

Local Criteria for Triangulation of Manifolds

Jean-Daniel Boissonnat

Université Côte d'Azur, Inria
Sophia Antipolis, France
Jean-Daniel.Boissonnat@inria.fr

Ramsay Dyer

Université Côte d'Azur, Inria
Sophia Antipolis, France
Ramsay.Dyer@inria.fr
 <https://orcid.org/0000-0001-5083-7181>

Arijit Ghosh¹

Indian Statistical Institute
Kolkata, India
arijitiitkgpster@gmail.com

Mathijs Wintraecken

Université Côte d'Azur, Inria
Sophia Antipolis, France
Mathijs.Wintraecken@inria.fr

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.

2012 ACM Subject Classification Mathematics of computing \rightarrow Geometric topology, Mathematics of computing \rightarrow Mesh generation

Keywords and phrases manifold, simplicial complex, homeomorphism, triangulation

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.9

Related Version A full version of the paper is available at <http://arxiv.org/abs/1803.07642>

Funding This work has been partially 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).

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

¹ 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).



say that \mathcal{A} triangulates M .

The purpose of this paper is to present criteria which ensure that a candidate map H is indeed a homeomorphism. This work is motivated by earlier investigations into the problem of algorithmically constructing a complex that triangulates a given manifold [6, 4]. It complements and is closely related to recent work that investigates a particular natural example of such a map [11].

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 16, are based on metric properties of H within “compatible” coordinate charts (Definition 4). The Euclidean metric in the local coordinate chart is central to the analysis, but no explicit metric on $|\mathcal{A}|$ or M is involved, and no explicit assumption of differentiability is required of H or M . However, our only examples that meet the required local criteria are in the differentiable setting. We do not know whether or not our criteria for homeomorphism implicitly imply that M admits a differentiable structure. They do imply that \mathcal{A} is piecewise linear (admits an atlas with piecewise linear transition functions).

Relation to other work

The first demonstrations that differentiable manifolds *can* always be triangulated were constructive. Cairns [9] 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 [15, 14], but it is not easily adapted to provide triangulation guarantees for complexes constructed by other algorithms.

An alternative approach was developed by Whitney [16] 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 [13] defined the restricted Delaunay complex of a subset M of Euclidean space as the nerve of the Voronoi diagram on M when the ambient Euclidean metric is used. They showed that if M is a compact manifold, then the restricted Delaunay complex is homeomorphic to M when the Voronoi diagram satisfies the *closed ball property (cbp)*: Voronoi faces are closed topological balls of the appropriate dimension.

Using the cbp, Amenta and Bern [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 cbp is an elegant topological result that applies in more general contexts. For example, it has been used to establish conditions for intrinsic Delaunay triangulations of surfaces [12], and Cheng et al. [10] have indicated how it can be applied for establishing weighted restricted Delaunay triangulations of smooth submanifolds of arbitrary dimension in Euclidean space.

However, the cbp is only applicable to Delaunay-like complexes that can be realized as the nerve of some kind of Voronoi diagram on the manifold. Thus, for example, it does not necessarily apply to the tangential Delaunay complex constructed by Boissonnat and Ghosh [6]. Secondly, even when a Delaunay-like complex is being constructed, it can be difficult to directly verify the properties of the associated Voronoi structure; sampling criteria and conditions on the complex under construction are desired, but may not be easy to obtain from the cbp. A third deficiency of the cbp is that, although it can establish that a complex \mathcal{A} triangulates the manifold M , it does not provide a specific triangulation $H : |\mathcal{A}| \rightarrow M$. Such a correspondence allows us to compare *geometric* properties of $|\mathcal{A}|$ and M .

In [6] 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 [8], guarantees a triangulation of a surface by any complex which satisfies a few easily verifiable properties. However, the argument relies heavily on the the codimension being 1.

