



Project number IST-25582

CGL

Computational Geometric Learning

Delaunay triangulation of manifolds

STREP

Information Society Technologies

Period covered: November 1, 2012–October 31, 2013
Date of preparation: November 12, 2013
Start date of project: November 1, 2010
Duration: 3 years
Project coordinator name: Joachim Giesen (FSU)
Project coordinator organisation: Friedrich-Schiller-Universität Jena
Jena, Germany

Delaunay triangulation of manifolds

Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh*

November 26, 2013

Abstract

We present an algorithmic framework for producing Delaunay triangulations of manifolds. The input to the algorithm is a set of sample points together with coordinate patches indexed by those points. The transition functions between nearby coordinate patches are required to be bi-Lipschitz with a constant close to 1. The primary novelty of the framework is that it can accommodate abstract manifolds that are not presented as submanifolds of Euclidean space. The output is a manifold simplicial complex that is the Delaunay complex of a perturbed set of points on the manifold. The guarantee of a manifold output complex demands no smoothness requirement on the transition functions, beyond the bi-Lipschitz constraint. In the smooth setting, when the transition functions are defined by common coordinate charts, such as the exponential map on a Riemannian manifold, the output manifold is homeomorphic to the original manifold, when the sampling is sufficiently dense.

*This research has been partially supported by the 7th Framework Programme for Research of the European Commission, under FET-Open grant number 255827 (CGL Computational Geometry Learning).

Contents

1	Introduction	4
2	Review of the algorithm for flat manifolds	5
3	Overview of the extended algorithm	8
3.1	Manifolds represented by transition functions	9
3.2	The setting and input data	9
3.3	Outline of the analysis	11
4	Output quality guarantees	14
4.1	The output complex is a Delaunay complex	14
4.2	A metric on the output complex	15
4.3	Homeomorphism guarantee in the smooth case	15
4.4	Discussion	16
A	Details of the analysis	16
A.1	Simplex distortion	16
A.2	Circumcentres and distortion maps	20
A.3	The size of the domains	23
A.4	Hoop distortion	25
	Bibliography	27

1 Introduction

We present an algorithm for triangulating a manifold that may be presented abstractly in terms of a coordinate atlas. In particular, there is no requirement that the manifold be an embedded submanifold of Euclidean space. This is the principal novelty in our approach. Earlier work [LL00] proposed a purely intrinsic algorithm for constructing Delaunay triangulations of manifolds, but this was based on an existence result that has been shown to be fundamentally flawed [BDG13a]. In fact, the point set must be bounded away from degeneracy by an amount that is not arbitrarily small. A refinement algorithm for constructing points that admit a Delaunay triangulation was presented in the latter work, but it applies only to sub-manifolds of Euclidean space. In this work we construct a Delaunay triangulation on abstract Riemannian manifolds.

Given a dense sample set on the manifold, we consider a coordinate chart centred at each sample point. A Delaunay triangulation is made locally, using the Euclidean metric of the coordinate chart, and these local triangulations assemble to create a global manifold simplicial complex. The idea has been proposed previously in the context of surfaces [GKS00], but it suffers from a fundamental problem: no matter how densely the surface is sampled, a Delaunay triangulation in one coordinate chart need not be isomorphic to the Delaunay triangulation of the corresponding points in a neighbouring coordinate chart. The problem is that points in a nearly degenerate, i.e., almost co-spherical, configuration make the Delaunay triangulation unstable with respect to small perturbations of the metric under consideration.

We overcome this difficulty by perturbing the sample points so that they form a configuration in which the local Delaunay triangulations are stable. The description of such point sets is based on the notion of δ -protected Delaunay simplices [BDG13c]. If $P \subset \mathbb{R}^m$, then a Delaunay m -simplex σ is δ -protected if there are no points in $P \setminus \sigma$ that are within a distance δ from the boundary of the Delaunay ball for σ . The stability of Delaunay triangulations with δ -protected m -simplices can be quantified in terms of δ .

We recently presented a simple perturbation algorithm to produce such point sets in Euclidean space [BDG13b]. It is an extension of an algorithm that had been previously proposed [ELM⁺00] for removing poorly shaped simplices from Delaunay triangulations in \mathbb{R}^3 . The algorithm proposed here extends the same idea to the context of non-Euclidean manifolds and adaptive sampling.

The essential observation that leads to the perturbation algorithm in Euclidean space is that if $P \subset \mathbb{R}^m$ is such that there exists a Delaunay m -simplex that is not δ -protected, then there is a *forbidden configuration*: a (possibly degenerate) simplex $\tau \subset P$ that has the property that every vertex is close to the circumsphere of the opposing facet. The perturbation algorithm guarantees that the Delaunay m -simplices will be δ -protected by ensuring that each point is perturbed to a position that is not too close to the circumsphere of any of the nearby simplices in the current (perturbed) point set. A volumetric argument shows that this can be achieved.

The current framework is built on the notion of local Euclidean coordinate patches and transition functions between them. We require that the transition functions have low metric distortion, i.e., they are quantifiably close to being isometries. We exploit another property of forbidden configurations: their facets have a lower bound on their *thickness*. This is a quality measure on a Euclidean simplex that indicates that it is not too close to being degenerate. We show that the change in the thickness and circumradius of a simplex under the influence of a transition function can be bounded. With these observations, together with previous results on the

stability of circumcentres with respect to perturbations [BDG13c], we show that if a point is close to the circumsphere of a thick simplex in one coordinate frame, then we can quantify how far the corresponding point will be from the circumsphere of the corresponding simplex in another coordinate frame.

This analysis permits us to use essentially the same perturbation algorithm as in the Euclidean case. We perturb each point in turn, sufficiently avoiding all relevant circumspheres. The perturbation of a point is performed in the Euclidean coordinate patch that is associated with that point. The metric on the manifold is never explicitly considered, and indeed it need not be explicitly given: it is approximately implicitly encoded by the local Euclidean metrics and the transition functions between them.

The output is a manifold abstract simplicial complex which is the Delaunay complex associated with the perturbed points on the manifold. We show that if a (smooth) Riemannian manifold, is sufficiently densely sampled, then the algorithm will produce an output mesh that is homeomorphic to the original manifold: a Delaunay triangulation.

Much recent work in computational geometry has focused on the reconstruction problem, which means inferring properties of a compact manifold based on a finite set of sample points. Giesen and Wagner [GW04] constructed a graph on the sample points in ambient space, from which they could infer the dimension and obtain a crude estimate of geodesic distances on the underlying submanifold. Cheng et al. [CDR05] presented the first algorithm to construct a simplicial complex that is homeomorphic to the submanifold from which the input point set was sampled. The running time of this algorithm depends exponentially on the ambient dimension, and Boissonnat and Ghosh [BG11] presented an algorithm that avoids this problem by working locally on estimates to the tangent spaces of the sample points.

These algorithms are all reconstruction algorithms, which means they only assume a finite point set as input. The point set is assumed to lie on or near a submanifold of Euclidean space. In contrast, we present here a triangulation algorithm which assumes complete knowledge of the manifold, but makes no reference to an ambient space. Our algorithm can be viewed as a generalisation of an earlier anisotropic meshing algorithm [BWY11] that required affine transition functions and a global reference coordinate patch.

We benefited from ideas and cautions presented in Gallier’s notes [Gal11].

We first review the algorithm in the Euclidean setting (Section 2), and then present a comprehensive overview of the algorithm extended to manifolds with curvature, its analysis, and its output guarantees (Section 3). The details of the calculations are presented in Section A.

2 Review of the algorithm for flat manifolds

We review the perturbation algorithm [BDG13b] for points in Euclidean space. The algorithm for more general manifolds is a natural extension of this.

We consider a finite set $P \subset \mathbb{R}^m$. A simplex $\sigma \subset P$ is a finite collection of points: $\sigma = \{p_0, \dots, p_k\}$. We work with abstract simplices, and in particular $x \in \sigma$ means x is a vertex of σ . The dimension of a simplex is one less than the number of vertices, and a j -dimensional simplex is a j -simplex. The dimension, if it is important, may be indicated with a superscript, as in σ^m . The notation $\sigma \leq \tau$ indicates that σ is a face of τ (i.e., a subset), and $\sigma < \tau$ means that it is a proper face. The join of two

simplices, $\tau * \sigma$, is the union of their vertices. By the standard abuse of notation, a point p may represent the 0-simplex $\{p\}$.

