



HAL
open science

Local Criteria for Triangulating General Manifolds

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

► **To cite this version:**

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh, Mathijs Wintraecken. Local Criteria for Triangulating General Manifolds. *Discrete and Computational Geometry*, 2023, 69, pp.156-191. 10.1007/s00454-022-00431-7 . hal-03794702

HAL Id: hal-03794702

<https://inria.hal.science/hal-03794702v1>

Submitted on 6 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Local Criteria for Triangulating General Manifolds

Jean-Daniel Boissonnat¹ · Ramsay Dyer¹ · Arijit Ghosh² ·
Mathijs Wintraecken³

Received: 20 February 2020 / Revised: 1 February 2022 / Accepted: 14 February 2022
© The Author(s) 2022

Abstract

We present criteria for establishing a triangulation of a manifold. Given a manifold M , a simplicial complex \mathcal{A} , and a map H from the underlying space of \mathcal{A} to M , our criteria are presented in local coordinate charts for M , and ensure that H is a homeomorphism. These criteria do not require a differentiable structure, or even an explicit metric on M . No Delaunay property of \mathcal{A} is assumed. The result provides a triangulation guarantee for algorithms that construct a simplicial complex by working in local coordinate patches. Because the criteria are easily verified in such a setting, they are expected to be of general use.

Keywords Triangulating/Meshing · Triangulation criteria · Local coordinate charts · Manifolds

Mathematics Subject Classification 57Q15 · 68U05 · 57R05

1 Introduction

A *triangulation* of a manifold M is a homeomorphism $H: |\mathcal{A}| \rightarrow M$, 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 .

Editor in Charge: Kenneth Clarkson

This work has been funded by the European Research Council under the European Union's ERC Grant Agreement number 339025 GUDHI (Algorithmic Foundations of Geometric Understanding in Higher Dimensions). Arijit Ghosh is supported by Ramanujan Fellowship (No. SB/S2/RJN-064/2015). Part of this work was done when Arijit Ghosh was a Researcher at Max-Planck-Institute for Informatics, Germany, supported by the IndoGerman Max Planck Center for Computer Science (IMPECS). Mathijs Wintraecken also received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 754411 and the Austrian Science Fund (FWF): M-3073. A part of the results described in this paper were presented at SoCG 2018 and in [3].

Extended author information available on the last page of the article

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 [5, 8]. It complements and is closely related to recent work that investigates a particular natural example of such a map [16].

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 2.17, are based on metric properties of H within “compatible” coordinate charts (Definition 2.4). That is, we make sure the linearization of H inside each coordinate chart is an embedding. If simplices are sufficiently thick and the original map is similar enough to the linearized one, this guarantees that H is a local homeomorphism. One needs an additional technical criterion to make sure that it is a homeomorphism rather than a covering map. 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 [13] 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 [21, 23], but it is not easily adapted to provide triangulation guarantees for complexes constructed by other algorithms.

Submanifolds of \mathbb{R}^N

An alternative approach was developed by Whitney [9, 24] 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. The argument is entwined with this specific construction, and is not easily adapted to other settings.

More recently, Edelsbrunner and Shah [18] (using general results on CW-complexes, see for example [20]) defined the restricted Delaunay complex of a subset M of Euclidean space as the nerve of the Voronoi diagram on M when the ambi-

ent 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*: Voronoi faces are closed topological balls of the appropriate dimension.

Using the closed ball property, Amenta and Bern [1] 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 [2].

Although it was established in the context of restricted Delaunay triangulations, the closed ball property 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 [17], and Cheng et al. [14] have indicated how it can be applied for establishing weighted restricted Delaunay triangulations of smooth submanifolds of arbitrary dimension in Euclidean space.

However, the closed ball property 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 [8]. Secondly, even when a Delaunay-like complex is being constructed, it can be difficult to directly verify the properties of the associated Voronoi structure. The closed ball property does not give us sampling criteria and conditions on the complex under construction. A third deficiency of the closed ball property is that, although it can establish that a complex \mathcal{A} triangulates the manifold M , it does not provide a specific triangulation $H: |\mathcal{A}| \rightarrow M$. Such a correspondence is required to compare *geometric* properties of $|\mathcal{A}|$ and M .

In [8] 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 [2], especially in the formulation presented by Boissonnat and Oudot [11], guarantees a triangulation of a surface by any complex which satisfies a few easily verifiable properties. However, the argument only holds for surfaces in \mathbb{R}^3 .

Riemannian Manifolds

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* $|\mathcal{A}| \rightarrow M$. In [16], conditions are presented which guarantee that this map is a triangulation. Although this map can also be applied to submanifolds of Euclidean space, it is often not easy to compute because it boils down to finding geodesics, which is far more complicated than computing closest points projections. This is perhaps the main reason why we have developed the

general approach of this paper. When revisiting the result for Riemannian manifolds we also generalize to smooth maps and simplify the exposition.

Outline

The general triangulation criteria (Theorem 2.17) are given in Sect. 2, without explicitly defining the map H . In Sect. 3, Theorem 2.17 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 [8], also correcting an error. Section 4 explores the differentiable setting, which is a generalization of the setting where M is a Riemannian manifold, and plugs a gap in [16]. In Appendix A, we discuss some lemmas on general distortion maps. In Appendix B we review degree theory and discuss the link with Whitney's work.

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 $H: |\mathcal{A}| \rightarrow M$ 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 2.1 (*simplices and stars*) In this section, a simplex σ 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 σ is the topological interior of σ considered as a subspace of its affine hull, and is denoted by $\text{relint}(\sigma)$. If σ is a simplex of \mathcal{A} , the subcomplex consisting of all simplices that have σ as a face, together with the faces of these simplices, is called the *star* of σ , denoted by $\text{St}(\sigma)$; the star of a vertex p is $\text{St}(p)$.

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

Notation 2.2 (*topology*) If $A \subseteq \mathbb{R}^m$, then the topological closure, interior, and boundary of A are denoted respectively by \overline{A} , $\text{int}(A)$, and $\partial A = \overline{A} \setminus \text{int}(A)$. We denote by $B_{\mathbb{R}^m}(c, r)$ the open ball in \mathbb{R}^m of radius r and centre c .

Notation 2.3 (*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 $G: |\mathcal{C}| \rightarrow \mathbb{R}^m$, 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 $x \in \sigma \in \mathcal{C}$ 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 (see diagram (2)).

Definition 2.4 (*compatible atlases*) We say that $|\mathcal{A}|$ and M have *compatible atlases* for $H : |\mathcal{A}| \rightarrow M$ if:

- (i) 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 each set U_p is connected.
- (ii) For each $p \in \mathcal{P}$, there is a subcomplex $\widetilde{\mathcal{C}}_p$ of \mathcal{A} that contains $\text{St}(p)$ 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 2.5 The requirement in Definition 2.4 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

$$F_p : |\mathcal{C}_p| \subset \mathbb{R}^m \rightarrow \mathbb{R}^m,$$

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:



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 $\widehat{\Phi}_p$ coincides with $\phi_p \circ H$ on vertices.

Remark 2.6 The setting described here encapsulates the one in the tangential complex work [8] and the one in the Riemannian triangulation paper [16]. In [8] one locally constructs a (weighted) Delaunay triangulation in the tangent space $T_p M$. 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 $\widehat{\mathcal{F}}_p$ 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 $\widehat{\mathcal{F}}_p$ really does coincide with the secant map of $\phi_p \circ H|_{|\widetilde{\mathcal{C}}_p|}$. In [16] the simplices are mapped to the Riemannian manifold using Riemannian centres of mass. In the Riemannian setting the coordinate charts are the natural concepts to work with from the outset.

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 $\widetilde{V}_p = \widehat{\mathcal{F}}_p^{-1}(V_p)$ 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 $H|_{\widetilde{V}_p} = \phi_p^{-1} \circ F_p \circ \widehat{\mathcal{F}}_p|_{\widetilde{V}_p}$, since ϕ_p and $\widehat{\mathcal{F}}_p$ are both embeddings. Since $|\mathcal{A}|$ is compact, Brouwer’s invariance of domain, together with Remark 2.5, 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 in $\text{St}(\hat{p})$ is an orientation preserving topological embedding. As discussed in Appendix B, we say that F_p preserves the orientation of an m -simplex $\sigma \subset \mathbb{R}^m$ if $F_p|_\sigma$ has degree 1 at any point in the image of the interior of σ , i.e., $\deg(F_p, \text{int}(\sigma), y) = 1$ for $y \in F_p(\text{int}(\sigma))$. (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 2.7 below, which is a particular case of an observation made by Whitney [24, Appx. II, Lem. 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 B, 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 2.7 (simplexwise positive embedding) *Assume \mathcal{C} is an oriented m -manifold finite simplicial complex with boundary embedded in \mathbb{R}^m . Let $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ 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., [4, Lems. 3.6 and 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 $w \in F(V) \setminus F(|\mathcal{C}^{m-1}|)$ (Lemma B.9). We need to show that F is also injective on $V \cap |\mathcal{C}^{m-1}|$.