If a set of vertices is contained within a sufficiently small neighbourhood on a Riemannian manifold, barycentric coordinates can be defined. So there is a natural map from a Euclidean simplex of the appropriate dimension to the manifold, assuming a correspondence between the vertices of the simplex and those on the manifold. Thus when a complex \mathcal{A} is appropriately defined with vertices on a Riemannian manifold M , there is a natural *barycentric coordinate map* $|\mathcal{A}| \rightarrow M$. In [11], conditions are presented which guarantee that this map is a triangulation. Although this map is widely applicable, the intrinsic criteria can be inconvenient, for example, in the setting of Euclidean submanifold reconstruction, and furthermore the closest-point projection map may be preferred for triangulation in that setting.

The argument in [11] is based on a general result [11, Proposition 16] for establishing that a given map is a triangulation of a differentiable manifold. However, the criteria include a bound on the differential of the map, which is not easy to obtain. The analysis required to show that the closest-point projection map meets this bound is formidable, and this motivated the current alternate approach. We have relaxed this constraint to a much more easily verifiable bound on the metric distortion of the map when viewed within a coordinate chart.

The sampling criteria for submanifolds imposed by our main result applied to the closest-point projection map (Theorem 17) are the most relaxed that we are aware of. The result could be applied to improve the sampling guarantees of previous works, e.g., [10, 6].

In outline, the argument we develop here is the same as that of [2], but extends the result to apply to abstract manifolds of arbitrary dimension and submanifolds of \mathbb{R}^N of arbitrary codimension. We first show that the map H is a local homeomorphism, and thus a covering map, provided certain criteria are met. Then injectivity is ensured when we can demonstrate that each component of M contains a point y such that $H^{-1}(y)$ is a single point. A core technical lemma from Whitney [16, Appendix II Lemma 15a] still lies at the heart of our argument.

Overview

The demonstration is developed abstractly without explicitly defining the map H . We assume that it has already been established that the restriction of H to any Euclidean simplex in $|\mathcal{A}|$ is an embedding. This is a nontrivial step that needs to be resolved from the specific properties of a particular choice of H . The criteria for local homeomorphism apply in a common coordinate chart (for $|\mathcal{A}|$ and M), and relate the size and quality of the simplices

with the metric distortion of H , viewed in the coordinate domain. The requirement that leads to injectivity is also expressed in a local coordinate chart; it essentially demands that the images of vertices behave in a natural and expected way.

After demonstrating the main result (Theorem 16), we discuss its application to submanifolds of Euclidean space: Theorem 17 presents criteria that ensure that the closest-point projection map from a complex provides a triangulation of a submanifold. The details can be found in the full version of this work [5].

2 The setting and notation

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 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* 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 $|\mathcal{C}|$.

► **Notation 2** (topology). If $A \subseteq \mathbb{R}^n$, then the topological closure, interior, and boundary of A are denoted respectively by \bar{A} , $\text{int}(A)$, and $\partial A = \bar{A} \setminus \text{int}(A)$. We denote by $B_{\mathbb{R}^m}(c, r)$ the open ball in \mathbb{R}^m of radius r and centre c .

► **Notation 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 \hat{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 \hat{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 the diagram on page 5).

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

1. There is a coordinate atlas $\{(U_p, \phi_p)\}_{p \in \mathcal{P}}$ for M , where the index set \mathcal{P} is the set of vertices of \mathcal{A} and each set U_p is connected.
2. For each $p \in \mathcal{P}$, $H(|\text{St}(p)|) \subset U_p$. Also, the secant map of $\Phi_p = \phi_p \circ H|_{|\text{St}(p)|}$ defines a piecewise linear embedding of $|\text{St}(p)|$ into \mathbb{R}^m . We denote this secant map by $\hat{\Phi}_p$. By definition, $\hat{\Phi}_p$ preserves the barycentric coordinates within each simplex, and thus the collection $\{(|\text{St}(p)|, \hat{\Phi}_p)\}_{p \in \mathcal{P}}$ provides a piecewise linear atlas for \mathcal{A} .