If $\sigma \subset \mathbb{R}^m$, then we freely talk about standard geometric properties, such as the diameter, $\Delta(\sigma)$, and the length of the shortest edge $L(\sigma)$. For $p \in \sigma$, σ_p is the facet opposite p , and $D(p, \sigma)$ is the *altitude* of p in σ , i.e., $D(p, \sigma) = d(p, \text{aff}(\sigma_p))$.

The *thickness* of σ^j is defined as

$$\Upsilon(\sigma^j) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p \in \sigma^j} \frac{D(p, \sigma^j)}{j \Delta(\sigma^j)} & \text{otherwise.} \end{cases}$$

It is a measure of the quality of the simplex. If $\Upsilon(\sigma) = 0$, then σ is *degenerate*. We say that σ is Υ_0 -thick, if $\Upsilon(\sigma) \geq \Upsilon_0$. If σ is Υ_0 -thick, then so are all of its faces.

A *circumscribing ball* for a simplex σ is any m -dimensional ball that contains the vertices of σ on its boundary. A degenerate simplex may not admit any circumscribing ball. If σ admits a circumscribing ball, then it has a *circumcentre*, $C(\sigma)$, which is the centre of the unique smallest circumscribing ball for σ . The radius of this ball is the *circumradius* of σ , denoted $R(\sigma)$. A degenerate simplex σ may or may not have a circumcentre and circumradius; we write $R(\sigma) < \infty$ to indicate that it does. In this case we can also define the *diametric sphere* as the boundary of the smallest circumscribing ball: $S^{m-1}(\sigma) = \partial B(C(\sigma), R(\sigma))$, and the *circumsphere*: $S(\sigma) = S^{m-1}(\sigma) \cap \text{aff}(\sigma)$. If $\sigma \leq \tau$, then $S(\sigma) \subseteq S(\tau)$, and if $\dim \sigma = m$, then $S(\sigma) = S^{m-1}(\sigma)$.

A ball $B(c, r)$ is open, and $\overline{B}(c, r)$ is its closure. The Delaunay complex, $\text{Del}(\mathbf{P})$ is the (abstract) simplicial complex defined by the criterion that a simplex belongs to $\text{Del}(\mathbf{P})$ if it has a circumscribing ball whose intersection with \mathbf{P} is empty. An m -simplex σ^m is δ -protected if $\overline{B}(C(\sigma^m), R(\sigma^m) + \delta) \cap \mathbf{P} = \sigma^m$.

A point set whose Delaunay m -simplices are δ -protected has quantifiable stability properties with respect to perturbations of the points, or of the metric [BDG13c]. In order for comparisons between radii, edge lengths, and δ , to be meaningful, we introduce a common “unit”. Let ϵ be an upper bound on the circumradii of the Delaunay balls centred in $\text{conv}(\mathbf{P})$, and let $\mu_0 \epsilon$ be the minimum distance between distinct points in \mathbf{P} . Since \mathbf{P} is finite, ϵ and μ_0 exist, with the only additional required assumptions on \mathbf{P} being that $\text{aff}(\mathbf{P}) = \mathbb{R}^m$, and that the points in \mathbf{P} are distinct.

If $D \subset \mathbb{R}^m$, then \mathbf{P} is ϵ -dense for D if $d(x, \mathbf{P}) < \epsilon$ for all $x \in D$. If no domain is specified, \mathbf{P} is ϵ -dense if it is ϵ -dense for

$$D_\epsilon(\mathbf{P}) = \{x \in \text{conv}(\mathbf{P}) \mid d(x, \partial \text{conv}(\mathbf{P})) \geq \epsilon\}.$$

We refer to ϵ as the *sampling radius*. The set \mathbf{P} is $\mu_0 \epsilon$ -separated if $d(p, q) \geq \mu_0 \epsilon$ for all $p, q \in \mathbf{P}$, and \mathbf{P} is a (μ_0, ϵ) -net (for D) if it is $\mu_0 \epsilon$ -separated, and ϵ -dense (for D).

The Delaunay stability results require an upper bound on the circumradii of the Delaunay balls. The *Delaunay complex of \mathbf{P} restricted to D* is the subcomplex consisting of simplices that have a Delaunay ball centred in D . We denote the Delaunay complex of \mathbf{P} restricted to $D_\epsilon(\mathbf{P})$ by $\text{Del}_\epsilon(\mathbf{P})$, and observe that the Delaunay balls for simplices in $\text{Del}_\epsilon(\mathbf{P})$ have radii smaller than ϵ , and in particular, $R(\sigma) < \epsilon$ if $\sigma \in \text{Del}_\epsilon(\mathbf{P})$.

The goal is to produce a perturbed point set \mathbf{P}' such that the Delaunay m -simplices of $\text{Del}_\epsilon(\mathbf{P}')$ are δ -protected. A ρ -perturbation of a (μ_0, ϵ) -net $\mathbf{P} \subset \mathbb{R}^m$ is a bijective application $\zeta : \mathbf{P} \rightarrow \mathbf{P}' \subset \mathbb{R}^m$ such that $d(\zeta(p), p) \leq \rho$ for all $p \in \mathbf{P}$. Unless otherwise

specified, a *perturbation* will always refer to a ρ -perturbation, with $\rho = \rho_0\epsilon$ for some

$$\rho_0 \leq \frac{\mu_0}{4}. \quad (1)$$

We also refer to P' itself as a perturbation of P . We generally use p' to denote the point $\zeta(p) \in P'$, and similarly, for any point $q' \in P'$ we understand q to be its preimage in P .

The ‘‘perturbations’’ constrained by Equation (1) are not particularly ‘‘small’’, and in particular, we do not expect a close relationship between $\text{Del}(P)$ and $\text{Del}(P')$. However, we do have the useful observation [BDG13b, Lemma 2.2] that P' is a (μ'_0, ϵ') -net:

Lemma 2.1 If $P \subset \mathbb{R}^m$ is a (μ_0, ϵ) -net, and P' is a perturbation of P , then P' is a (μ'_0, ϵ') -net, where

- $\epsilon' = (1 + \rho_0)\epsilon \leq \frac{5}{4}\epsilon$, and
- $\mu'_0 = \frac{\mu_0 - 2\rho_0}{1 + \rho_0} \geq \frac{2}{5}\mu_0$.

The construction of the algorithm is based on the concept of *forbidden configurations*, which are specific simplices that indicate the presence Delaunay m -simplices that are not sufficiently thick, or not δ -protected. Forbidden configurations possess the α_0 -hoop property. A simplex has the α_0 -hoop property if every vertex lies close to the circumsphere of the opposing facet. Specifically, for $\alpha_0 > 0$, a simplex τ has the α_0 -hoop property if, for every $p \in \tau$,

$$d(p, S(\tau_p)) \leq \alpha_0 R(\tau_p) < \infty.$$

We will first sketch the argument that shows that Delaunay m -simplices that are not δ -protected indicate the presence of forbidden configurations. Then we will describe how the hoop property is exploited in the perturbation algorithm.

A forbidden configuration is a specific kind of poorly shaped simplex that has the property that *all* its altitudes are small. A simplex that has an upper bound on its thickness will also have an altitude that is subjected to an upper bound, but in order to obtain a useful upper bound on all of the altitudes, we need a more refined measure of simplex quality. Given a positive parameter $\Gamma_0 \leq 1$, we say that σ is Γ_0 -good if for all $\sigma^j \leq \sigma$, we have $\Upsilon(\sigma^j) \geq \Gamma_0^j$, where Γ_0^j is the j^{th} power of Γ_0 . A simplex that is Γ_0 -good is necessarily Γ_0^m -thick, but the converse is not generally true. A Γ_0 -flake is a simplex that is not Γ_0 -good, but whose facets all are. The altitudes of a flake are subjected to an upper bound proportional to Γ_0 .

