{s_0 t_0}{2} \leq \frac{a_0}{2m},$$

where a_0 is a lower bound on the altitudes of $\hat{p}$, as in Lemma 12. Since $F_p(\hat{x}) = \phi_p \circ H(x) = \phi_p \circ H(q)$ is at least as far away from $\hat{x}$ as $\partial \underline{\mathrm{St}}(\hat{p})$, Lemma 12 implies that the barycentric coordinate of $\hat{x} \in \hat{\boldsymbol{\sigma}}$ with respect to $\hat{p}$ must be strictly less than $\frac{1}{2m} \leq \frac{1}{m+1}$. Since $\widehat{\Phi}_p$ preserves barycentric coordinates, and the argument works for any vertex p of $\boldsymbol{\sigma}$, we conclude that all the barycentric coordinates of x in $\boldsymbol{\sigma}$ are strictly less than $\frac{1}{m+1}$. We have reached a contradiction with the fact that the barycentric coordinates of x must sum to 1. □

Remark 58 Notice that the constraint on ξ in Proposition 57 is only linear in t_0, whereas it is quadratic in t_0 in Theorem 17.

C.1 Erratum for *Riemannian simplices and triangulations*

This subsection is an erratum for [DVW15], and we will employ here the notation and conventions of that paper, which differ slightly from those in the rest of the current document. In particular, the dimension of the Riemannian manifold we are triangulating is n, and σ denotes an abstract simplex, i.e., a set of vertices, usually on the manifold, M. When these vertices are lifted via the inverse of the exponential map to T_pM, the resulting vertex set is denoted $\sigma(p)$. A "filled in" Euclidean simplex is denoted $\boldsymbol{\sigma}_{\mathbb{E}}$. The vertex set of $\mathcal{A}$ is denoted by S instead of $\mathcal{P}$.

The proof of the generic triangulation criteria, Proposition 16 in [DVW15], is flawed; the criteria presented do not guarantee that the map H is injective. This problem infects the results in [DVW15] which rely on Proposition 16: Theorem 2, Proposition 26, and Theorem 3.

These results all hold true without any further modifications if the following hypothesis is added to Proposition 16, and Theorem 2 (the other results simply require that the hypotheses of Theorem 2 are met):

Hypothesis 59 (simple injectivity assumption) If q is a vertex of $\mathcal{A}$ and $q \in H(\boldsymbol{\sigma}_{\mathbb{E}})$, then q is a vertex of $\boldsymbol{\sigma}$.

The injectivity of the map H follows trivially from Hypothesis 59, since it has been established that H is a covering map. However, this assumption is not easy to verify, at least not in some applications of interest, e.g., [BDG17]. We provide here corrected statements of the affected results, obtained by replacing [DVW15, Proposition 16] with Proposition 57, which uses the vertex sanity assumption, Definition 15. Unfortunately, the bound required on the differential in prop:strong.bnd.triang(3) is qualitatively different from that imposed in [DVW15, Proposition 16]. In particular, we need to impose a lower bound s_0 on the diameters of the simplices. It is convenient to define $\mu_0 = s_0/L_0$.

In the proof of [DVW15, Theorem 2], the Proposition 16 is employed at the bottom of page 23, just before the statement of the theorem. Reworking that short calculation, and incorporating the vertex sanity hypothesis, we obtain:

Theorem 60 ([DVW15, Theorem 2] corrected) *Suppose M is a compact n-dimensional Riemannian manifold with sectional curvatures K bounded* *by $|K| \leq \Lambda$, and $\mathcal{A}$ is an abstract simplicial complex with finite vertex set $S \subset M$. Define quality parameters $t_0 > 0$, $0 < \mu_0 \leq 1$, and let*

$$h = \min\left\{\frac{\iota_M}{4}, \frac{\sqrt{\mu_0}t_0}{6\sqrt{\Lambda}}\right\}. \tag{33}$$

Suppose

(1) *For every $p \in S$, the vertices of $\underline{\mathrm{St}}(p)$ are contained in $B_M(p,h)$, and the balls $\{B_M(p,h)\}_{p\in S}$ cover M.*

(2) *For every $p \in S$, the restriction of the inverse of the exponential map $\exp_p^{-1}$ to the vertices of $\underline{\mathrm{St}}(p) \subset \mathcal{A}$ defines a piecewise linear embedding of $|\underline{\mathrm{St}}(p)|$ into T_pM, realising $\underline{\mathrm{St}}(p)$ as a full star, $\widehat{\underline{\mathrm{St}}(p)}$, such that every simplex $\sigma(p)$ has thickness $t(\sigma(p)) \geq t_0$ and diameter $\mu_0 L_0 \leq L(\sigma(p)) \leq L_0$.*

(3) *For all vertices $p, q \in S$, if $(\exp_p|_{B_M(p,h)})^{-1}(q) \in |\widehat{\underline{\mathrm{St}}(p)}|$, then q is a vertex of $\underline{\mathrm{St}}(p)$.*

Then $\mathcal{A}$ triangulates M, and the triangulation is given by the barycentric coordinate map on each simplex.