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


For convenience, we define $\hat{p} = \widehat{\Phi}_p(p)$, so that $\widehat{\Phi}_p(\text{St}(p)) = \text{St}(\hat{p})$. We will work within the compatible local coordinate charts. Thus we are studying a map of the form

$$F_p: |\text{St}(\hat{p})| \subset \mathbb{R}^m \rightarrow \mathbb{R}^m,$$

where

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

as shown in the following diagram (recall M and \mathcal{A} are of dimension m):



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

3 Local homeomorphism

Our goal is to ensure that there is some open $V_p \subset |\text{St}(\hat{p})|$ such that $F_p|_{V_p}$ is an embedding and that the sets $\tilde{V}_p = \widehat{\Phi}_p^{-1}(V_p)$ are sufficiently large to cover $|\mathcal{A}|$. This will imply that H is a local homeomorphism. Indeed, if V_p is embedded by F_p , then \tilde{V}_p is embedded by $H|_{\tilde{V}_p} = \phi_p^{-1} \circ F_p \circ \widehat{\Phi}_p|_{\tilde{V}_p}$, since ϕ_p and $\widehat{\Phi}_p$ are both embeddings. Since $|\mathcal{A}|$ is compact, Brouwer’s invariance of domain, together with Observation 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 require that it be *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 [5, Appendix A], we can use degree theory to talk about orientation-preserving maps even if the maps are not differentiable. The other requirement we have for F_p is that when it is restricted to an m -simplex it does not distort distances very much, as discussed below.

The local homeomorphism demonstration is based on Lemma 6 below, which is a particular case of an observation made by Whitney [16, Appendix II Lemma 15a]. Whitney demonstrated a more general result from elementary first principles. The proof we give here is roughly the same as Whitney’s, except that we exploit elementary degree theory, as discussed in [5, Appendix A], in order to avoid the differentiability assumptions Whitney made.

In the statement of the lemma, \mathcal{C}^{m-1} refers to the $(m-1)$ -skeleton of the complex \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 6** (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 underlying space of the boundary complex (see, e.g., [3, Lemmas 3.6, 3.7]): $\partial|\mathcal{C}| = |\partial\mathcal{C}|$. Let $\Omega = |\mathcal{C}| \setminus |\partial\mathcal{C}|$. Since F is simplexwise positive, and $F(V)$ lies within a connected component of $\mathbb{R}^m \setminus F(\partial\Omega)$, the fact that $F^{-1}(y)$ is a single point implies that $F^{-1}(w)$ is a single point for any $w \in F(V) \setminus F(|\mathcal{C}^{m-1}|)$ (see [5, Lemma 49]). We need to show that F is also injective on $V \cap |\mathcal{C}^{m-1}|$.

Let $\sigma \in \mathcal{C}^{m-1} \setminus \partial\mathcal{C}$, and observe that the open star $\text{st}(\sigma)$ (Notation 1) is open in \mathbb{R}^m . We now show that $F(\text{st}(\sigma))$ is open. 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 underlying space of our usual star: $\overline{\text{st}(\tau)} = |\text{St}(\tau)|$, By [5, Lemma 49], every point in $U \setminus F(|\text{St}(\tau)^{m-1}|)$ 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 it follows from Brouwer's invariance of domain that $F|_V$ is an embedding. ◀

Our strategy for employing Lemma 6 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 $V_p \subset |\text{St}(\hat{p})|$ is not intersected by the image of the boundary of $|\text{St}(\hat{p})|$, i.e., we will establish that $F_p(V_p) \cap F_p(|\partial\text{St}(\hat{p})|) = \emptyset$. We need to also establish that there is a point y in $F_p(V_p) \setminus F_p(|\text{St}(\hat{p})^{m-1}|)$ such that $F^{-1}(y)$ is a single point. The metric distortion bound will help us here as well.