We now show that $F(\text{st}(\sigma))$ is open for any $\sigma \in \mathcal{C}^{m-1} \setminus \partial\mathcal{C}$, where $\text{st}(\sigma)$ is the open star of σ , defined in Notation 2.1. Suppose $x \in \text{relint}(\tau)$ for some $\tau \in \mathcal{C} \setminus \partial\mathcal{C}$. 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 $U \cap F(\partial\text{st}(\tau)) = \emptyset$. Since the closure of the open star is equal to the carrier of our usual star,

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

Lemma B.9 implies that every point in $U \setminus F(|\underline{\text{St}}(\tau)^{m-1}|)$ 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 $U \subseteq F(\text{st}(\tau))$.

If $x \in \text{st}(\sigma)$, then $x \in \text{relint}(\tau)$ for some $\tau \in \mathcal{C} \setminus \partial\mathcal{C}$ that has σ as a face. Since $\text{st}(\tau) \subseteq \text{st}(\sigma)$, we have $U \subseteq F(\text{st}(\sigma))$, and we conclude that $F(\text{st}(\sigma))$ 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 σ, τ such that $w \in \text{relint}(\sigma)$ and $z \in \text{relint}(\tau)$. So there is an open neighbourhood U of $F(w) = F(z)$ that is contained in $F(\text{st}(\sigma)) \cap F(\text{st}(\tau))$.

We must have $\text{st}(\sigma) \cap \text{st}(\tau) = \emptyset$, because if $x \in \text{st}(\sigma) \cap \text{st}(\tau)$, then $x \in \text{relint}(\mu)$ for some μ that has both σ and τ as faces. But this means that both w and z belong to μ , contradicting the injectivity of $F|_\mu$. 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 2.7 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 2.8 (ξ -distortion map) A map $F : U \subset \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a ξ -distortion map if for all $x, y \in U$ we have

$$\left| \|F(x) - F(y)\| - \|x - y\| \right| \leq \xi \|x - y\|. \tag{3}$$

We are interested in ξ -distortion maps with small ξ . (3) can be equivalently written as

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

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 $(\xi/(1 - \xi))$ -distortion map (see Lemma A.1 (i)).

Assuming that $F_p|_\sigma$ is a ξ -distortion map for each m -simplex $\sigma \in \mathcal{C}_p$, 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 σ 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 2.9 (*simplex quality*) If p is a vertex of σ , the *altitude* of p is the distance from p to the opposing facet of σ and is denoted $a_p(\sigma)$. The *thickness* of σ , denoted $t(\sigma)$ (or just t if there is no risk of confusion), is given by $a/(mL)$, where $a = a(\sigma)$ is the smallest altitude of σ , and $L = L(\sigma)$ is the length of the longest edge. We set $t(\sigma) = 1$ if σ has dimension 0.

Lemma 2.10 (*trilateration*) Suppose $\sigma \subset \mathbb{R}^m$ is an m -simplex, and $F: \sigma \rightarrow \mathbb{R}^m$ is a ξ -distortion map that leaves the vertices of σ fixed. If $\xi \leq 1$, then for any $x \in \sigma$,

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

where L is the length of the longest edge of σ , and t is its thickness.

Proof Let $\{p_0, \dots, p_m\}$ be the vertices of σ . For $x \in \sigma$, let $\tilde{x} = F(x)$. We choose p_0 as the origin, and observe that

$$p_i^\top x = \frac{\|x\|^2 + \|p_i\|^2 - \|x - p_i\|^2}{2}, \tag{4}$$

which we write in matrix form as $P^\top x = b$, 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 (4). Similarly, we have $P^\top \tilde{x} = \tilde{b}$ with the obvious definition of \tilde{b} . Then

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

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 [4, Lem. 2.4] we have $\|(P^\top)^{-1}\| \leq (\sqrt{m} tL)^{-1}$, and the stated bound follows. □




Fig. 1 The open set V_p (shaded) is a homothetically shrunk copy of the interior of $\underline{\text{St}}(\hat{p})$

2.1.1 Using $\underline{\text{St}}(p)$ as $\tilde{\mathcal{C}}_p$

For the local complex $\tilde{\mathcal{C}}_p \subset \mathcal{A}$ introduced in Definition 2.4, we now make a specific choice: $\tilde{\mathcal{C}}_p = \underline{\text{St}}(p)$. This is the smallest complex allowed by the definition. For convenience, we define $\hat{p} = \hat{\Phi}_p(p)$, so that $\mathcal{C}_p = \hat{\Phi}_p(\tilde{\mathcal{C}}_p) = \underline{\text{St}}(\hat{p})$.

We define V_p to be the open set obtained by homothetically “shrinking” $|\underline{\text{St}}(\hat{p})|$ such that it is just large enough to contain the barycentres of the simplices that have \hat{p} as a vertex (see Fig. 1). To be more specific we define V_p to be the open set consisting of the points in $|\underline{\text{St}}(\hat{p})|$ whose barycentric coordinate with respect to \hat{p} is strictly larger than $1/(m + 1) - \delta$, where $\delta > 0$ is arbitrarily small. Since the barycentric coordinates in each m -simplex sum to 1, and the piecewise linear maps $\hat{\Phi}_p$ preserve barycentric coordinates, this ensures that the sets $\hat{\Phi}_p^{-1}(V_p)$ cover $|\mathcal{A}|$.

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 V_p . In order to employ the simplexwise positive embedding lemma (Lemma 2.7), 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 2.11 (a point covered once) *Suppose \mathcal{C} is a pure m -dimensional finite simplicial complex embedded in \mathbb{R}^m , and that for each $\sigma \in \mathcal{C}$ we have $t(\sigma) \geq t_0$. Let $\sigma \in \mathcal{C}$ be an m -simplex with the largest diameter, i.e., $L(\sigma) \geq L(\tau)$ for all $\tau \in \mathcal{C}$, and let b be the barycentre of σ . If $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ leaves the vertices of \mathcal{C} fixed, and its restriction to any m -simplex in \mathcal{C} is a ξ -distortion map with*

$$\xi \leq \frac{1}{6} \cdot \frac{mt_0^2}{m + 1}, \tag{5}$$

then $F^{-1}(F(b)) = \{b\}$.

Proof Since $\xi < 1$, the restriction of F to σ is injective. Suppose $x \in |\mathcal{C}|$ is such that $F(x) = F(b)$, but $x \neq b$. Then x belongs to some m -simplex $\tau \in \mathcal{C}$ different from σ .

Since the distance from b to the boundary of σ is

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

it follows that $\|x - b\| > mt(\sigma)L(\sigma)/(m + 1)$. But using Lemma 2.10 and the constraint (5) we arrive at a contradiction:

$$\|x - b\| \leq \|b - F(b)\| + \|x - F(x)\| \leq \frac{3\xi L(\sigma)}{t(\sigma)} + \frac{3\xi L(\tau)}{t(\tau)} \leq \frac{mt(\sigma)L(\sigma)}{m + 1}.$$

□

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 $\underline{\text{St}}(\hat{p})$. We say that $\underline{\text{St}}(\hat{p})$ is a *full star* if its carrier is an m -manifold with boundary and \hat{p} does not belong to $\partial\underline{\text{St}}(\hat{p})$.

Lemma 2.12 (barycentric boundary separation) *Suppose $\underline{\text{St}}(\hat{p})$ is a full m -dimensional star embedded in \mathbb{R}^m . Let $a_0 = \min_{\sigma \in \underline{\text{St}}(\hat{p})} a_{\hat{p}}(\sigma)$ be the smallest altitude of \hat{p} in the m -simplices in $\underline{\text{St}}(\hat{p})$. Suppose $x \in \sigma \in \underline{\text{St}}(\hat{p})$, where σ is an m -simplex, and $\lambda_{\sigma, \hat{p}}(x)$, the barycentric coordinate of x with respect to \hat{p} in σ , satisfies $\lambda_{\sigma, \hat{p}}(x) \geq \alpha$. Then $d_{\mathbb{R}^m}(x, |\partial\underline{\text{St}}(\hat{p})|) \geq \alpha a_0$. If t_0 is a lower bound on the thicknesses of the simplices in $\underline{\text{St}}(\hat{p})$, and s_0 is a lower bound on their diameters, then $d_{\mathbb{R}^m}(x, |\partial\underline{\text{St}}(\hat{p})|) \geq \alpha m t_0 s_0$.*

Proof Since we are interested in the distance to the boundary, consider a point $y \in |\partial\underline{\text{St}}(\hat{p})|$ such that the segment $[x, y]$ lies in $|\underline{\text{St}}(\hat{p})|$. The segment passes through a sequence of m -simplices, $\sigma_0 = \sigma, \sigma_1, \dots, \sigma_n$, that partition it into subsegments $[x_i, y_i] \subset \sigma_i$ with $x_0 = x, y_n = y$, and $x_i = y_{i-1}$ for all $i \in \{1, \dots, n\}$.

Observe that $\lambda_{\sigma_i, \hat{p}}(x_i) = \lambda_{\sigma_{i-1}, \hat{p}}(y_{i-1})$, and that

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

Thus

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

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

Lemma 2.12 allows us to quantify the distortion bound that we need to ensure that the boundary of $\underline{\text{St}}(\hat{p})$ does not get mapped by F_p into the image of the open set V_p . The argument is the same as for Lemma 2.11, 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 (5) and (6)).