If a simplex is not Γ_0 -good, then it necessarily contains a face that is a flake. This follows easily from the observation that $\Upsilon(\sigma) = 1$ if σ is a 1-simplex. If $\sigma^m \in \text{Del}(P)$ is not δ -protected, then there is a $q \in P \setminus \sigma^m$ that is within a distance δ of the circumsphere of σ^m . Since $q * \sigma^m$ is $(m + 1)$ -dimensional, it is degenerate, and therefore has a face τ that is a Γ_0 -flake. Such a τ is a forbidden configuration.

The bound on the altitudes, together with the stability property of circumscribing balls of thick simplices, allows us to demonstrate that forbidden configurations have the hoop property. In addition to the two parameters that describe a (μ_0, ϵ) -net, forbidden configurations depend on the flake parameter Γ_0 , as well as the parameter δ_0 , which governs the protection via the requirement $\delta = \delta_0\mu_0\epsilon$. We do not need to explicitly define forbidden configurations here. Instead, we make reference to the following summary [BDG13b, Theorem 3.10] of properties of forbidden configurations in P' :

Lemma 2.2 (Properties of forbidden configurations) Suppose that $P \subset \mathbb{R}^m$ is a (μ_0, ϵ) -net and that P' is a perturbation of P . If

$$\delta_0 \leq \Gamma_0^{m+1} \quad \text{and} \quad \Gamma_0 \leq \frac{2\mu_0^2}{75}, \quad (2)$$

then every forbidden configuration $\tau \subset P'$ satisfies *all* of the following properties:

$\mathcal{P}1$ Simplex τ has the α_0 -hoop property, with $\alpha_0 = \frac{2^{13}\Gamma_0}{\mu_0^3}$.

$\mathcal{P}2$ For all $p \in \tau$, $R(\tau_p) < 2\epsilon$.

$\mathcal{P}3$ $\Delta(\tau) < \frac{5}{2}(1 + \frac{1}{2}\delta_0\mu_0)\epsilon$.

$\mathcal{P}4$ Every facet of τ is Γ_0 -good.

The algorithm focuses on Property $\mathcal{P}1$ of forbidden configurations. A critical aspect of this property is its symmetric nature; if we can ensure that τ has one vertex that is not too close to its opposite facet, then τ cannot be a forbidden configuration.

Using Property $\mathcal{P}3$, we can find, for each $p \in P$, a complex \mathcal{S}_p consisting all simplices $\sigma \in P$ such that after perturbations $p * \sigma$ could be a forbidden configuration.

The algorithm proceeds by perturbing each point $p \in P$ in turn, such that each point is only visited once. The perturbation $p \mapsto p'$ is found by randomly trying perturbations $p \mapsto x$ until it is found that x is a *good perturbation*. A good perturbation is one in which $d(x, S(\sigma)) > 2\alpha_0\epsilon$ for all $\sigma \in \mathcal{S}_p(P')$, where $\mathcal{S}_p(P')$ is the complex in the current perturbed point set whose simplices correspond to those in \mathcal{S}_p .

A volumetric argument based on the finite number of simplices in \mathcal{S}_p , the small size of α_0 , and the volume of the ball $B(p, \rho_0\epsilon)$ of possible perturbations of p , reveals a high probability that $p \mapsto x$ will be a good perturbation, and thus ensures that the algorithm will terminate.

Upon termination there will be no forbidden configurations in P' , because every perturbation $p \mapsto p'$ ensures that there are no forbidden configurations incident to p' in the current point set, and no new forbidden configurations are introduced.

3 Overview of the extended algorithm

The extension of the perturbation algorithm to the curved setting is accomplished by performing the perturbations, and the analysis, in local Euclidean coordinate patches.

We assume we have a finite set of points P in a compact manifold \mathcal{M} . It is convenient to employ an index set \mathcal{N} of unique (integer) labels for P , thus we employ a bijection $\iota : \mathcal{N} \rightarrow P \subset \mathcal{M}$. We assume that P is sufficiently dense that we may define an atlas $\{(W_i, \varphi_i)\}_{i \in \mathcal{N}}$ for \mathcal{M} such that the coordinate charts $\varphi_i : W_i \rightarrow U_i \subset \mathbb{R}^m$ have low metric distortion, as defined in Section 3.2. We refer to U_i as a *coordinate patch*.

We work exclusively in the Euclidean coordinate patches U_i , exploiting the transition functions $\varphi_{ji} = \varphi_j \circ \varphi_i^{-1}$ to translate between them. We define $P_i = \varphi_i(W_i \cap P)$, but given these sets, the algorithm itself makes no explicit reference to either P or the coordinate charts φ_i , except to keep track of the labels of the points. We employ the discrete map $\phi_i = \varphi_i \circ \iota$ to index the elements of the set P_i .

The idea is to perturb $p_i = \phi_i(i) \in U_i$ in such a way, $p_i \mapsto p'_i$, that not only are there no forbidden configurations incident to p'_i in $P'_i = \phi_i(W_i \cap P')$, but there are no forbidden configurations incident to $\varphi_{ji}(p'_i) \in P'_j \subset U_j$ either, where j is the index of any sample point near p_i .

Before detailing the requirements of the input data in Section 3.2, we briefly discuss the implicit and explicit properties of the underlying manifold \mathcal{M} in Section 3.1. The analysis of the algorithm is summarised in Section 3.3, and the quality of the output complex is discussed in Section 4.

3.1 Manifolds represented by transition functions

The essential input data for the algorithm are the transition functions, and the sample points in the coordinate patches; we do not explicitly use the coordinate charts or the metric on the manifold. However, given that the transition functions can be defined by an atlas on a manifold, this manifold is essentially unique.

If $\tilde{\mathcal{M}}$ has an atlas $\{(\tilde{W}_i, \tilde{\varphi}_i)\}_{i \in \mathcal{N}}$ such that $\tilde{\varphi}_i(\tilde{W}_i) = U_i$ and $\varphi_{ji} = \tilde{\varphi}_j \circ \tilde{\varphi}_i^{-1}$ for all $i, j \in \mathcal{N}$, then $\tilde{\mathcal{M}}$ and \mathcal{M} are homeomorphic. Indeed, we define the map $f : \mathcal{M} \rightarrow \tilde{\mathcal{M}}$ by $f(x) = \tilde{\varphi}_i^{-1} \circ \varphi_i(x)$ if $x \in W_i$. The map is well defined, because $\varphi_j = \varphi_{ji} \circ \varphi_i$ on $W_i \cap W_j$, and $\tilde{\varphi}_j^{-1} = \tilde{\varphi}_i^{-1} \circ \varphi_{ji}^{-1}$ on $U_{ji} = \varphi_j(W_i \cap W_j)$. It can be verified directly from the definition that f is a homeomorphism, since it is bijective and locally a homeomorphism.

Although the algorithm does not make explicit reference to a metric on the manifold \mathcal{M} , the metric distortion bounds required on the transition functions imply a metric constraint. Implicitly we are using a metric on the manifold for which the coordinate charts have low metric distortion. If a metric on the manifold is not explicitly given then, at least in the case where the transition functions are smooth, we can be sure that such a metric exists: Given the coordinate charts an appropriate Riemmanian metric on the manifold can be obtained from the coordinate patches by the standard construction employing a partition of unity subordinate to the atlas (e.g., [Boo86, Thm. V.4.5]).

Thus, although the manifold may be presented abstractly in terms of coordinate patches and transition functions between them, this information essentially characterises the manifold. The algorithm we present is not a reconstruction algorithm, it is an algorithm to triangulate a known manifold.

3.2 The setting and input data

We take as input a finite index set $\mathcal{N} = \{1, \dots, n\}$, which we might think of as an abstract set of points (without geometry), together with the geometric data we will now introduce. The details of the arguments that lead to our choices in the size of the domains are given in Section A.3.

Local charts. For each $i \in \mathcal{N}$ we have a neighbourhood set $\mathcal{N}_i \subset \mathcal{N}$, a sampling radius $\epsilon_i > 0$, and an injective application

$$\phi_i : \mathcal{N}_i \rightarrow U_i \subset \mathbb{R}^m,$$

such that $P_i = \phi_i(\mathcal{N}_i)$ is an ϵ_i -dense sample set for $B(p_i, 8\epsilon_i) \subset U_i$, where we adopt the notation $p_j = \phi_i(j)$ for any $j \in \mathcal{N}_i$. We call the standard metric on $U_i \subset \mathbb{R}^m$ the *local Euclidean metric* for i , and we will denote it by d_i to distinguish between

the different local Euclidean metrics. Similarly, $B_i(c, r)$ denotes a ball with respect to the metric d_i .