► **Definition 7** (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

$$||F(x) - F(y)| - |x - y|| \leq \xi |x - y|. \quad (2)$$

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

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

and it is clear that when $\xi < 1$, a ξ -distortion map is a bi-Lipschitz map. For our purposes the metric distortion constant ξ is more convenient than a bi-Lipschitz constant. It is easy

to show that if F is a ξ -distortion map, with $\xi < 1$, then F is a homeomorphism onto its image, and F^{-1} is a $\frac{\xi}{1-\xi}$ -distortion map (see [5, Lemma 19(1)]).

Assuming that $F_p|_{\sigma}$ is a ξ -distortion map for each m -simplex $\sigma \in \text{St}(\hat{p})$, we can bound how much it displaces points. Specifically, for any point $x \in |\text{St}(\hat{p})|$, 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 8** (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 $\frac{a}{mL}$, 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 9** (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

$$p_i^\top x = \frac{1}{2} (|x|^2 + |p_i|^2 - |x - p_i|^2), \tag{3}$$

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 (3). 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 [3, Lemma 2.4] we have $\|(P^\top)^{-1}\| \leq (\sqrt{m}tL)^{-1}$, and the stated bound follows. ◀

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



■ **Figure 1** The open set V_p .

We assume that F_p is a ξ -distortion map on each simplex.

The idea is to show that F_p is an embedding on V_p . In order to employ the simplexwise positive embedding lemma (Lemma 6), we need to establish that there is a point in $V_p \setminus |\text{St}(\hat{p})|^{m-1}$ that is not mapped to the image of any other point in $|\text{St}(\hat{p})|$. 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 10** (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*

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

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 9 and the constraint (4) we arrive at a contradiction:

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

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

► **Lemma 11** (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

$$\begin{aligned} |x - y| &= \sum_{i=0}^n |x_i - y_i| \geq \sum_{i=0}^n a_{\hat{p}}(\sigma_i) |\lambda_{\sigma_i, \hat{p}}(x_i) - \lambda_{\sigma_i, \hat{p}}(y_i)| \geq a_0 \sum_{i=0}^n (\lambda_{\sigma_i, \hat{p}}(x_i) - \lambda_{\sigma_i, \hat{p}}(y_i)) \\ &= a_0 (\lambda_{\sigma_0, \hat{p}}(x) - \lambda_{\sigma_n, \hat{p}}(y)) = a_0 \lambda_{\sigma_0, \hat{p}}(x) \geq a_0 \alpha. \end{aligned}$$

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

Lemma 11 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 10, but there we were only concerned with the barycentre of the largest simplex, so the relative sizes of the simplices were not relevant as they are here (compare the bounds (4) and (5)).

► **Lemma 12** (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*

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

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 $\frac{1}{m+1} - \delta$, 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 $\frac{1}{m+1} - \delta$, and $y \in \tau \in \partial \text{St}(\hat{p})$, then Lemmas 9 and 11 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 (5) when $\delta > 0$ satisfies

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

When inequality (5) (and therefore also inequality (4)) is satisfied, we can employ the embedding lemma (Lemma 6) to guarantee that V_p is embedded:

► **Lemma 13** (local homeomorphism). *Suppose M is an m -manifold and \mathcal{A} is a simplicial complex with vertex set \mathcal{P} . A map $H: |\mathcal{A}| \rightarrow M$ is a covering map if the following criteria are satisfied:*

1. **manifold complex** \mathcal{A} is a compact m -manifold complex (without boundary).
2. **compatible atlases** There are compatible atlases for H (Definition 4).
3. **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 8).
4. **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 (Definition 7) with

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

4 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 Observation 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 coordinate patches $\text{St}(\hat{p})$, and it is with respect to these patches that the vertices behave well.

9:10 Local Criteria for Triangulation of Manifolds

► **Definition 14** (vertex sanity). If $H: |\mathcal{A}| \rightarrow M$ has compatible atlases (Definition 4), then H exhibits *vertex sanity* if: for all vertices $p, q \in \mathcal{P}$, if $\phi_p \circ H(q) \in |\underline{\text{St}}(\hat{p})| = \widehat{\Phi}_p(|\underline{\text{St}}(p)|)$, then q is a vertex of $\underline{\text{St}}(p)$.