Lemma 2.13 (boundary separation for V_p) *Suppose $\text{St}(\hat{p})$ is a full star embedded in \mathbb{R}^m , and every m -simplex σ in $\text{St}(\hat{p})$ satisfies $s_0 \leq L(\sigma) \leq L_0$, and $t(\sigma) \geq t_0$. If the restriction of F_p to any m -simplex in $\text{St}(\hat{p})$ is a ξ -distortion map, with*

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

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

Proof If $x \in \sigma \in \text{St}(\hat{p})$ has barycentric coordinate with respect to \hat{p} larger than $1/(m+1) - \delta$, and $y \in \tau \in \partial \text{St}(\hat{p})$, then Lemmas 2.10 and 2.12 ensure that $F_p(x) \neq F_p(y)$ provided

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

which is satisfied by (6) when $\delta > 0$ satisfies

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

When inequality (6) (and therefore also (5)) is satisfied, we can employ the embedding lemma (Lemma 2.7) to guarantee that V_p is embedded:

Lemma 2.14 (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 $H : |\mathcal{A}| \rightarrow M$ is a covering map if the following criteria are satisfied:*

(i) **compatible atlases** *There are compatible atlases for H , with $\mathcal{C}_p = \text{St}(p)$ for each $p \in \mathcal{P}$ (Definition 2.4).*

For the charts in the compatible atlases we have the following:

- (ii) **simplex quality** *For each $p \in \mathcal{P}$, every simplex $\sigma \in \text{St}(\hat{p}) = \widehat{\Phi}_p(\text{St}(p))$ satisfies $s_0 \leq L(\sigma) \leq L_0$ and $t(\sigma) \geq t_0$ (Notation 2.9).*
- (iii) **distortion control** *For each $p \in \mathcal{P}$, the map*

$$F_p = \phi_p \circ H \circ \widehat{\Phi}_p^{-1} : |\text{St}(\hat{p})| \rightarrow \mathbb{R}^m,$$

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

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

(Definitions B.5 and 2.8).

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 $H(q) \in H(\sigma)$, then q is a vertex of σ . This ensures that $H^{-1}(H(q)) = \{q\}$, and by Remark 2.5 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 2.15 (*vertex sanity*) If $H: |\mathcal{A}| \rightarrow M$ has compatible atlases (Definition 2.4), then H exhibits *vertex sanity* if: for all vertices $p, q \in \mathcal{P}$, if $\phi_p \circ H(q) \in |\text{St}(\hat{p})| = \widehat{\Phi}_p(|\text{St}(p)|)$, then q is a vertex of $\text{St}(p)$.

Together with the distortion bounds that are imposed on F_p , Definition 2.15 ensures that the image of a vertex cannot lie in the image of a simplex to which it does not belong:

Lemma 2.16 (*injectivity*) If $H: |\mathcal{A}| \rightarrow M$ satisfies the hypotheses of Lemma 2.14 as well as Definition 2.15, then H is injective, and therefore a homeomorphism.

Proof Towards a contradiction, suppose that $H(q) \in H(\sigma)$ and that q is not a vertex of the m -simplex σ . This means there is some $x \in \sigma$ such that $H(x) = H(q)$. Let p be a vertex of σ . The vertex sanity hypothesis (Definition 2.15) implies that $\phi_p \circ H(q)$ must be either outside of $|\text{St}(\hat{p})|$, or belong to its boundary. Thus Lemmas 2.12 and 2.10, and the bound on ξ from Lemma 2.14 (iii) imply that the barycentric coordinate of x with respect to p must be smaller than $\frac{1}{\gamma}(m + 1)$: Let $\hat{x} = \widehat{\Phi}_p(x)$, and $\hat{\sigma} = \widehat{\Phi}_p(\sigma)$. Lemma 2.10 says that

$$|F_p(\hat{x}) - \hat{x}| \leq \frac{3\xi L_0}{t_0} < \frac{ms_0 t_0}{2(m + 1)} \leq \frac{a_0}{2(m + 1)},$$

where a_0 is a lower bound on the altitudes of \hat{p} , as in Lemma 2.12. Since $F_p(\hat{x}) = \phi_p \circ H(x)$ is at least as far away from \hat{x} as $\partial\text{St}(\hat{p})$, Lemma 2.12 implies that the barycentric coordinate of $\hat{x} \in \hat{\sigma}$ with respect to \hat{p} must be no larger than $1/(2(m + 1))$. Since $\widehat{\Phi}_p$ preserves barycentric coordinates, and the argument works for any vertex p of σ , we conclude that all the barycentric coordinates of x in σ are strictly less than $1/(m + 1)$. We have reached a contradiction with the fact that the barycentric coordinates of x must sum to 1. □

2.3 Main Result

To recap, Lemmas 2.14 and 2.16 yield the following triangulation result. In the bound on ξ from Lemma 2.14 (iii), we replace the factor $m/(m + 1)$ with $1/2$, the lower bound attained when $m = 1$.

Theorem 2.17 (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 $H: |\mathcal{A}| \rightarrow M$ is a homeomorphism if the following criteria are satisfied:*

(i) **compatible atlases** *There are compatible atlases*

$$\{(\tilde{\mathcal{C}}_p, \hat{\Phi}_p)\}_{p \in \mathcal{P}}, \quad \tilde{\mathcal{C}}_p \subset \mathcal{A}, \quad \text{and} \quad \{(U_p, \phi_p)\}_{p \in \mathcal{P}}, \quad U_p \subset M,$$

for H , with $\tilde{\mathcal{C}}_p = \underline{\text{St}}(p)$ for each $p \in \mathcal{P}$, the vertex set of \mathcal{A} (Definition 2.4).

For the charts in the compatible atlases we have the following:

- (ii) **simplex quality** *For each $p \in \mathcal{P}$, every simplex $\sigma \in \underline{\text{St}}(\hat{p}) = \hat{\Phi}_p(\underline{\text{St}}(p))$ satisfies $s_0 \leq L(\sigma) \leq L_0$ and $t(\sigma) \geq t_0$ (Notation 2.9).*
- (iii) **distortion control** *For each $p \in \mathcal{P}$, the map*

$$F_p = \phi_p \circ H \circ \hat{\Phi}_p^{-1}: |\underline{\text{St}}(\hat{p})| \rightarrow \mathbb{R}^m,$$

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

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

(Definitions B.5 and 2.8).

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

Heuristically speaking the first criterion ensures, among others, that all the stars of vertices completely lie in a coordinate neighbourhood. Without this condition there could be parts that are not covered by the simplicial complex. The criteria (ii) and (iii) prevent the orientation of a simplex to change from one map to another, while the vertex sanity condition prevents multiple covers.

Remark 2.18 The constants L_0 , s_0 , and t_0 that constrain the simplices in the local complex $\underline{\text{St}}(\hat{p})$, and the metric distortion of F_p in Theorem 2.17 can be considered to be local, i.e., they may depend on $p \in \mathcal{P}$.

3 Submanifolds of Euclidean Space

As a specific application of Theorem 2.17, 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 [8]. In that case \mathcal{A} is constructed as a subcomplex of a weighted Delaunay triangulation of \mathcal{P} in the ambient space \mathbb{R}^N , so it is necessarily embedded. However, in the general setting discussed in this paper we do not need to assume *a priori* that \mathcal{A} is embedded

in \mathbb{R}^N . (This does not force us to consider \mathcal{A} to be abstract in the combinatorial sense. In particular, the simplices are Euclidean simplices, not just sets of vertices.) Instead, we assume only that the embedding of the vertex set $\mathcal{P} \hookrightarrow \mathbb{R}^N$ defines an *immersion* $\iota: |\mathcal{A}| \rightarrow \mathbb{R}^N$. By this we mean that for any vertex $p \in \mathcal{P}$ we have that the restriction of ι to $|\text{St}(p)|$ is an embedding.

At each point $x \in M$, the tangent space $T_x M \subset T_x \mathbb{R}^N$ is naturally viewed as an m -dimensional 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 2.17 will be defined by projection maps. The coordinate charts are defined using the orthogonal projection $\text{pr}_{T_p M}: \mathbb{R}^N \rightarrow T_p M$. As discussed in Sect. 3.3, for a sufficiently small neighbourhood $U_p \subset M$, we obtain an embedding

$$\phi_p = \text{pr}_{T_p M}|_{U_p}: U_p \subset M \rightarrow T_p M \cong \mathbb{R}^m,$$

which will define our coordinate maps for M .

For the map $H: |\mathcal{A}| \rightarrow M$, we will employ the closest point projection map defined in Sect. 3.1 and discussed further in Sect. 3.4. 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 $\text{pr}_M: U_M \rightarrow M$ is well defined. We define $H = \text{pr}_M \circ \iota$.

As demanded by Definition 2.4, 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 $\widetilde{\mathcal{C}}_p = |\text{St}(p)|$, and since $\text{pr}_{T_p M}$ is already a linear map, and pr_M is the identity on the vertices, this means $\widehat{\Phi}_p = \text{pr}_{T_p M} \circ \iota|_{|\text{St}(p)|}$.

In Sects. 3.1 and 3.2 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 A.1, we are free to choose any convenient metric on M . We employ here the metric of the ambient space \mathbb{R}^N , rather than the intrinsic metric of geodesic distances.