For [DVW15, Proposition 26], the affected argument is in the paragraph preceding the statement of the proposition. We replace that paragraph with (where it is understood that references to Theorem 2 are to the corrected version, Theorem 60). The result is that the new (stronger) bounds on h are always sufficient to ensure a piecewise flat metric on $\mathcal{A}$:

> Thus in order to guarantee that the ℓ_{ij} describe a non-degenerate Euclidean simplex, we require that $\Lambda h^2 = \eta t_0^2/2$, for some non-negative $\eta < 1$.
>
> Under the conditions of Theorem 2 we may have $h^2\Lambda = \frac{\mu_0 t_0^2}{36}$, which gives us $\eta = \frac{\mu_0}{18} \leq \frac{1}{18} < 1$. Thus the requirements of Theorem 2 are sufficient to ensure the existence of a piecewise flat metric on $\mathcal{A}$, and we obtain:

Proposition 61 ([DVW15, Proposition 26] corrected) *If the requirements of Theorem 2, are satisfied then the geodesic distances between the* *endpoints of the edges in $\mathcal{A}$ define a piecewise flat metric on $\mathcal{A}$ such that each simplex $\sigma \in \mathcal{A}$ satisfies*

$$t(\sigma) > \frac{3}{4\sqrt{n}} t_0.$$

Likewise, for [DVW15, Theorem 3], it is sufficient to impose the new bounds introduced in Theorem 60 to obtain the same metric distortion bound:

Theorem 62 ([DVW15, Theorem 3] corrected) *If the requirements of Theorem 2 are satisfied, then $\mathcal{A}$ is naturally equipped with a piecewise flat metric $d_{\mathcal{A}}$ defined by assigning to each edge the geodesic distance in M between its endpoints.*

With this metric on $\mathcal{A}$, if $H\colon |\mathcal{A}| \to M$ is the triangulation defined by the barycentric coordinate map, then the metric distortion induced by H is quantified as

$$|d_M(H(x), H(y)) - d_{\mathcal{A}}(x, y)| \leq \frac{50\Lambda h^2}{t_0^2} d_{\mathcal{A}}(x, y),$$

for all $x, y \in |\mathcal{A}|$.

Finally, we remark that [BDG17, Theorem 3] employed [DVW15, Theorem 3], but although some of the discussion leading up to the statement of the result should be modified to account for the new bound imposed by Theorem 62, the actual result [BDG17, Theorem 3] stands as stated, since the bound on the sampling density required there is manifestly sufficient to accommodate (33), and the construction via local Delaunay triangulations in the coordinate domains automatically ensures that the vertex sanity criterion is satisfied.

References

[AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22(4):481–504, 1999. 3, 23

[ACDL02] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Int. J. Computational Geometry and Applications*, 12(2):125–141, 2002. 3, 4

[AEM07] Dominique Attali, Herbert Edelsbrunner, and Yuriy Mileyko. Weak witnesses for Delaunay triangulations of submanifolds. In *Symp. Solid and Physical Modeling*, pages 143–150, 2007. 22

[BDG13] J.-D. Boissonnat, R. Dyer, and A. Ghosh. The stability of Delaunay triangulations. *International Journal of Computational Geometry & Applications*, 23(04n05):303–333, 2013. (arXiv:1304.2947). 9, 11, 25, 39

[BDG17] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay triangulation of manifolds. *Foundations of Computational Mathematics*, 2017. (doi:10.1007/s10208-017-9344-1; arXiv:1311.0117). 2, 47, 49

[BG14] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. *Discrete and Computational Geometry*, 51(1):221–267, 2014. 2, 3, 4, 5, 7, 20, 28

[BLW17a] J.-D. Boissonnat, A. Lieutier, and M. Wintraecken. The reach, metric distortion, geodesic convexity and the variation of tangent spaces. preprint, Inria, Sophia-Antipolis, 2017. `http://ramsaydyer.com/tmp/tanvar.pdf`. 22, 24, 25

[BLW17b] J.-D. Boissonnat, A. Lieutier, and M. Wintraecken. Some properties of submanifolds with positive reach: metric distortion, and geodesic convexity of intersections with spheres. preprint supplement, Inria, Sophia-Antipolis, 2017. `http://ramsaydyer.com/tmp/convsup.pdf`. 40, 41, 42

[BO05] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, 2005. 4

[Bro12] L. E. J. Brouwer. Über Abbildung von Mannigfaltigkeiten. *Math. Ann.*, 71(4):598, 1912. 40

[Cai34] S. S. Cairns. On the triangulation of regular loci. *Annals of Mathematics. Second Series*, 35(3):579–587, 1934. 2

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

[DVW15] R. Dyer, G. Vegter, and M. Wintraecken. Riemannian simplices and triangulations. *Geometriae Dedicata*, 2015. DOI:10.1007/s10711-015-0069-5. (Also: arXiv:1406.3740). 2, 4, 5, 43, 44, 45, 47, 48, 49

[DZM08] R. Dyer, H. Zhang, and T. Möller. Surface sampling and the intrinsic Voronoi diagram. 27(5):1393–1402, 2008. 3

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

[Fed59] H. Federer. Curvature measures. *Trans. Amer. Math. Soc.*, 93(3):418–491, 1959. 22, 24, 29

[Foo84] Robert L. Foote. Regularity of the distance function. *Proc. Amer. Math. Soc.*, 92(1):153–155, 1984. 22

[Mun68] J. R. Munkres. *Elementary differential topology.* Princton University press, second edition, 1968. 3

[OR09] E. Outerelo and J. M. Ruiz. *Mapping degree theory*, volume 108. American Mathematical Soc., 2009. 37, 38, 40, 44

[Whi40] J. H. C. Whitehead. On C^1-complexes. *Ann. of Math*, 41(4):809–824, 1940. 3

[Whi57] H. Whitney. *Geometric Integration Theory.* Princeton University Press, 1957. 3, 4, 8, 25, 37, 40, 44