Transition functions. For each $p_j \in B_i(p_i, 6\epsilon_i) \subset U_i$ we require a neighbourhood $U_{ij} \subset U_i$ such that

$$B_i(p_i, 6\epsilon_i) \cap B_i(p_j, 9\epsilon_i) \subset U_{ij}$$

and

$$U_{ij} \cap \phi_i(\mathcal{N}_i) = \phi_i(\mathcal{N}_i \cap \mathcal{N}_j).$$

The set U_{ij} is the domain of the *transition function* φ_{ji} , which is a homeomorphism

$$\varphi_{ji} : U_{ij} \rightarrow U_{ji} \subset U_j,$$

such that $\varphi_{ji} = \varphi_{ij}^{-1}$ and

$$\varphi_{ji} \circ \phi_i = \phi_j \quad \text{on} \quad \mathcal{N}_i \cap \mathcal{N}_j.$$

These transition functions are required to have *low metric distortion*:

$$|d_i(x, y) - d_j(\varphi_{ji}(x), \varphi_{ji}(y))| \leq \xi_0 d_i(x, y) \quad \text{for all } x, y \in U_{ij}, \quad (3)$$

where $\xi_0 > 0$ is a small positive parameter that quantifies the metric distortion. We say that φ_{ji} is a ξ_0 -*distortion map*.

In order to ease the notational burden, $\phi_i(j) \in U_{ik}$, and $\phi_k(j) \in U_{ki}$, are denoted by the same symbol, p_j . Ambiguities are avoided by distinguishing between the Euclidean metrics d_i and d_k . Although d_i is the canonical metric on U_{ik} , we may consider the pullback of d_k from the homeomorphic domain U_{ki} . Thus for $x, y \in U_{ik}$ the expression $d_k(x, y)$ is understood to mean $d_k(\varphi_{ki}(x), \varphi_{ki}(y))$, but we also occasionally employ the latter, redundant, notation.

Using symmetry, we observe that Equation (3) implies that

$$|d_i(x, y) - d_j(x, y)| \leq \xi_0 \min\{d_i(x, y), d_j(x, y)\}.$$

Our analysis will require that ξ_0 be very small. For standard coordinate charts, ξ_0 can be shown to be $O(\epsilon)$, where ϵ is a sampling radius on the manifold. For example, this is the case when considering a smooth submanifold of \mathbb{R}^N , and using the orthogonal projection onto the tangent space as a coordinate chart [BDG13a, Lemma 3.7]. Thus ξ_0 may be made as small as desired by increasing the sampling density.

Adaptive sampling We will further require a constraint on the difference between neighbouring sampling radii:

$$|\epsilon_i - \epsilon_j| \leq \epsilon_0 \min\{\epsilon_i, \epsilon_j\},$$

whenever $d_i(p_i, p_j) \leq 6\epsilon_i$. This allows us to work with a constant sampling radius in each coordinate frame, while accommodating a globally adaptive sampling radius.

For example, suppose $\epsilon : \mathcal{M} \rightarrow \mathbb{R}$ is a positive, ν -Lipschitz function, with respect to the metric $d_{\mathcal{M}}$ on the manifold. Then ϵ may be used as an adaptive sampling radius on \mathcal{M} , i.e., $\mathbf{P} \subset \mathcal{M}$ is ϵ -dense if $d_{\mathcal{M}}(x, \mathbf{P}) < \epsilon(x)$ for all $x \in \mathcal{M}$. A popular example of such a function is $\epsilon(x) = \nu f(x)$, where f is the (1-Lipschitz) *local feature size* [AB99].

Using the ν -Lipschitz continuity of ϵ , we can define, for any $p_i \in \mathcal{M}$, a constant ϵ_i , such that \mathbf{P} is ϵ_i dense in some neighbourhood of p_i . In fact, given $c > 0$, with $c < \nu^{-1}$, we find that \mathbf{P} is ϵ_i dense within the ball $B_{\mathcal{M}}(p_i, c\epsilon_i)$, where

$$\epsilon_i = \frac{\epsilon(p_i)}{1 - c\nu}.$$

For any $p_j \in B_{\mathcal{M}}(p_i, \epsilon_i)$, we obtain $|\epsilon_i - \epsilon_j| \leq \epsilon_0 \epsilon_i$, where

$$\epsilon_0 = \frac{c\nu}{1 - c\nu}, \quad (4)$$

and if $\nu \leq \frac{1}{2c}$, then $\epsilon_0 \leq 1$.

Similarly, if \mathbf{P} is $\hat{\mu}_0 \epsilon$ -sparse, then it will be $\mu_0 \epsilon_i$ -sparse on $B_{\mathcal{M}}(p_i, c\epsilon_i)$, provided $\mu_0 \leq (1 - 2c\nu)\hat{\mu}_0$. The constant $\hat{\mu}_0$ itself may be constrained to satisfy $\hat{\mu}_0 \leq (1 + \nu)^{-1} \leq \frac{1}{2}$.

In our framework here, the local constant sampling radii are applied to the local Euclidean metric, rather than the metric on the manifold, but the same idea applies. Although Equation (4) indicates that ϵ_0 is expected to become small as the sampling radius decreases, our analysis does not demand this. As explained in Section 3.3, we only require that ϵ_0 be mildly bounded.

3.3 Outline of the analysis

We have defined the point sets $\mathbf{P}_i = \phi_i(\mathcal{N}_i)$ in the local frame for p_i . We will let \mathbf{P}'_i denote the corresponding perturbed point set at any stage in the algorithm: \mathbf{P}'_i changes during the course of the algorithm, and we do not rename it according to the iteration as was done in the original description of the algorithm for flat manifolds [BDG13b]. The perturbation of a point $p_i \mapsto p'_i$ is performed in the coordinate patch U_i , and then all the coordinate charts must be updated so that if $i \in \mathcal{N}_j$, then $\phi'_j(i) = \varphi_{ji}(p'_i)$. However, we will refer to the point as p'_i regardless of which coordinate frame we are considering. The discrete maps ϕ'_i will change as the algorithm progresses, but ϕ_i will always refer to the initial map.

In order to exploit Lemma 2.1, we need to constrain the perturbation so that Equation (1) effectively applies in all local Euclidean metrics. If $p_i \mapsto \zeta(p_i)$ such that $d_i(\zeta(p_i), p_i) \leq \rho = \rho_0 \epsilon_i$, then $d_j(\zeta(p_i), p_i) \leq (1 + \xi_0)\rho \leq (1 + \xi_0)(1 + \epsilon_0)\rho_0 \epsilon_j$, and we effectively have

$$\tilde{\rho}_0 = (1 + \xi_0)(1 + \epsilon_0)\rho_0.$$

Thus we demand that

$$\tilde{\rho}_0 \leq \frac{\mu_0}{4}.$$

We will assume that

$$\epsilon_0 \leq \frac{1 - \xi_0}{1 + \xi_0},$$

and observe that this implies that

$$(1 + \epsilon_0)(1 + \xi_0) \leq 2. \quad (5)$$

We will keep the definition of forbidden configuration as in the flat case. In other words a forbidden configuration is that which satisfies the four properties described in Lemma 2.2, where ϵ refers to the local sampling radius ϵ_i .

We do not attempt to remove the forbidden configurations from all of \mathbf{P}'_i . Rather, we define $\mathbf{Q}'_i = \mathbf{P}'_i \cap B_i(p_i, 6\epsilon_i)$ as our region of interest. The reasoning behind this choice appears in Section A.3, where we also show (Lemma A.12) that [BDG13b, Lemma 3.5] implies:

Lemma 3.1 (Protected stars) If there are no forbidden configurations in \mathbf{Q}'_i , then all the m -simplices in $\text{star}(p'_j; \text{Del}(\mathbf{Q}'_i))$ are Γ_0 -good and δ -protected, with $\delta = \delta_0 \mu'_0 \epsilon'_i$.