3.1 The Reach

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 $\text{ax}(M) \subset \mathbb{R}^N$ that have more than one such closest point on M . In other words, if $x \in \text{ax}(M)$, 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 $\overline{\text{ax}}(M)$. The *reach* of M is defined by $\text{rch}(M) := d_{\mathbb{R}^N}(M, \overline{\text{ax}}(M))$. We observe below that for compact C^2 submanifolds, $\text{rch}(M) > 0$. Thus, by definition, every point x in the open neighbourhood $U_M := \mathbb{R}^N \setminus \overline{\text{ax}}(M)$ of M has a unique closest point $\check{x} \in M$. The *closest-point projection* map

$$\text{pr}_M: U_M \rightarrow M,$$

takes x to this closest point: $\text{pr}_M(x) = \check{x}$. To each point $x \in M$, we associate a *normal space*

$$N_x M = \{n \in T_x \mathbb{R}^N \mid n \cdot v = 0 \text{ for all } v \in T_x M\}$$

of vectors orthogonal to $T_x M$. Thus $T_x \mathbb{R}^N = N_x M \oplus T_x M$. 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 $T_x M$.

Definition 3.1 The reach of M is the distance between the medial axis and M itself.

Remark 3.2 Most of the results below can be localized in terms of the local feature size [1, 19] (or similar quantities), see [7] for details.

3.2 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, $\text{pr}_K(u)$ is the orthogonal projection of u into K . We define $\angle(u, K)$ to be $\pi/2$ if $\text{pr}_K(u) = 0$, and otherwise $\angle(u, K) = \angle(u, \text{pr}_K(u))$ (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 σ is a simplex with $\dim \sigma \leq \dim L$, then $\angle(\sigma, L) := \angle(\text{aff}(\sigma), L)$, where $\text{aff}(\sigma)$ is the affine hull of σ . 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 3.3

- (i) If K, L are subspaces of \mathbb{R}^N , then $\angle(L^\perp, K^\perp) = \angle(K, L)$.
- (ii) If $Q \subset \mathbb{R}^N$ is a subspace of codimension 1, then $\angle(Q^\perp, K) = \pi/2 - \angle(K, Q)$.

We now recall a number of results starting with [19, Thm. 4.8(7)].

Lemma 3.4 Given any two points $x, y \in M \subset \mathbb{R}^N$, we have

- (i) $\sin \angle([x, y], T_x M) \leq \frac{|x - y|}{2 \text{rch}(M)}$,
- (ii) $d_{\mathbb{R}^N}(y, T_x M) \leq \frac{|x - y|^2}{2 \text{rch}(M)}$.

We will also use the following result, Lemma 6 and Corollary 3 of [10]:

Lemma 3.5 (tangent space variation) Let $x, y \in M$. Then

$$\sin \angle(T_x M, T_y M) \leq \frac{|y - x|}{\text{rch}(M)} \quad \text{and} \quad \angle(T_x M, T_y M) \leq \frac{\pi |y - x|}{2 \text{rch}(M)}.$$

Remark 3.6 In the proof of the lemma above the fact that M is C^2 is used, however alternatives with worse constants exist for general manifolds with positive reach [10]. In particular all results in this section hold for general manifolds of positive reach albeit with worse constants.

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 [24, p. 127] in the formulation presented in [4, Lem. 2.1]:

Lemma 3.7 (Whitney’s angle bound) *Suppose σ 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*

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

where t is the thickness of σ and L is the length of its longest edge.

The following lemma can be found in [3, Lem. 7.9].

Lemma 3.8 (simplices lie close to M) *Let σ be a simplex with vertices on M . Then for all $x, y \in \sigma$,*

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

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

Lemma 3.9 (simplex-tangent space angle bounds) *Suppose $\sigma \subset \mathbb{R}^N$ is a simplex of dimension $\leq m$ with vertices on M . If p is a vertex of σ , then*

- (i) *if p is a vertex of σ , $\sin \angle(\sigma, T_pM) \leq \frac{L}{t \text{rch}(M)}$;*
- (ii) *if $x \in \sigma$ and $\check{x} = \text{pr}_M(x)$, then $\sin \angle(\sigma, T_{\check{x}}M) \leq \frac{3L}{t \text{rch}(M)}$.*

Proof (i) By Lemma 3.4(ii), all the vertices of σ are within a distance $\eta = L^2/(2 \text{rch}(M))$ from T_pM , and so Lemma 3.7 ensures

$$\sin \angle(\sigma, T_pM) \leq \frac{2\eta}{tL} = \frac{L}{t \text{rch}(M)}.$$

(ii) Because $|x - \check{x}| < L(\sigma)$, and $|x - p| < L(\sigma)$, we see that $|p - \check{x}| < 2L(\sigma)$. Lemma 3.5 now yields

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

and the result follows using part (i). □

3.3 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_pM . 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 $\xi < 1$. The proof of the following lemma can be found in [3, Lem. 7.14].

Lemma 3.10 (definition and distortion of ϕ_p) *Let $U_p = B_{\mathbb{R}^N}(p, r) \cap M$, where $r = \rho \operatorname{rch}(M)$, with $\rho < 1/2$. Define $\phi_p := \operatorname{pr}_{T_pM}|_{U_p}$. Then ϕ_p is a ξ -distortion map with $\xi = 4\rho^2$.*

Remark 3.11 (differential of ϕ_p) It is straightforward to verify from the definitions that for any $x \in U_p$,

$$d(\operatorname{pr}_{T_pM}|_M)_x = \operatorname{pr}_{T_pM}|_{T_xM}.$$

The domain of the map $\widehat{\Phi}_p$, i.e., an upper bound on the allowable size of the simplices in $\iota(|\underline{\operatorname{St}}(p)|)$, is governed by the following bound on the metric distortion of the projection from a simplex. Here $d(f)_x$ denotes the differential of the function f at the point x .

Lemma 3.12 (simplexwise distortion of $\widehat{\Phi}_p$) *Suppose $\sigma \subset \mathbb{R}^N$ is a simplex of dimension $\leq m$ with vertices on M . If p is a vertex of σ , and*

$$L(\sigma) < t(\sigma) \operatorname{rch}(M),$$

then the restriction of pr_{T_pM} to σ is a ξ -distortion map with

$$\xi = \left(\frac{L(\sigma)}{t(\sigma) \operatorname{rch}(M)} \right)^2.$$

Proof Let $x, y \in \sigma$ and set $\hat{x} = \operatorname{pr}_{T_pM}(x)$, $\hat{y} = \operatorname{pr}_{T_pM}(y)$. By Lemma 3.9 (i),

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

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

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

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

Therefore, in order to use our framework, we need to ensure that for each simplex $\sigma \in \mathcal{A}$, the simplex $\iota(\sigma) \subset \mathbb{R}^N$ satisfies $L < t \operatorname{rch}(M)$. Then, in conformance with Sect. 2.1.1, we set $\widehat{\Phi}_p = \operatorname{pr}_{T_p M} \circ \iota|_{|\underline{\operatorname{St}}(p)|}$, and we *require* that it be an embedding. Although we have established that $\operatorname{pr}_{T_p M}$ is an embedding on each simplex, and we have assumed that $\iota|_{|\underline{\operatorname{St}}(p)|}$ 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 3.16), i.e., something that needs to be established in context. For the case of the tangential Delaunay complex, the embedding follows naturally because $\phi_p(|\underline{\operatorname{St}}(p)|)$ is seen as a weighted Delaunay triangulation in $T_p M$ [8].

3.4 Distortion of the Closest-Point Projection Map: H

Recall that $H: |\mathcal{A}| \rightarrow M$ is the map that we wish to show is a homeomorphism. In our current context H is based on the closest-point projection map $\operatorname{pr}_M: \mathbb{R}^N \rightarrow M$. As discussed at the beginning of this section, we define $H = \operatorname{pr}_M \circ \iota$, where $\iota: |\mathcal{A}| \rightarrow \mathbb{R}^N$ 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 .

We will be needing the following result from [19, Thm. 4.8(8)]:

Lemma 3.13 (upper bound for pr_M distortion) *Let $x, y \in U_M$ and $\check{x} = \operatorname{pr}_M(x)$, and $\check{y} = \operatorname{pr}_M(y)$, as usual. If $a \geq \max\{\delta_M(x), \delta_M(y)\}$ for some $a < \operatorname{rch}(M)$, then*

$$|\check{y} - \check{x}| \leq \left(1 - \frac{a}{\operatorname{rch}(M)}\right)^{-1} |y - x|. \tag{7}$$

Lemma 3.14 (simplexwise distortion of H) *Suppose $\sigma \subset U_M$ is a simplex of dimension $\leq m$ whose vertices lie in M . If $L(\sigma) < t(\sigma) \operatorname{rch}(M)/3$, then the restriction of pr_M to σ is a ξ -distortion map with*

$$\xi = \frac{12L^2}{t^2 \operatorname{rch}(M)^2}.$$

Proof By Lemma 3.8,

$$\delta_M(x) < a = \frac{2L^2}{\operatorname{rch}(M)} \quad \text{for any } x \in \sigma. \tag{8}$$

Thus it follows from Lemma 3.13 that

$$|\check{y} - \check{x}| \leq \left(1 - \frac{2L^2}{\operatorname{rch}(M)^2}\right)^{-1} |y - x| \leq \left(1 + \frac{4L^2}{\operatorname{rch}(M)^2}\right) |y - x|, \tag{9}$$

for any $x, y \in \sigma$.