Together with the distortion bounds that are imposed on F_p , Definition 14 ensures that the image of a vertex cannot lie in the image of a simplex to which it does not belong:

► **Lemma 15** (injectivity). *If $H: |\mathcal{A}| \rightarrow M$ satisfies the hypotheses of Lemma 13 as well as Definition 14, 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 14) implies that $\phi_p \circ H(q)$ must be either outside of $|\underline{\text{St}}(\hat{p})|$, or belong to its boundary. Thus Lemmas 11 and 9, and the bound on ξ from Lemma 13(3) imply that the barycentric coordinate of x with respect to p must be smaller than $\frac{1}{m+1}$: Let $\hat{x} = \widehat{\Phi}_p(x)$, and $\hat{\sigma} = \widehat{\Phi}_p(\sigma)$. Lemma 9 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 11. Since $F_p(\hat{x}) = \phi_p \circ H(x)$ is at least as far away from \hat{x} as $\partial \underline{\text{St}}(\hat{p})$, Lemma 11 implies that the barycentric coordinate of $\hat{x} \in \hat{\sigma}$ with respect to \hat{p} must be no larger than $\frac{1}{2(m+1)}$. Since $\widehat{\Phi}_p$ preserves barycentric coordinates, and the argument works for any vertex p of σ , we conclude that all the barycentric coordinates of x in σ are strictly less than $\frac{1}{m+1}$. We have reached a contradiction with the fact that the barycentric coordinates of x must sum to 1. ◀

5 Main result

To recap, Lemmas 13 and 15 yield the following triangulation result. In the bound on ξ from Lemma 13(3), we replace the factor $\frac{m}{m+1}$ with $\frac{1}{2}$, the lower bound attained when $m = 1$.

► **Theorem 16** (triangulation). *Suppose M is an m -manifold, and \mathcal{A} is a simplicial complex with vertex set \mathcal{P} . A map $H: |\mathcal{A}| \rightarrow M$ is a homeomorphism if the following criteria are satisfied:*

1. **manifold complex** \mathcal{A} is a compact m -manifold complex (without boundary).
2. **compatible atlases** There are compatible atlases for H (Definition 4):

$$\{(\underline{\text{St}}(p), \widehat{\Phi}_p)\}_{p \in \mathcal{P}}, \quad \underline{\text{St}}(p) \subset \mathcal{A}, \quad \text{and} \quad \{(U_p, \phi_p)\}_{p \in \mathcal{P}}, \quad U_p \subset M.$$

3. **simplex quality** For each $p \in \mathcal{P}$, every simplex $\sigma \in \underline{\text{St}}(\hat{p}) = \widehat{\Phi}_p(\underline{\text{St}}(p))$ satisfies $s_0 \leq L(\sigma) \leq L_0$ and $t(\sigma) \geq t_0$ (Notation 8).
4. **distortion control** For each $p \in \mathcal{P}$, the map

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

when restricted to any m -simplex in $\underline{\text{St}}(\hat{p})$, is an orientation-preserving ξ -distortion map (Definition 7) with

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

5. **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 $\underline{\text{St}}(p)$.



► **Remark.** The constants L_0 , s_0 , and t_0 that constrain the simplices in the local complex $\text{St}(\hat{p})$, and the metric distortion of F_p in Theorem 16 can be considered to be local, i.e., they may depend on $p \in \mathcal{P}$. This result applies for any dimension m , but for $m \leq 2$, triangulation criteria already exist which demand neither a lower bound on the size (s_0), nor on the quality (t_0) of the simplices [2, 8].

Application: submanifolds of Euclidean space

As a specific application of Theorem 16, 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 [6]. 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 general 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 16 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 [5, Section 4.3], 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 employ the closest point projection map: The *medial axis*, $\text{ax}(M)$ is the set of points $x \in \mathbb{R}^N$ that have more than one closest point on M , and $\overline{\text{ax}}(M)$ is its closure. The open set $U_M = \mathbb{R}^N \setminus \overline{\text{ax}}(M)$ contains M and each point in it has a unique closest point on M , so the *closest-point projection map*, $\text{pr}_M: U_M \rightarrow M$, is well-defined (see, e.g., [5, Section 4.1]). We define $H = \text{pr}_M \circ \iota$.