This allows us to exploit the Delaunay metric stability result [BDG13c, Theorem 4.17], which we show (Lemma A.13) may be stated in our current context as:

Lemma 3.2 (Stable stars) If

$$\xi_0 \leq \frac{\Gamma_0^{2m+1} \mu_0^2}{2^{12}},$$

and there are no forbidden configurations in \mathbf{Q}'_i , then for all $p'_j \in \text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$, we have

$$\text{star}(p'_i; \text{Del}(\mathbf{P}'_i)) \cong \text{star}(p'_i; \text{Del}(\mathbf{P}'_j)).$$

The main technical result we develop in the current analysis is the bound on the distortion of the hoop property (Lemma A.15) due to the transition functions:

Lemma 3.3 If

$$\xi_0 \leq \left(\frac{\Gamma_0^{2m+1}}{4} \right)^2,$$

then for any forbidden configuration $\tau = p'_j * \sigma \subset \mathbf{Q}'_i$, there is a simplex $\tilde{\sigma} = \varphi_{ji}(\sigma) \subset \mathbf{P}'_j$ such that $d_j(p'_j, S^{m-1}(\tilde{\sigma})) \leq 2\tilde{\alpha}_0 \epsilon_j$, where

$$\tilde{\alpha}_0 = \frac{2^{16} m^{\frac{3}{2}} \Gamma_0}{\mu_0^3}.$$

The proof of Lemma 3.3 relies heavily on the thickness bound (Property $\mathcal{P}4$) for the facets of a forbidden configuration. In Section A.1 we show bounds on the changes of the intrinsic properties, such as thickness and circumradius, of a Euclidean simplex subjected to the influence of a transition function. This leads, as shown in Section A.2, to bounds on circumcentre displacement under small changes of a Euclidean metric. These bounds could not be recovered directly from earlier work [BDG13c], because they involve simplices that are not full dimensional. With these results in place, the proof of Lemma 3.3 is assembled in Section A.4.

By considering the diameter of a forbidden configuration subjected to metric distortion, we can determine the size of the neighbourhood of p_i that must be considered when checking whether a perturbation $p_i \mapsto p'_i$ creates conflicts.

Suppose $\tau \subset \mathbf{Q}'_j \subset U_j$ is a forbidden configuration with $p'_i \in \tau$. By Lemma 2.2, Property $\mathcal{P}3$, we have $\Delta(\tau) < \frac{5}{2}(1 + \frac{1}{2}\delta_0 \mu_0) \epsilon_j$. It follows then that if $\tilde{\tau} = \varphi_{ij}(\tau) \subset \mathbf{P}'_i$, then

$$\begin{aligned} \Delta(\tilde{\tau}) &< (1 + \xi_0)(1 + \epsilon_0) \frac{5}{2} \left(1 + \frac{1}{2} \delta_0 \mu_0 \right) \epsilon_i \\ &\leq 5 \left(1 + \frac{1}{2} \delta_0 \mu_0 \right) \epsilon_i, \end{aligned}$$

and we find, as in [BDG13b, Lemma 4.4], that if $\delta_0 \leq \frac{2}{5}$, then

$$(\phi'_j)^{-1}(\tau) \subset \phi_i^{-1}(B_i(p_i, r) \cap \mathbf{P}_i), \quad \text{where } r = \left(5 + \frac{3\mu_0}{2} \right) \epsilon_i.$$

Indeed, this is ensured by the fact that \mathbf{P}_i is a (μ_0, ϵ_i) -net for $B_i(p_i, 8\epsilon_i)$, and $8\epsilon_i - r > \epsilon_i$.

Let \mathcal{S}_i denote all the m -simplices in \mathcal{N}_i whose vertices are contained in

$$\phi_i^{-1}(B_i(p_i, r) \cap \mathbf{P}_i) \setminus \{i\} \quad \text{where} \quad r = \left(5 + \frac{3\mu_0}{2}\right) \epsilon_i.$$

Then the simple packing argument demonstrated in [BDG13b, Lemma 4.4] yields

$$\#\mathcal{S}_i < \left(\frac{14}{\mu_0}\right)^{m^2+m}. \quad (6)$$

We strengthen the definition of a good perturbation:

Definition 3.4 For the extended algorithm, we say that $p_i \mapsto x$ is a *good perturbation* of $p_i \in U_i$ if there are no m -simplices $\sigma \in \phi'_i(\mathcal{S}_i)$ such that $d_i(x, S^{m-1}(\sigma)) \leq 2\tilde{\alpha}_0\epsilon_i$, where $\tilde{\alpha}_0$ is defined in Lemma 3.3.

It is sufficient to only consider the m -simplices, because if σ is a non-degenerate j -simplex, with $j < m$, then it is the face of some non-degenerate m -simplex τ , and $S(\sigma) \subset S^{m-1}(\tau)$. With this definition of a good perturbation, the extended algorithm yields the analogue of [BDG13b, Lemma 4.3]:

Lemma 3.5 After the extended algorithm terminates, for every $i \in \mathcal{N}$ there will be no forbidden configurations in \mathbf{Q}'_i .

Proof We argue by induction that after the i^{th} iteration, for any $j \leq i$, and any $k \in \mathcal{N}$, there are no forbidden configurations in \mathbf{Q}'_k that have p'_j as a vertex. For $i = 1$, the assertion follows directly from Definition 3.4, and Lemma 3.3. Assume the assertion is true for $i - 1$. Suppose τ is a forbidden configuration in \mathbf{Q}'_k , after the i^{th} iteration. Then since p'_i is a good perturbation, according to Definition 3.4, τ cannot contain p'_i . Also, τ cannot contain any p'_j with $j < i$, for that would contradict the induction hypothesis. Thus the hypothesis holds for all $i \geq 1$. \square

We need to quantify the conditions under which the algorithm is guaranteed to terminate. We use the same volumetric analysis that is demonstrated in the proof of [BDG13b, Lemma 5.4], with the only modifications being a change in two of the constants involved in the calculation. In particular, the number of simplices involved is now given by Equation (6), and we use the bound on $\tilde{\alpha}_0$ given by Lemma 3.3, which is $2^3 m^{\frac{3}{2}}$ times the bound on α_0 used in the original calculation. Thus, using Lemma 3.1, the main result [BDG13b, Theorem 4.1] of the original perturbation algorithm can be adapted to the context of the extended algorithm as:

Lemma 3.6 (Algorithm guarantee) Let

$$\tilde{\rho}_0 = (1 + \epsilon_0)(1 + \xi_0)\rho_0.$$

If

$$\epsilon_0 \leq \frac{1 - \xi_0}{1 + \xi_0}, \quad \text{and} \quad \tilde{\rho}_0 \leq \frac{\mu_0}{4}, \quad \text{and} \quad \xi_0 \leq \frac{1}{2^4} \left(\frac{\rho_0}{C}\right)^{4m+2},$$

where $C = m^{\frac{3}{2}} \left(\frac{2}{\mu_0}\right)^{4m^2+5m+21}$, then the extended algorithm terminates, and for every $i \in \mathcal{N}$, the set \mathbf{Q}'_i is a (μ'_0, ϵ'_i) -net such that there are no forbidden configurations with

$$\Gamma_0 = \frac{\rho_0}{C}, \quad \text{and} \quad \delta = \Gamma_0^{m+1} \mu'_0 \epsilon'_i,$$

where $\mu'_0 = \frac{\mu_0 - 2\tilde{\rho}_0}{1 + \tilde{\rho}_0}$, and $\epsilon'_i = (1 + \tilde{\rho}_0)\epsilon_i$.