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

$$|\check{y} - \check{x}| \geq |\widehat{y} - \check{x}| \geq |\widehat{y} - \check{x}| - |\widehat{y} - \widehat{y}|. \tag{10}$$

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

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

and by Lemma 3.9(ii),

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

Thus,

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

and so

$$|\widehat{y} - \check{x}| \geq \left(1 - \frac{9L^2}{t^2 \text{rch}(M)^2}\right) |y - x|. \tag{11}$$

To get an upper bound on the second term on the right side of (10), let Q_y be the hyperplane through \check{y} and orthogonal to $[\check{y}, y]$. We have

$$\begin{aligned} |\widehat{y} - \widehat{y}| &= |\check{y} - y| \cos \angle([\check{y}, y], Q_x) \\ &= |\check{y} - y| \sin \angle(Q_x, Q_y) \text{ by Lemma 3.3 (ii)} \\ &\leq |\check{y} - y| \sin \angle(T_{\check{x}}, T_{\check{y}}) \\ &\leq |\check{y} - y| \frac{|\check{y} - \check{x}|}{\text{rch}(M)} \text{ by Lemma 3.5} \\ &\leq \frac{2L^2}{\text{rch}(M)^2} |\check{y} - \check{x}| \text{ by (8)}. \end{aligned}$$

Now, using (9), we have

$$|\widehat{y} - \widehat{y}| \leq \frac{2L^2}{\text{rch}(M)^2} \left(1 + \frac{4L^2}{\text{rch}(M)^2}\right) |y - x| < \frac{3L^2}{\text{rch}(M)^2} |y - x|, \tag{12}$$

since the hypothesis $3L/\text{rch}(M) < t$ implies $4L^2/\text{rch}(M)^2 < 1/2$. Finally, plugging (11) and (12) back into (10), we get

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

Comparing this lower bound with (9), we arrive at the stated value for the metric distortion ξ . □

3.5 Triangulation Criteria for Submanifolds

In order to employ Theorem 2.17 we first ensure that we meet the compatible atlases criteria (Definition 2.4). In Lemma 3.10 we defined

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

We need to ensure that $H(|\underline{\operatorname{St}}(p)|) \subseteq U_p$. In our context, this means that we require $\operatorname{pr}_M(\iota(|\underline{\operatorname{St}}(p)|)) \subset U_p$, and Lemma 3.8 ensures that if $L(\sigma) \leq L_0$ for all $\sigma \in \iota(\underline{\operatorname{St}}(p))$, then it is sufficient to choose $r = L_0 + 2L_0^2/\operatorname{rch}(M)$, or

$$\rho = \frac{L_0}{\operatorname{rch}(M)} \left(1 + \frac{2L_0}{\operatorname{rch}(M)} \right). \tag{13}$$

Our bound on L_0 itself will be much smaller than $\operatorname{rch}(M)$.

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 $\widehat{\Phi}_p(\iota(\underline{\operatorname{St}}(p)))$, and ensure that it meets the distortion control criterion of Theorem 2.17. Anticipating the bound we will need to meet the distortion-control criterion of Theorem 2.17, we impose the constraint

$$\frac{L_0^2}{t_0^2 \operatorname{rch}(M)^2} \leq \frac{1}{16^2}, \tag{14}$$

where t_0 is a lower bound on the thickness: $t(\sigma) \geq t_0$ for all $\sigma \in \iota(\underline{\operatorname{St}}(p))$. We remark that L_0 and $\operatorname{rch}(M)$ 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 $L_0/\operatorname{rch}(M)$ will be global constants.

Using (14) together with Lemmas 3.12 and A.1 (i), we can bound the metric distortion of $\widehat{\Phi}_p^{-1}$ as

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

Lemma 3.14 gives us the distortion of H :

$$\xi_2 = 12 \frac{L_0^2}{t_0^2 \operatorname{rch}(M)^2}.$$

For ϕ_p , using Lemma 3.10 and (13) we get the distortion bound

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

Lemma A.1 (ii) 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 (14) we find that F_p is a ξ -distortion map with

$$\xi = \frac{19L_0^2}{t_0^2 \operatorname{rch}(M)^2}. \tag{15}$$

Observe that (14) 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 2.17. We have chosen to use here the properties of the Euclidean simplices in the ambient space \mathbb{R}^N , but in Theorem 2.17 we are considering simplices in the local coordinate space; for us these are the projected simplices, e.g., $\hat{\sigma} = \widehat{\Phi}_p(\sigma) = \operatorname{pr}_{T_p M}(\sigma)$. Using Lemma 3.12, the distortion properties of the affine map $\operatorname{pr}_{T_p M}$ imply

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

and therefore

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

We can thus set

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

where s_0 is a lower bound for the diameters of the simplices in $t(\operatorname{St}(p))$. Then, in order to meet the distortion control criterion of Theorem 2.17, we require

$$\frac{19L_0^2}{t_0^2 \operatorname{rch}(M)^2} < \frac{\hat{s}_0 \hat{t}_0^2}{12\hat{L}_0} = \left(1 - \frac{L_0^2}{t_0^2 \operatorname{rch}(M)^2}\right)^3 \frac{s_0 t_0^2}{12L_0}. \tag{16}$$

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

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

we see that (16) is satisfied if

$$L_0^2 \leq \frac{\mu_0 t_0^4 \operatorname{rch}(M)^2}{16^2}. \tag{17}$$

Our constraint (17) on L_0 induces a bound on ρ , as defined in (13),

$$\rho \operatorname{rch}(M) \leq \frac{9}{8} L_0 \leq \frac{9}{27} \operatorname{rch}(M). \tag{18}$$

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

Lemma 3.15 *F_p is simplexwise positive on $\operatorname{pr}_{T_p M}(\iota(|\underline{\operatorname{St}}(p)|))$.*

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 $\widehat{\sigma} \in \operatorname{pr}_{T_p M}(\iota(\underline{\operatorname{St}}(p)))$ is differentiable. The bound (17) together with (15) implies F_p is a ξ -distortion map with $\xi < 1$, so Lemma A.2 ensures that for any m -simplex $\sigma \in \operatorname{pr}_{T_p M}(\iota(\underline{\operatorname{St}}(p)))$, the differential of F_p does not vanish on σ .

We choose an orientation on U_p ; thus we have an orientation on each tangent space $T_x M$, $x \in U_p$. The projection map $\operatorname{pr}_{T_p M}|_{T_x M} = d(\operatorname{pr}_{T_p M}|_M)_x = d(\phi_p)_x$ (see Remark 3.11) is then orientation preserving, because of continuity: it is certainly true when $x = p$, and Lemmas 3.10 and A.2 imply that the differential $d(\phi_p)$ is nondegenerate on U_p . Thus ϕ_p is simplexwise positive.

For an m -simplex $\sigma \in \iota(\underline{\operatorname{St}}(p))$ we define the orientation such that $\operatorname{pr}_{T_p M}|_\sigma$ is positive. Equivalently, we define the orientation on $\iota(|\underline{\operatorname{St}}(p)|)$ 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 $\operatorname{pr}_M|_\sigma$ is positive for each m -simplex $\sigma \in \iota(\underline{\operatorname{St}}(p))$. First observe that for any $x \in \mathbb{R}^N$, the kernel of $d(\operatorname{pr}_M)_x$ 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 $\check{x} = \operatorname{pr}_M(x)$). Also, observe that if $x \in M$, then $d(\operatorname{pr}_M|_{T_x M})_x = \operatorname{id}_{T_x M}$. This is easily seen by noting that $d(\operatorname{pr}_M|_{T_x M})_x = d(\operatorname{pr}_M)_x|_{T_x M}$ and using curves on M to apply the definition of the differential.

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

$$d(\operatorname{pr}_M|_\sigma)_p = \operatorname{id}_{T_p M} \circ \operatorname{pr}_{T_p M}|_\sigma.$$

Since we have established that $\operatorname{pr}_{T_p M}|_\sigma$ is positive, it follows that $d(\operatorname{pr}_M|_\sigma)_p$ is positive, and since $d(\operatorname{pr}_M|_\sigma)$ is nondegenerate on σ (because $d(F_p)$ is), it follows that $H|_{\iota(|\underline{\operatorname{St}}(p)|)}$ is simplexwise positive. Thus F_p is orientation preserving on each simplex of $\underline{\operatorname{St}}(\hat{p})$. □

Theorem 3.16 (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 $\mathcal{A}^0 \rightarrow \mathcal{P}$ such that the resulting piecewise linear map $\iota: |\mathcal{A}| \rightarrow \mathbb{R}^N$ is an immersion, i.e., $\iota|_{|\underline{\operatorname{St}}(p)|}$ is an embedding for each vertex p . Assume that:*

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

- (b) There are constants $0 < t_0 \leq 1$, $0 < \mu_0 \leq 1$, and $\varepsilon_0 > 0$ such that for each simplex $\sigma \in \iota(\mathcal{A})$, and each vertex $p \in \sigma$,

$$t(\sigma) \geq t_0, \quad \mu_0 \varepsilon_0 \operatorname{rch}(M) \leq L(\sigma) \leq \varepsilon_0 \operatorname{rch}(M), \quad \varepsilon_0 \leq \frac{\mu_0^{1/2} t_0^2}{16}.$$

- (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 $\operatorname{pr}_{T_p M}(q) \in \operatorname{pr}_{T_p M}(\iota(\underline{\operatorname{St}}(p)))$ if and only if q is a vertex of $\underline{\operatorname{St}}(p)$.