The *local feature size* is the function $\text{lfs}: M \rightarrow \mathbb{R}_{>0}$ defined by $\text{lfs}(x) = d_{\mathbb{R}^N}(x, \overline{\text{ax}}(M))$. It is easily verified that this function is continuous. It plays an important role as a sizing function that governs the density of sample points (vertices of \mathcal{A}) that are required to construct a triangulation.

As demanded by Definition 4, for each $p \in \mathcal{P}$ the coordinate map $\hat{\Phi}_p$ for \mathcal{A} is the secant map of $\phi_p \circ H$ restricted to $|\text{St}(p)|$, and since $\text{pr}_{T_p M}$ is already a linear map, and pr_M is the identity on the vertices, this means $\hat{\Phi}_p = \text{pr}_{T_p M} \circ \iota|_{|\text{St}(p)|}$ (recall that the secant map $\hat{\Phi}_p$ is defined by the action of $\phi_p \circ H|_{|\text{St}(p)|}$ on the vertices alone).

In order to employ Theorem 16, we need to analyze the projection maps to obtain metric distortion bounds on the restriction to m -simplices. For the projection onto the tangent spaces, this analysis [5, Section 4.3] yields bounds on the size of the simplices and coordinate neighbourhoods such that the compatible atlas criterion will be automatically satisfied. Thus the compatible atlas criterion does not appear explicitly in the statement of Theorem 17; it has been subsumed by the proof.

9:12 Local Criteria for Triangulation of Manifolds

The metric distortion of pr_M restricted to an m -simplex is analyzed in [5, Section 4.4], using recent bounds on the angle between nearby tangent spaces [7]. The distortion bounds on the different projection maps yield a distortion bound on their composition, F_p . In order to obtain a bound in this fashion, it is necessary to choose a metric on M . We employed the metric of the ambient space \mathbb{R}^N restricted to M , rather than the intrinsic metric of geodesic distances.

The simplex quality and sampling criteria of Theorem 17(b) are tailored to ensure that the metric distortion criterion of Theorem 16(4) is met. Thus an explicit metric distortion condition does not appear in the statement of Theorem 17. Some adjustment was also made so that the simplex quality conditions in Theorem 17(b) refer to the ambient simplices of \mathcal{A} , rather than the projected simplices in the tangent spaces, as required by Theorem 16.

We then arrive at the following specific incarnation of Theorem 16:

► **Theorem 17** (triangulation for submanifolds). *Let $M \subset \mathbb{R}^N$ be a compact C^2 -continuous 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., $\iota|_{|\underline{\text{St}}(p)|}$ is an embedding for each vertex p .*

If:

- (a) **manifold complex** *For each vertex $p \in \mathcal{P}$, the projection $\text{pr}_{T_p M}|_{\iota(|\underline{\text{St}}(p)|)}$ is an embedding and p lies in the interior of $\text{pr}_{T_p M}(\iota(|\underline{\text{St}}(p)|))$.*
- (b) **simplex quality** *There are constants $0 < t_0 \leq 1$, $0 < \mu_0 \leq 1$, and $\epsilon_0 > 0$ such that for each simplex $\sigma \in \iota(\mathcal{A})$, and each vertex $p \in \sigma$,*

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

- (c) **vertex sanity** *For any vertices $p, q \in \mathcal{P}$, if $q \in U_p = B_{\mathbb{R}^N}(p, r) \cap M$, where $r = \text{lfs}(p)/15$, then $\text{pr}_{T_p M}(q) \in \text{pr}_{T_p M}(\iota(|\underline{\text{St}}(p)|))$ if and only if q is a vertex of $\underline{\text{St}}(p)$.*

Then:

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

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

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

Using the reach, which is a global bound on the local feature size, $\text{rch}(M) = \inf_{x \in M} \text{lfs}(x)$, we obtain the following variation of Theorem 17, which is a corollary in the sense that it follows from essentially the same proof, even though it does not follow from the statement of Theorem 17.