This allows us to apply Lemma 3.2, and we can define the abstract complex $\text{Del}(\mathbf{P}')$ by the criterion that $\phi'_i(\text{star}(i; \text{Del}(\mathbf{P}'))) = \text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$ for all $i \in \mathcal{N}$. This is a manifold piecewise linear simplicial complex. The bound on ξ_0 imposed by Lemma 3.2 is met by the one imposed by Lemma 3.6, and we arrive at our main result:

Theorem 3.7 (Manifold mesh) Let

$$\tilde{\rho}_0 = (1 + \epsilon_0)(1 + \xi_0)\rho_0.$$

If

$$\epsilon_0 \leq \frac{1 - \xi_0}{1 + \xi_0}, \quad \text{and} \quad \tilde{\rho}_0 \leq \frac{\mu_0}{4}, \quad \text{and} \quad \xi_0 \leq \frac{1}{2^4} \left(\frac{\rho_0}{C} \right)^{4m+2},$$

where $C = m^{\frac{3}{2}} \left(\frac{2}{\mu_0} \right)^{4m^2+5m+21}$, then the extended algorithm produces a manifold abstract simplicial complex $\text{Del}(\mathbf{P}')$ defined by

$$\text{star}(i; \text{Del}(\mathbf{P}')) \cong \text{star}(p'_i; \text{Del}(\mathbf{P}'_i)).$$

4 Output quality guarantees

Given appropriate constraints on the input, Theorem 3.7 guarantees that the output mesh, $\text{Del}(\mathbf{P}')$, will be a manifold simplicial complex. In this section, we justify the name $\text{Del}(\mathbf{P}')$ and examine the fitness of $\text{Del}(\mathbf{P}')$ as a representative of the input manifold \mathcal{M} .

4.1 The output complex is a Delaunay complex

Given $\mathbf{P} \subset \mathcal{M}$, we define the set $\mathbf{P}' \subset \mathcal{M}$ to be the perturbed point set produced by the algorithm, i.e., $\mathbf{P} \rightarrow \mathbf{P}'$ is given by $p \mapsto p' = \varphi_i^{-1}(p'_i)$, where $i \in \mathcal{N}$ is the label associated with $p \in \mathbf{P}$. If the metric on \mathcal{M} is such that the coordinate maps φ_i themselves have low metric distortion, then the constructed complex $\text{Del}(\mathbf{P}')$ is in fact the Delaunay complex of $\mathbf{P}' \subset \mathcal{M}$. This follows from the fact that in the local Euclidean coordinate frames we have ensured that the points have stable Delaunay triangulations. Thus, using $\Gamma_0 = \frac{\rho_0}{C}$ given by Lemma 3.6, the stability result Lemma A.10 leads, by the same reasoning that yields Lemma 3.2, to the following:

Theorem 4.1 (Delaunay complex) Suppose that $\{(W_i, \varphi_i)\}_{i \in \mathcal{N}}$ is an atlas for the compact m -manifold \mathcal{M} , and the finite set $\mathbf{P} \subset \mathcal{M}$ is such that the conditions of Section 3.2 are satisfied. Suppose also that \mathcal{M} is equipped with a metric $d_{\mathcal{M}}$, such that

$$|d_i(\varphi_i(x), \varphi_i(y)) - d_{\mathcal{M}}(x, y)| \leq \eta d_i(\varphi_i(x), \varphi_i(y)),$$

whenever x and y belong to $\varphi_i^{-1}(B(p_i, 6\epsilon_i))$. If

$$\eta \leq \frac{\mu_0^2}{2^{12}} \left(\frac{\rho_0}{C} \right)^{2m+1},$$

and the conditions of Theorem 3.7 are met, then $\text{Del}(\mathbf{P}')$ is the Delaunay complex of $\mathbf{P}' \subset \mathcal{M}$ with respect to $d_{\mathcal{M}}$.

Thus we construct a Delaunay complex for a metric that is locally well approximated by the local Euclidean metrics. This is the standard scenario: the manifold is sampled densely enough that natural local coordinate charts have very low distortion and thus yield transition functions which satisfy the requirements of Theorem 3.7.

4.2 A metric on the output complex

The output mesh $\text{Del}(P')$ is an abstract simplicial complex, but it admits a natural piecewise flat metric. This means that each simplex in $\text{Del}(P')$ is identified with a geometric Euclidean simplex. This is achieved by assigning lengths to the edges in $\text{Del}(P')$. If $\{i, j\}$ is an edge in $\text{Del}(P')$, then we give it the length ℓ_{ij} , where $\ell_{ij} = \frac{1}{2}(d_i(p'_i, p'_j) + d_j(p'_i, p'_j))$.

We observe that with this assignment of edge lengths every simplex in $\text{Del}(P')$ can be isometrically identified with the vertices of a geometric Euclidean simplex. This is a consequence of our observations in Section A.1. Specifically, suppose $\sigma = \{p_0, \dots, p_k\}$ is a simplex of points in \mathbb{R}^m and let $\{\ell_{ij}\}_{0 \leq i < j \leq k}$ be a set of positive numbers such that $\ell_{ij} = \ell_{ji}$ and $|\ell_{ij} - \|p_i - p_j\|| \leq \xi_0 \Delta(\sigma)$. Consider the $k \times k$ symmetric matrix G defined by $[G]_{ij} = \frac{1}{2}(\ell_{0i}^2 + \ell_{0j}^2 - \ell_{ij}^2)$. Then using the same notation and arguments as in Lemma A.2, we have that $G = P^\top P + E$, and that $s_1(E) \leq 4k\xi_0\Delta(\sigma)^2$ if $\xi_0 \leq \frac{2}{3}$. It follows then, using Lemma A.1, that G is positive definite if $\xi_0 < \left(\frac{\Upsilon(\sigma)}{2}\right)^2$, just as was argued in Lemma A.3. This means that we may write $G = \tilde{P}^\top \tilde{P}$, where \tilde{P} defines a non-degenerate Euclidean simplex whose edge lengths are given by the numbers ℓ_{ij} . The choice of \tilde{P} is not important: by considering the singular value decompositions it is evident that if $Q^\top Q = \tilde{P}^\top \tilde{P}$, then there is an orthogonal matrix O such that $Q = O\tilde{P}$.

For any $\tilde{\sigma} \in \text{Del}(P')$ we may compare its edge lengths to those of the corresponding simplex $\sigma \in \text{star}(p'_i; \text{Del}(P'_i))$, where i is the label of one of the vertices of $\tilde{\sigma}$. Under the assumptions of Lemma 3.6, the condition $\xi_0 < \left(\frac{\Upsilon(\sigma)}{2}\right)^2$ is satisfied, and the above argument shows that $\tilde{\sigma}$ may be identified with a Euclidean simplex.

4.3 Homeomorphism guarantee in the smooth case

If the coordinate charts have small metric distortion, then the sampling radii in the coordinate patches imply that P' satisfies some constant sampling radius, ϵ , with respect to the metric on the compact manifold. For a Riemannian manifold, a natural choice for a coordinate chart is the inverse of the exponential map, e.g., $\varphi_i = \exp_{p_i}^{-1}|_{W_i}$. Then the metric distortion of the coordinate chart φ_i decreases as the diameter of the domain W_i decreases; this is a consequence of a theorem due to Rauch [Cha06, Thm IX.2.3, p. 390]. This means that for ϵ sufficiently small, $\text{Del}(P')$ will be homeomorphic to \mathcal{M} . This can be argued by appealing to the homeomorphism demonstration detailed by Boissonnat and Ghosh [BG11], for example.

We employ the Nash embedding theorem, which ensures that \mathcal{M} may be isometrically embedded in some ambient Euclidean space \mathbb{R}^N . When the neighbourhoods W_i are small, the restricted ambient metric $d_{\mathbb{R}^N|_{\mathcal{M}}}$ is a close approximation to the intrinsic metric $d_{\mathcal{M}}$ on the manifold. (See [BDG13a, Lemma 3.7], for example.) In other words the identity map $(W_i, d_{\mathbb{R}^N|_{\mathcal{M}}}) \rightarrow (W_i, d_{\mathcal{M}})$ is an η -distortion map, where the magnitude of η decreases with the diameter of W_i .

This means that when ϵ is small enough, the stability theorem [BDG13c, Theorem 4.17] ensures that not only is $\text{Del}(P')$ the Delaunay complex of $P' \subset \mathcal{M}$ with respect to the intrinsic metric $d_{\mathcal{M}}$, but also with respect to the restricted ambient metric $d_{\mathbb{R}^N|_{\mathcal{M}}}$. We are also assured, by Lemma A.3, that the simplices of $\text{Del}(P')$ satisfy some lower bound $\tilde{\Upsilon}_0$ on their thickness when they are realised as Euclidean simplices in the ambient space, i.e., when the edge lengths between vertices are given

by the Euclidean distance in \mathbb{R}^N . These observations place $\text{Del}(P')$ in the setting considered by Boissonnat and Ghosh [BG11] in the context of the tangential Delaunay complex, and the same argument demonstrates that $\text{Del}(P')$ is homeomorphic to \mathcal{M} . The homeomorphism construction of Boissonnat and Ghosh requires the submanifold \mathcal{M} to be smooth.