Then:

- (i) ι is an embedding, so the complex \mathcal{A} may be identified with $\iota(\mathcal{A})$.
- (ii) The closest-point projection map $\operatorname{pr}_M|_{|\mathcal{A}|}$ is a homeomorphism $|\mathcal{A}| \rightarrow M$.
- (iii) For any $x \in \sigma \in \mathcal{A}$,

$$\delta_M(x) = |\check{x} - x| \leq 2\varepsilon_0^2 \operatorname{rch}(M), \quad \sin \angle(\sigma, T_{\check{x}}) \leq \frac{3\varepsilon_0}{t_0},$$

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

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 (17). Condition (c) is the vertex sanity condition of Theorem 2.17, we only have to verify that we have chosen our neighbourhood large enough. Indeed using (18) we see $r = 9 \operatorname{rch}(M)/128 < \operatorname{rch}(M)/14$.

Thus, from our argument above, the criteria of Theorem 2.17 are satisfied, and $H = \operatorname{pr}_M \circ \iota$ is a homeomorphism. It follows that ι is injective, and since $|\mathcal{A}|$ is compact, ι must be an embedding. The second consequence, that $\operatorname{pr}_M|_{|\mathcal{A}|}$ is a homeomorphism, is now immediate. For the third consequence using Lemma 3.8 and condition (b) gives

$$\delta_M(x) \leq \frac{2L^2}{\operatorname{rch}(M)} \leq \frac{2\varepsilon_0^2 \operatorname{rch}(M)^2}{\operatorname{rch}(M)} \leq 2\varepsilon_0^2 \operatorname{rch}(M).$$

The second inequality follows from Lemma 3.9(ii). □

4 Strong Differential Bounds and Riemannian Manifolds

In this section we will fit the triangulation criteria for Riemannian manifolds in the general framework provided by Theorem 2.17. To do so we first develop triangulation criteria for maps where we have control on the distortion thanks to bounds on the differential of the map F_p , see point (iv) of Proposition 4.5 for details. We believe that this abstract view of the triangulation criteria may be of independent interest.

The triangulation criteria of Theorem 2.17 presented in Sect. 2 have been inspired by the criteria in the Riemannian setting [16, Prop. 16], so the many intermediate results can be used without modification. However one of the conditions in [16, Prop. 16] was too weak and injectivity was in fact not assured. The need to close this gap further motivates this section.

4.1 Strong Differential Bounds

It is not difficult to derive bounds on the metric distortion, assuming one has bounds on the differential of a map. More precisely, for F_p to be 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 ,

$$(1 - \xi)|w| \leq |d(F_p)_u w| \leq (1 + \xi)|w|, \quad (19)$$

see Lemma A.3. Following [16, Prop. 16] we are however interested in a stronger bound on the differential. We require that

$$\|d(F_p)_u - \text{id}\| \leq \xi \quad (20)$$

for all u in the domain of F_p (with $p \in \mathcal{P}$ fixed). We stress that the bound (20) is strictly stronger than (19). It is not difficult to establish that (20) implies that F_p is a ξ -distortion map [16, Lem. 11] on each simplex. However, whereas (19) only constrains how much dF_p can change the magnitude of a vector, (20) 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 2.10):

Lemma 4.1 [16, Lem. 12] *Suppose $\omega \subseteq \mathbb{R}^m$ is a convex set and $F: \omega \rightarrow \mathbb{R}^m$ is a smooth map with a fixed point $p \in \omega$. If*

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

then

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

The added control obtained by the strong bound (20) on the differential enables us to ensure that F_p embeds the boundary of $\text{St}(p)$. This in turn implies that F_p embeds all of $\text{St}(p)$. This follows from the following corollary to Whitney's lemma [24, Lemma AII.15a] (which is demonstrated as Lemmas B.9 and 2.7 in this work). We include the proof here since the proof in [16] 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 4.2 (*smooth on \mathcal{C}*) Given a simplicial complex \mathcal{C} , we say that a map $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ is *smooth on \mathcal{C}* if for each $\sigma \in \mathcal{C}$ the restriction $F|_\sigma$ is smooth.

Lemma 4.3 [16, Lem. 13] *Let \mathcal{C} be a (finite) simplicial complex embedded in \mathbb{R}^m such that $\text{int}(|\mathcal{C}|)$ is nonempty and connected, and $\partial|\mathcal{C}|$ is a compact, connected $(m-1)$ -manifold. Suppose $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ 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 $\text{int}(|\mathcal{C}|)$ 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 2.7, $F(\text{int}(|\mathcal{C}|))$ is open. By the Jordan–Brouwer separation theorem [22, Sect. 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 B.9 implies that $F(\text{int}(|\mathcal{C}|))$ must coincide with the bounded component, and in particular $F(\text{int}(|\mathcal{C}|)) \cap F(\partial|\mathcal{C}|) = \emptyset$, so $F(\text{int}(|\mathcal{C}|))$ is a single connected component.

We need to show that F is injective. First we observe that the set of points in $F(\text{int}(|\mathcal{C}|))$ that have exactly one point in the preimage is nonempty. It suffices to look in a neighbourhood of a point $y \in F(\partial|\mathcal{C}|)$. Choose $y = F(x)$, where x is in the relative interior of $\sigma^{m-1} \subset \partial|\mathcal{C}|$. 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 σ^m be the unique m -simplex that has σ^{m-1} as a face. Then $F^{-1}(V \cap F(|\mathcal{C}|)) \subset \sigma^m$, and it follows that every point in $V \cap \text{int}(|\mathcal{C}|)$ has a unique point in its image. Now the injectivity of F follows from Lemma 2.7. \square

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

Lemma 4.4 [16, Lem. 14] *Suppose $\mathcal{C} = \underline{\text{St}}(\hat{p})$ 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 $\hat{p} \in \text{int}(|\mathcal{C}|)$ (i.e., $\underline{\text{St}}(\hat{p})$ is a full star). If $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ is smooth on \mathcal{C} , and satisfies*

$$\|dF - \text{Id}\| < mt_0 \tag{21}$$

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

If the requirements of Lemma 4.4 are met, then H is a local homomorphism, 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 2.15), and we arrive at the following triangulation result, which can replace the flawed [16, Prop. 16]:

Proposition 4.5 (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 $H: |\mathcal{A}| \rightarrow M$ is a homeomorphism if the following criteria are satisfied:*

(i) **compatible atlases** *There are compatible atlases*

$$\{(\tilde{\mathcal{C}}_p, \hat{\Phi}_p)\}_{p \in \mathcal{P}}, \quad \tilde{\mathcal{C}}_p \subset \mathcal{A}, \quad \text{and} \quad \{(U_p, \phi_p)\}_{p \in \mathcal{P}}, \quad U_p \subset M,$$

for H , with $\tilde{\mathcal{C}}_p = \underline{\text{St}}(p)$ for each $p \in \mathcal{P}$, the vertex set of \mathcal{A} (Definition 2.4).

or the charts in the compatible atlases we have the following:

- (ii) **simplex quality** *For each $p \in \mathcal{P}$, every simplex $\sigma \in \underline{\text{St}}(\hat{p}) = \hat{\Phi}_p(\underline{\text{St}}(p))$ satisfies $s_0 \leq L(\sigma) \leq L_0$ and $t(\sigma) \geq t_0$ (Notation 2.9).*
- (iii) **distortion control** *For each $p \in \mathcal{P}$, the map*

$$F_p = \phi_p \circ H \circ \hat{\Phi}_p^{-1}: |\underline{\text{St}}(\hat{p})| \rightarrow \mathbb{R}^m,$$

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

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

(Definitions 4.2 and B.8).

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

Proof Since the requirements of Lemma 4.4 are met, we only need to show that H is injective. The argument the same as for Lemma 2.16, with minor modifications.

Towards a contradiction, suppose that $H(q) \in H(\sigma)$ and that q is not a vertex of the m -simplex σ . This means there is some $x \in \sigma$ such that $H(x) = H(q)$. Let p be a vertex of σ . The vertex sanity hypothesis (Definition 2.15) implies that $\phi_p \circ H(q)$ must be either outside of $|\underline{\text{St}}(\hat{p})|$, or belong to its boundary. Thus Lemmas 2.12 and 4.1, and the bound on ξ imply that the barycentric coordinate of x with respect to p must be smaller than $1/(m + 1)$: Let $\hat{x} = \widehat{\Phi}_p(x)$, and $\hat{\sigma} = \widehat{\Phi}_p(\sigma)$. Lemma 4.1 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 2.12. 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 $\partial \underline{\text{St}}(\hat{p})$, Lemma 2.12 implies that the barycentric coordinate of $\hat{x} \in \hat{\sigma}$ with respect to \hat{p} must be strictly less than $1/(2m) \leq 1/(m + 1)$. Since $\widehat{\Phi}_p$ preserves barycentric coordinates, and the argument works for any vertex p of σ , we conclude that all the barycentric coordinates of x in σ are strictly less than $1/(m + 1)$. We have reached a contradiction with the fact that the barycentric coordinates of x must sum to 1. □

Remark 4.6 Notice that the constraint on ξ in Proposition 4.5 is only linear in t_0 , whereas it is quadratic in t_0 in Theorem 2.17.