► **Corollary 18.** *If the conditions (b) and (c) in Theorem 17 are replaced by*

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

$$t(\sigma) \geq t_0, \quad \mu_0 \epsilon_0 \text{rch}(M) \leq L(\sigma) \leq \epsilon_0 \text{rch}(M), \quad \epsilon_0 \leq \frac{\mu_0^{\frac{1}{2}} t_0^2}{16}.$$

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

then the conclusions of Theorem 17 hold, and consequence (3) can be tightened to:

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

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

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

Discussion

Theorem 16 can be applied to any map that could serve as a triangulation. However, we only know of two specific examples: the closest-point projection map featured in Theorem 17, and the barycentric coordinate map [11] mentioned in the introduction.

Conditions guaranteeing that the barycentric coordinate map is a triangulation were already established in [11], however the generic triangulation theorem [11, Proposition 16] on which the result is based had a flaw (discovered during the preparation of this current work), so that injectivity of the map is not guaranteed. This flaw has been repaired, using the vertex sanity criterion (Definition 14); this is described in [5, Appendix C], where the corrected statement of the theorem for the barycentric coordinate map can be found [5, Theorem 60].

Theorem 16 also leads to triangulation criteria for the barycentric coordinate map. However, the bound on the diameter of the simplices (the sampling radius) is proportional to the square of the thickness bound t_0 when Theorem 16 is employed, but it is only linear in t_0 when [11, Proposition 16] is used. This indicates that there is also room to improve the bound on ϵ_0 in Theorem 17(b) from quadratic to linear in t_0 .

References

- 1 N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22(4):481–504, 1999.
- 2 N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Int. J. Computational Geometry and Applications*, 12(2):125–141, 2002.
- 3 J.-D. Boissonnat, R. Dyer, and A. Ghosh. The stability of Delaunay triangulations. *International Journal of Computational Geometry & Applications*, 23(4-5):303–333, 2013. (arXiv:1304.2947).
- 4 J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay triangulation of manifolds. *Foundations of Computational Mathematics*, 2017. (arXiv:1311.0117).
- 5 J.-D. Boissonnat, R. Dyer, A. Ghosh, and M. Wintraecken. Local criteria for triangulation of manifolds. Technical Report 1803.07642, arXiv, 2017. URL: <http://arxiv.org/abs/1803.07642>.
- 6 J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. *Discrete and Computational Geometry*, 51(1):221–267, 2014.
- 7 J.-D. Boissonnat, A. Lieutier, and M. Wintraecken. The reach, metric distortion, geodesic convexity and the variation of tangent spaces. Technical Report hal-01661227, Inria, Sophia-Antipolis, 2017. Accepted for SoCG 2018. URL: <https://hal.inria.fr/hal-01661227>.
- 8 J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, 2005.
- 9 S. S. Cairns. On the triangulation of regular loci. *Annals of Mathematics. Second Series*, 35(3):579–587, 1934.
- 10 S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA*, pages 1018–1027, 2005.

9:14 Local Criteria for Triangulation of Manifolds

- 11 R. Dyer, G. Vegter, and M. Wintraecken. Riemannian simplices and triangulations. *Geometriae Dedicata*, 179:91–138, 2015.
- 12 R. Dyer, H. Zhang, and T. Möller. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum (Special Issue of Symp. Geometry Processing)*, 27(5):1393–1402, 2008.
- 13 H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Int. J. Comput. Geometry Appl.*, 7(4):365–378, 1997.
- 14 J. R. Munkres. *Elementary differential topology*. Princeton University press, second edition, 1968.
- 15 J. H. C. Whitehead. On C^1 -complexes. *Ann. of Math*, 41(4):809–824, 1940.
- 16 H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957.