4.4 Discussion

We are able to quantify the conditions under which $\text{Del}(P')$ will be a manifold Delaunay complex. These are already stringent conditions on the metric, although smoothness is not required. However, we make no claim that $\text{Del}(P')$ is homeomorphic to \mathcal{M} , except for certain smooth transition functions, and for “sufficiently dense” sampling. The specification of concrete sampling criteria that are sufficient to guarantee a homeomorphic complex is still an open problem. A careful construction of a homeomorphism will also enable us to quantify its metric distortion, i.e., to quantify a bound on the Gromov-Hausdorff distance between \mathcal{M} and $\text{Del}(P')$, when the latter is equipped with the metric described in Section 4.2, for example.

When the differences in the metrics are sufficiently small, the Delaunay stability results imply that not only is $\text{Del}(P')$ a Delaunay complex for $P' \subset \mathcal{M}$, but it is also the Delaunay complex of its vertices associated with its own intrinsic metric. Such meshes are of interest in discrete differential geometry [Gli05, Dye10, HKV12].

A Details of the analysis

In this section we provide details to support the argument made in Section 3.3.

A.1 Simplex distortion

Our transition functions introduce a metric distortion when we move from one coordinate chart to another. The geometric properties of a simplex will be slightly different if we consider it with respect to the Euclidean metric d_i than they would be if we are using a different Euclidean metric d_j . We wish to bound the magnitude of the change of such properties as the thickness and the circumradius of a simplex that is subjected to such a distortion. This is an exercise in linear algebra.

We wish to compare two Euclidean simplices with corresponding vertices, but whose corresponding edge lengths differ by a relatively small amount. The embedding of the simplex in Euclidean space (i.e., the coordinates of the vertices) is not relevant to us. Previous results often only consider the case where the vertices of a given simplex are perturbed a small amount to obtain a new simplex. Lemma A.5 demonstrates the existence of an isometry that allows us to also consider the general situation in terms of vertex displacements.

We will exploit observations on the linear algebra of simplices developed in previous work [BDG13c]. A k -simplex $\sigma = \{p_0, \dots, p_k\}$ in \mathbb{R}^m can be represented by an $m \times k$ matrix P , whose i^{th} column is $p_i - p_0$. We let $s_i(A)$ denote the i^{th} singular value of a matrix A , and observe that $\|P\| = s_1(P) \leq \sqrt{k}\Delta(\sigma)$.

We are particularly interested in bounds on the smallest singular value of P , which is the inverse of the largest singular value of the pseudo-inverse $P^\dagger = (P^\top P)^{-1} P^\top$. If the columns of P are viewed as a basis for $\text{aff}(\sigma)$, then the rows of P^\dagger may be viewed as the dual basis. The magnitude of a dual vector is equal to the inverse of

the corresponding altitudes in σ , and this leads directly to the desired bound on the smallest singular value of P , which is expressed in the following Lemma [BDG13c, Lemma 2.4]:

Lemma A.1 (Thickness and singular value) Let $\sigma = [p_0, \dots, p_k]$ be a non-degenerate k -simplex in \mathbb{R}^m , with $k > 0$, and let P be the $m \times k$ matrix whose i^{th} column is $p_i - p_0$. Then the i^{th} row of P^\dagger is given by w_i^\top , where w_i is orthogonal to $\text{aff}(\sigma_{p_i})$, and

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

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

$$s_k(P) \geq \sqrt{k} \Upsilon(\sigma) \Delta(\sigma).$$

We will also have use for a lower bound on the thickness of σ , given the smallest singular value for the representative matrix P . We observe that P was constructed by arbitrarily choosing one vertex, p_0 , to serve as the origin. If there is a vertex p_i , different from p_0 , such that $D(p_i, \sigma)$ is minimal amongst all the altitudes of σ , then according to Lemma A.1, $\|w_i\| = (k \Upsilon(\sigma) \Delta(\sigma))^{-1}$, and it follows then that $s_1(P^\dagger) \geq (k \Upsilon(\sigma) \Delta(\sigma))^{-1}$, and therefore

$$s_k(P) \leq k \Upsilon(\sigma) \Delta(\sigma), \quad (7)$$

in this case.

We are going to be interested here in purely intrinsic properties of simplices in \mathbb{R}^m ; properties that are not dependent on the choice of embedding in \mathbb{R}^m . In this context it is convenient to make use of the *Gram matrix* $P^\top P$, because if $Q^\top Q = P^\top P$, then there is an orthogonal transformation O such that $P = OQ$. This assertion becomes evident when considering the singular value decompositions of P and Q . Indeed, the entries of the Gram matrix can be expressed in terms of squared edge lengths, as observed in the proof of the following:

Lemma A.2 Suppose that $\sigma = \{p_0, \dots, p_k\}$ and $\tilde{\sigma} = \{\tilde{p}_0, \dots, \tilde{p}_k\}$ are two k -simplices in \mathbb{R}^m such that

$$\| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \| \leq \xi_0 \Delta(\sigma),$$

for all $0 \leq i < j \leq k$. Let P be the matrix whose i^{th} column is $p_i - p_0$, and define \tilde{P} similarly. Consider the Gram matrices, and let E be the matrix that records their difference:

$$\tilde{P}^\top \tilde{P} = P^\top P + E.$$

If $\xi_0 \leq \frac{2}{3}$, then the entries of E are bounded by $|E_{ij}| \leq 4\xi_0 \Delta(\sigma)^2$, and in particular

$$\|E\| \leq 4k\xi_0 \Delta(\sigma)^2. \quad (8)$$

Proof Let $v_i = p_i - p_0$, and $\tilde{v}_i = \tilde{p}_i - \tilde{p}_0$. Expanding scalar products of the form $(v_j - v_i)^\top (v_j - v_i)$, we obtain a bound on the magnitude of the coefficients of E :

$$\begin{aligned} |\tilde{v}_i^\top \tilde{v}_j - v_i^\top v_j| &\leq \frac{1}{2} (|\|\tilde{v}_i\|^2 - \|v_i\|^2| + |\|\tilde{v}_j\|^2 - \|v_j\|^2| + |\|\tilde{v}_j - \tilde{v}_i\|^2 - \|v_j - v_i\|^2|) \\ &\leq \frac{3}{2} (2 + \xi_0) \xi_0 \Delta(\sigma)^2 \\ &\leq 4\xi_0 \Delta(\sigma)^2. \end{aligned}$$

This leads us to a bound on $s_1(E) = \|E\|$. Indeed, the magnitude of the column vectors of E is bounded by \sqrt{k} times a bound on the magnitude of their coefficients, and the magnitude of $s_1(E)$ is bounded by \sqrt{k} times a bound on the magnitude of the column vectors. We obtain Equation (8). \square

Lemma A.2 enables us to bound the thickness of a distorted simplex:

Lemma A.3 (Thickness under distortion) Suppose that $\sigma = \{p_0, \dots, p_k\}$ and $\tilde{\sigma} = \{\tilde{p}_0, \dots, \tilde{p}_k\}$ are two k -simplices in \mathbb{R}^m such that

$$\| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \| \leq \xi_0 \Delta(\sigma)$$

for all $0 \leq i < j \leq k$. Let P be the matrix whose i^{th} column is $p_i - p_0$, and define \tilde{P} similarly.