4.2 Triangulations of Riemannian Manifolds

We now apply the theory developed in the previous subsection to Riemannian simplices and triangulations, as we did in [16]. Because of the close relationship with [16] 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 $T_p M$, the resulting vertex set is denoted $\sigma(p)$. A “filled in” Euclidean simplex is denoted $\sigma_{\mathbb{E}}$. The vertex set of \mathcal{A} is denoted by S instead of \mathcal{P} .

As we mentioned in the introduction of this section an extra assumption needs to be added to some statements of [16]. If the following hypothesis is added to Proposition 16 and Theorem 2, all results of [16] hold:

Hypothesis 4.7 (*simple injectivity assumption*) If q is a vertex of \mathcal{A} and $H(q) \in H(\sigma_{\mathbb{E}})$, then q is a vertex of σ .

The injectivity of the map H follows trivially from Hypothesis 4.7, 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., [5]. We provide here corrected statements of the affected results, obtained by replacing [16, Prop. 16] with Proposition 4.5, which uses the vertex sanity assumption, Definition 2.15. Unfortunately, the bound required on the differential in Proposition 4.5 (iv) is qualitatively different from that imposed in [16, Prop. 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 [16, Thm. 2], 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 4.8 ([16, Thm. 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*

$$h = \min \left\{ \frac{t_M}{4}, \frac{\sqrt{\mu_0 t_0}}{6\sqrt{\Lambda}} \right\}. \tag{22}$$

Suppose:

- (i) For every $p \in S$, the vertices of $\underline{\text{St}}(p)$ are contained in $B_M(p, h)$, and the balls $\{B_M(p, h)\}_{p \in S}$ cover M .
- (ii) For every $p \in S$, the restriction of the inverse of the exponential map \exp_p^{-1} to the vertices of $\underline{\text{St}}(p) \subset \mathcal{A}$ defines a piecewise linear embedding of $|\underline{\text{St}}(p)|$ into $T_p M$, realising $\underline{\text{St}}(p)$ as a full star, $\widehat{\underline{\text{St}}}(p)$, such that every simplex $\sigma(p)$ has thickness $t(\sigma(p)) \geq t_0$ and diameter $\mu_0 L_0 \leq L(\sigma(p)) \leq L_0$.
- (iii) For all vertices $p, q \in S$, if $(\exp_p|_{B_M(p,h)})^{-1}(q) \in |\widehat{\underline{\text{St}}}(p)|$, then q is a vertex of $\underline{\text{St}}(p)$.

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

For [16, Prop. 26], the affected argument is in the paragraph preceding the statement of the proposition. New (stronger) bounds on h ensuring a piecewise flat metric on \mathcal{A} are given when we replace that paragraph with (where it is understood that references to Theorem 2 are to the corrected version, Theorem 4.8):

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 $h^2 \Lambda = \mu_0 t_0^2/36$, which gives us $\eta = \mu_0/18 \leq 1/18 < 1$. Thus the requirements of Theorem 2 are sufficient to ensure the existence of a piecewise flat metric on \mathcal{A} , and we obtain:

Proposition 4.9 ([16, Prop. 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*

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

Likewise, for [16, Thm. 3], it is sufficient to impose the new bounds introduced in Theorem 4.8 to obtain the same metric distortion bound:

Theorem 4.10 ([16, Thm. 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 $H: |\mathcal{A}| \rightarrow M$ 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 [5, Thm. 3] employed [16, Thm. 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 4.10, the actual result [5, Thm. 3] stands as stated, since the bound on the sampling density required there is manifestly sufficient to accommodate (22), and the construction via local Delaunay triangulations in the coordinate domains automatically ensures that the vertex sanity criterion is satisfied.

5 Conclusion

We have developed general triangulation criteria for manifolds that are relatively simple to verify for an explicit candidate homeomorphism H . Our two main applications are submanifolds of Euclidean space and Riemannian manifolds, for which we improve on the state-of-the-art.

Our conditions are natural and have been used to prove the correctness of triangulation algorithms [8, 16]. Moreover, variants of these conditions may be derived, so as to obtain conditions that are even more convenient to prove the correctness of other algorithms. In a companion paper with André Lieutier [6, Thm. 5] that was motivated by the analysis of a specific algorithm to triangulate point clouds of Euclidean spaces [15], our criteria (1) (compatible atlases) and (3) (vertex sanity) are replaced by the requirement that the projection on local tangent planes of *all* top dimensional simplices inside a ball of small enough radius have disjoint interiors. Note that the requirement involves all the simplices around a vertex, not simply those of the star of a vertex as we assumed in criterion (1). With this stronger condition, one can exploit some specific properties of submanifolds of Euclidean spaces to optimize the constants governing the density of the point sample, obtaining even more relaxed sampling requirements.

Acknowledgements We thank André Lieutier for discussion and the reviewers for suggestions that improved the exposition.

Funding Open access funding provided by the Austrian Science Fund (FWF).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Metric and Differentiable Distortion Maps

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

A.1 Metric Distortion Maps

The definition of a distortion map (Definition 2.8) makes sense in a more general context: a map $F : (X, d_X) \rightarrow (Y, d_Y)$ between metric spaces is a ξ -distortion map if

$$|d_Y(F(x), F(y)) - d_X(x, y)| \leq \xi d_X(x, y) \quad \text{for all } x, y \in X. \quad (23)$$

Lemma A.1 (inverse and composition of distortion maps)

- (i) If $F : (X, d_X) \rightarrow (Y, d_Y)$ is a ξ -distortion map with $\xi < 1$, then F^{-1} is a $(\xi/(1 - \xi))$ -distortion map.
- (ii) Suppose $F_i : (X_i, d_{X_i}) \rightarrow (X_{i+1}, d_{X_{i+1}})$, $1 \leq i \leq k$, are respectively ξ_i -distortion maps. Then

$$F_k \circ F_{k-1} \circ \dots \circ F_1 : X_1 \rightarrow X_{k+1}$$

is a $(\sum_{W \in \wp(\{k\})} \prod_{i \in W} \xi_i)$ -distortion map, where $\wp(\{k\})$ is the set of nonempty subsets of $\{1, \dots, k\}$. In particular, $F_2 \circ F_1$ is a $(\xi_1 + \xi_2 + \xi_1 \xi_2)$ -distortion map.

Proof (i) Let $u = F(x)$ and $v = F(y)$. Then (23) becomes

$$|d_X(F^{-1}(u), F^{-1}(v)) - d_Y(u, v)| \leq \xi d_X(F^{-1}(u), F^{-1}(v)). \quad (24)$$

But it follows from (24) that $d_X(F^{-1}(u), F^{-1}(v)) \leq d_Y(u, v)/(1 - \xi)$, and plugging this back into (24) yields the result.

(ii) Using the observation that

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

we find

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

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 $G = F_{k-1} \circ \dots \circ F_1$, then $F_k \circ G$ is a $(\xi_k + \eta + \xi_k \eta)$ -distortion map. Now notice that

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

□

A.2 Differentiable Distortion Maps

Although the homeomorphism demonstration that yields Theorem 2.17 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 $F : U \subseteq \mathbb{R}^m \rightarrow \mathbb{R}^m$ is differentiable, then the *differential* of F at x is the linear map defined by

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

where $\alpha : I \subset \mathbb{R} \rightarrow \mathbb{R}^m$ 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 $A : U \subseteq \mathbb{R}^m \rightarrow \mathbb{R}^m$ 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 $s_i(A)$, $1 \leq i \leq m$, ordered such that $s_i(A) \geq s_j(A)$ if $i \leq j$. 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 A.2 *If $F : U \subseteq \mathbb{R}^m \rightarrow \mathbb{R}^m$ 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,

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

so

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

which yields the claimed result. □

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 A.3 *Suppose ω is a convex set, $\omega \subseteq U \subseteq \mathbb{R}^m$, and $F: U \rightarrow \mathbb{R}^m$ is a differentiable map such that $F(\omega) \subseteq V \subseteq F(U)$ 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 $F|_\omega$ is a ξ -distortion map.

Proof Since ω 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 :

$$|F(y) - F(x)| \leq \int_0^1 |dF_{x+t(y-x)}(y - x)| dt \leq (1 + \xi)|y - x|. \tag{25}$$

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

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

Since, for $x, y \in \omega$ the segment $[F(x), F(y)]$ is contained in $F(U)$, we can use the same argument as in (25), but using F^{-1} instead of F , so

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

Therefore, combining with (25) we have

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

Appendix B: Elementary Degree Theory

We recall here some basic ideas in degree theory. Our primary motivation is to facilitate the statement and proof of Lemma 2.7, recovering what we need of a result of Whitney [24, Appendix II, Lemma 15a], without using differentiability assumptions.

We are interested in the degree of continuous maps $F: \overline{\Omega} \rightarrow \mathbb{R}^m$, where $\Omega \subset \mathbb{R}^m$ is an open, bounded, and nonempty domain in \mathbb{R}^m , as can be found in Chapter IV, Sects. 1 and 2 of [22], 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

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

where $\text{sgn}_x(F)$ 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 $F: \overline{\Omega} \rightarrow \mathbb{R}^m$ by means of the Weierstrass approximation theorem (here $\|\cdot\|_\infty$ denotes the supremum norm):

Lemma B.1 [22, Prop. and Defn. IV.2.1] *Let $F: \overline{\Omega} \rightarrow \mathbb{R}^m$ be a continuous map, and let $y \in \mathbb{R}^m \setminus F(\partial\Omega)$. Then there exists a smooth map $G: \overline{\Omega} \rightarrow \mathbb{R}^m$ such that $\|F - G\|_\infty < d_{\mathbb{R}^m}(y, F(\partial\Omega))$. For all such G , the degree $\text{deg}(G, \Omega, y)$ is defined ($y \in \mathbb{R}^m \setminus G(\partial\Omega)$) and is the same, and we define the degree of F by*

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

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

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

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

Lemma B.2 [22, Prop. IV.2.3] *Let $F: \overline{\Omega} \rightarrow \mathbb{R}^m$ 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 B.3 [22, Cor. IV.2.5 (3)] *Given a continuous mapping $F: \overline{\Omega} \rightarrow \mathbb{R}^m$, two disjoint open subsets $\Omega_1, \Omega_2 \subset \Omega$, and a point $y \notin F(\overline{\Omega} \setminus (\Omega_1 \cup \Omega_2))$,*

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

Since no connectedness assumptions are made on the open sets in question, a straightforward inductive argument allows us to strengthen the statement of Lemma B.3:

Lemma B.4 *Suppose $\Omega_1, \dots, \Omega_n$ are mutually disjoint open subsets of the open domain $\Omega \subset \mathbb{R}^m$, and $F: \overline{\Omega} \rightarrow \mathbb{R}^m$ is a continuous map. If $y \notin F(\overline{\Omega} \setminus \bigcup_{i=1}^n \Omega_i)$, then*

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

B.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 $\sigma \subset \mathbb{R}^m$ is positively oriented if its vertices $v_i, 0 \leq i \leq m$, are ordered such that the basis $\{v_1 - v_0, \dots, v_m - v_0\}$ 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 B.5 (*orientation preserving map*) If $\sigma \subset \mathbb{R}^m$ is an m -simplex, we say that a continuous topological embedding $F: \sigma \rightarrow \mathbb{R}^m$ is *orientation preserving*, or *positive*, if $\deg(F, \text{relint}(\sigma), y) = 1$, for any point $y \in \text{int}(F(\sigma))$, otherwise F is *orientation reversing*.

Definition B.6 (*cogent maps*) Suppose \mathcal{C} is a pure m -complex. We call a map $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ *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 $x \in \text{relint}(\sigma)$, where σ is an m -simplex in \mathcal{C} , we define

$$\text{sgn}_x(F) := \deg(F, \text{relint}(\sigma), F(x)). \tag{26}$$

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

Lemma B.7 (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 $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ is cogent with respect to \mathcal{C} , then for any $y \in \mathbb{R}^m \setminus F(|\mathcal{C}^{m-1}|)$ (recall that $\partial|\mathcal{C}| \subseteq |\mathcal{C}^{m-1}|$ [4, Lemmas 3.6 and 3.7]),*

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

A very minor modification to the proof of Lemma B.1 would allow us to sharpen the bound on $\|F - G\|_\infty$ in our context, so that

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

and Lemma B.7 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 B.4.

Proof of Lemma B.7 The preimage of y is a finite set of points: $F^{-1}(y) = \{x_i\}$, $i \in \{1, \dots, n\}$. For each i , let σ_i be the m -simplex that contains x_i in its interior, and set $\Omega_i = \text{relint}(\sigma_i)$. Then the statement follows from Lemma B.4 and the definition (26) of $\text{sgn}_{x_i}(F)$. □

Definition B.8 (*simplexwise positive*) A map is *simplexwise positive* if it is cogent and its restriction to any m -simplex is orientation preserving.

Lemmas B.2 and B.7 yield

Lemma B.9 *If \mathcal{C} is a finite pure m -complex embedded in \mathbb{R}^m , and $F: |\mathcal{C}| \rightarrow \mathbb{R}^m$ 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 B.10 Whitney’s Lemma [24, Appx. II, Lem. 15] is a combination of Lemmas B.9 and 2.7, 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 $\overline{\Omega} \rightarrow \mathbb{R}^m$ 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 two or one m -simplices (those that are the face of only one m -simplex define the boundary complex).

Brouwer's original exposition of degree theory [12] 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 [22, p. 23], so that approach would also allow us to recover the pseudomanifold aspect of Whitney's lemma.

References

1. Amenta, N., Bern, M.: Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* **22**(4), 481–504 (1999)
2. Amenta, N., Choi, S., Dey, T.K., Leekha, N.: A simple algorithm for homeomorphic surface reconstruction. *Int. J. Comput. Geom. Appl.* **12**(2), 125–141 (2002)
3. Boissonnat, J.-D., Chazal, F., Yvinec, M.: *Geometric and Topological Inference*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge (2018)
4. Boissonnat, J.-D., Dyer, R., Ghosh, A.: The stability of Delaunay triangulations. *Int. J. Comput. Geom. Appl.* **23**(4–5), 303–333 (2013)
5. Boissonnat, J.-D., Dyer, R., Ghosh, A.: Delaunay triangulation of manifolds. *Found. Comput. Math.* **18**(2), 399–431 (2018)
6. Boissonnat, J.-D., Dyer, R., Ghosh, A., Lieutier, A., Wintraecken, M.: Local conditions for triangulating submanifolds of Euclidean space. *Discrete Comput. Geom.* **66**(2), 666–686 (2021)
7. Boissonnat, J., Dyer, R., Ghosh, A., Wintraecken, M.: Local criteria for triangulation of manifolds (2018). [arXiv:1803.07642](https://arxiv.org/abs/1803.07642)
8. Boissonnat, J.-D., Ghosh, A.: Manifold reconstruction using tangential Delaunay complexes. *Discrete Comput. Geom.* **51**(1), 221–267 (2014)
9. Boissonnat, J.-D., Kachanovich, S., Wintraecken, M.: Triangulating submanifolds: an elementary and quantified version of Whitney's method. *Discrete Comput. Geom.* **66**(1), 386–434 (2021)
10. Boissonnat, J.-D., Lieutier, A., Wintraecken, M.: The reach, metric distortion, geodesic convexity and the variation of tangent spaces. *J. Appl. Comput. Topol.* **3**(1–2), 29–58 (2019)
11. Boissonnat, J.-D., Oudot, S.: Provably good sampling and meshing of surfaces. *Graph. Models* **67**(5), 405–451 (2005)
12. Brouwer, L.E.J.: Über Abbildung von Mannigfaltigkeiten. *Math. Ann.* **71**(4), 598 (1912)
13. Cairns, S.S.: On the triangulation of regular loci. *Ann. Math.* **35**(3), 579–587 (1934)
14. Cheng, S.-W., Dey, T.K., Ramos, E.A.: Manifold reconstruction from point samples. In: 16th Annual ACM-SIAM Symposium on Discrete Algorithms (Vancouver 2005), pp. 1018–1027. ACM, New York (2005)
15. Cohen-Steiner, D., Lieutier, A., Vuillamy, J.: Lexicographic optimal homologous chains and applications to point cloud triangulations. In: 36th International Symposium on Computational Geometry. Leibniz Int. Proc. Inform., vol. 164, # 32. Leibniz-Zent. Inform., Wadern (2020)
16. Dyer, R., Vegter, G., Wintraecken, M.: Riemannian simplices and triangulations. *Geom. Dedicata* **179**, 91–138 (2015)
17. Dyer, R., Zhang, H., Möller, T.: Surface sampling and the intrinsic Voronoi diagram. *Comput. Graph. Forum* **27**(5), 1393–1402 (2008)
18. Edelsbrunner, H., Shah, N.R.: Triangulating topological spaces. *Int. J. Comput. Geom. Appl.* **7**(4), 365–378 (1997)
19. Federer, H.: Curvature measures. *Trans. Am. Math. Soc.* **93**(3), 418–491 (1959)
20. Lundell, A.T., Weingram, S.: *The Topology of CW Complexes*. The University Series in Higher Mathematics. Van Nostrand Reinhold, New York (1969)
21. Munkres, J.R.: *Elementary Differential Topology*, 2nd edn. Princeton University Press, Princeton (1968)
22. Outerelo, E., Ruiz, J.M.: *Mapping Degree Theory*, vol. 108. American Mathematical Society, Providence (2009)
23. Whitehead, J.H.C.: On C^1 -complexes. *Ann. Math.* **41**(4), 809–824 (1940)
24. Whitney, H.: *Geometric Integration Theory*. Princeton University Press, Princeton (1957)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Jean-Daniel Boissonnat¹ · Ramsay Dyer¹  · Arijit Ghosh² ·
Mathijs Wintraecken³ 

Jean-Daniel Boissonnat
jean-daniel.boissonnat@inria.fr

Ramsay Dyer
yasmarr@gmail.com

Arijit Ghosh
arijitiitkgpster@gmail.com

Mathijs Wintraecken
m.h.m.j.wintraecken@gmail.com

¹ INRIA Sophia Antipolis, Université Côte d'Azur, Sophia Antipolis, France

² Indian Statistical Institute, Kolkata, India

³ IST Austria, Klosterneuburg, Austria