If

$$\xi_0 \leq \left(\frac{\eta \Upsilon(\sigma)}{2} \right)^2 \quad \text{with } \eta^2 \leq 1,$$

then

$$s_k(\tilde{P}) \geq (1 - \eta^2) s_k(P),$$

and

$$\Upsilon(\tilde{\sigma}) \Delta(\tilde{\sigma}) \geq \frac{1}{\sqrt{k}} (1 - \eta^2) \Upsilon(\sigma) \Delta(\sigma),$$

and

$$\Upsilon(\tilde{\sigma}) \geq \frac{4}{5\sqrt{k}} (1 - \eta^2) \Upsilon(\sigma).$$

Proof The equation $\tilde{P}^\top \tilde{P} = P^\top P + E$ implies that

$$|s_k(\tilde{P})^2 - s_k(P)^2| \leq s_1(E),$$

and so

$$|s_k(\tilde{P}) - s_k(P)| \leq \frac{s_1(E)}{s_k(\tilde{P}) + s_k(P)} \leq \frac{s_1(E)}{s_k(P)}.$$

Thus

$$s_k(\tilde{P}) \geq s_k(P) - \frac{s_1(E)}{s_k(P)} = s_k(P) \left(1 - \frac{s_1(E)}{s_k(P)^2} \right).$$

From Lemma A.2 and the bound on ξ_0 we have

$$s_1(E) \leq \eta^2 k \Upsilon(\sigma)^2 \Delta(\sigma)^2,$$

and so $\frac{s_1(E)}{s_k(P)^2} \leq \eta^2$ by Lemma A.1, and we obtain $s_k(\tilde{P}) \geq (1 - \eta^2) s_k(P)$.

For the thickness bound we assume, without loss of generality, that there is some vertex different from \tilde{p}_0 that realises the minimal altitude in $\tilde{\sigma}$ (our choice of ordering of the vertices is unimportant, other than to establish the correspondence between σ and $\tilde{\sigma}$). Thus Equation (7) and Lemma A.1, give the inequalities

$$k \Upsilon(\tilde{\sigma}) \Delta(\tilde{\sigma}) \geq s_k(\tilde{P}), \quad \text{and} \quad s_k(P) \geq \sqrt{k} \Upsilon(\sigma) \Delta(\sigma),$$

and we obtain

$$k \Upsilon(\tilde{\sigma}) \Delta(\tilde{\sigma}) \geq (1 - \eta^2) \sqrt{k} \Upsilon(\sigma) \Delta(\sigma).$$

The final result follows since $\frac{\Delta(\sigma)}{\Delta(\tilde{\sigma})} \geq \frac{1}{1 + \xi_0} \geq \frac{4}{5}$. \square

In order to obtain a bound on the circumradius of $\tilde{\sigma}$ with respect to that of σ , it is convenient to find an isometry that maps the vertices of σ close to the vertices of $\tilde{\sigma}$. Choosing \tilde{p}_0 and p_0 to coincide at the origin, the displacement error for the remaining vertices is minimised by taking the orthogonal polar factor of the linear transformation $A = \tilde{P}P^{-1}$ that maps σ to $\tilde{\sigma}$. In other words, if the singular value decomposition of A is $A = U_A \Sigma_A V_A^\top$, then $A = \Phi S$, where $S = V_A \Sigma_A V_A^\top$, and $\Phi = U_A V_A^\top$ is the desired linear isometry. We have the following result, which is a special case of a theorem demonstrated by Jiménez and Petrova [JP13]:

Lemma A.4 (Close alignment of bases) Suppose that P and \tilde{P} are non-degenerate $k \times k$ matrices such that

$$\tilde{P}^\top \tilde{P} = P^\top P + E. \quad (9)$$

Then there exists a linear isometry $\Phi : \mathbb{R}^k \rightarrow \mathbb{R}^k$ such that

$$\|\tilde{P} - \Phi P\| \leq \frac{s_1(P)s_1(E)}{s_k(P)^2}.$$

Proof Multiplying by $P^{-T} := (P^\top)^{-1}$ on the left, and by P^{-1} on the right, we rewrite Equation (9) as

$$A^\top A = I + F, \quad (10)$$

where $A = \tilde{P}P^{-1}$, and $F = P^{-T}EP^{-1}$. Using the singular value decomposition $A = U_A \Sigma_A V_A^\top$, we let $\Phi = U_A V_A^\top$, and we find

$$\tilde{P} - \Phi P = (A - \Phi)P = U_A(\Sigma_A - I)V_A^\top P. \quad (11)$$

From Equation (10) we deduce that $s_1(A)^2 \leq 1 + s_1(F)$, and also that $s_k(A)^2 \geq 1 - s_1(F)$. It follows that

$$\max_i |s_i(A) - 1| \leq \frac{s_1(F)}{1 + s_i(A)} \leq s_1(F),$$

and thus

$$\|\Sigma_A - I\| \leq s_1(F) \leq s_1(P^{-1})^2 s_1(E) = s_k(P)^{-2} s_1(E).$$

The result now follows from Equation (11). \square

Recalling that an upper bound on the norm of a matrix also serves as an upper bound on the norm of its column vectors, we obtain the following immediate consequence of Lemma A.4, using Lemma A.2 and Lemma A.1:

Lemma A.5 (Close alignment of simplices) Suppose that $\sigma = \{p_0, \dots, p_k\}$ and $\tilde{\sigma} = \{\tilde{p}_0, \dots, \tilde{p}_k\}$ are two k -simplices in \mathbb{R}^m such that

$$\| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \| \leq \xi_0 \Delta(\sigma),$$

for all $0 \leq i < j \leq k$. Let P be the matrix whose i^{th} column is $p_i - p_0$, and define \tilde{P} similarly. If $\xi_0 \leq \frac{2}{3}$, then there exists an isometry $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}^m$ such that

$$\|\tilde{P} - \Phi P\| \leq \frac{4\sqrt{k}\xi_0\Delta(\sigma)}{\Upsilon(\sigma)^2},$$

and if $\hat{\sigma} = \Phi\sigma = \{\hat{p}_0, \dots, \hat{p}_k\}$, then $\hat{p}_0 = \tilde{p}_0$, and

$$\|\hat{p}_i - \tilde{p}_i\| \leq \frac{4\sqrt{k}\xi_0\Delta(\sigma)}{\Upsilon(\sigma)^2} \quad \text{for all } 1 \leq i \leq k.$$

Using Lemma A.5 together with [BDG13c, Lemma 4.3] we obtain a bound on the difference in the circumradii of two simplices whose edge lengths are almost the same:

Lemma A.6 (Circumradii under distortion) Suppose that $\sigma = \{p_0, \dots, p_k\}$ and $\tilde{\sigma} = \{\tilde{p}_0, \dots, \tilde{p}_k\}$ are two k -simplices in \mathbb{R}^m such that

$$\| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \| \leq \xi_0 \Delta(\sigma),$$

for all $0 \leq i < j \leq k$. If

$$\xi_0 \leq \left(\frac{\Upsilon(\sigma)}{4} \right)^2,$$

then

$$|R(\tilde{\sigma}) - R(\sigma)| \leq \frac{16k^{\frac{3}{2}}R(\sigma)\xi_0}{\Upsilon(\sigma)^3}.$$

Proof We define $\hat{\sigma} = \Phi\sigma$, where $\Phi : \sigma \rightarrow \text{aff}(\tilde{\sigma})$ is the isometry described in Lemma A.5. Since $\hat{p}_0 = \tilde{p}_0$, and $R(\hat{\sigma}) = R(\sigma)$, we have $|R(\tilde{\sigma}) - R(\sigma)| \leq \|C(\hat{\sigma}) - C(\tilde{\sigma})\|$. By Lemma A.4, the distances between $C(\sigma)$ and the vertices of $\tilde{\sigma}$ are all bounded by

$$R(\sigma) + \frac{4\sqrt{k}\xi_0\Delta(\sigma)}{\Upsilon(\sigma)^2} \leq \left(1 + \frac{\sqrt{k}}{2}\right)R(\sigma) \leq \frac{3\sqrt{k}}{2}R(\sigma),$$

and these distances differ by no more than

$$\frac{8\sqrt{k}\xi_0\Delta(\sigma)}{\Upsilon(\sigma)^2}.$$

It follows then from [BDG13c, Lemma 4.3] that

$$\begin{aligned} \|C(\hat{\sigma}) - C(\tilde{\sigma})\| &\leq \frac{\frac{3\sqrt{k}}{2}R(\sigma)}{\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma})} \left(\frac{8\sqrt{k}\xi_0\Delta(\sigma)}{\Upsilon(\sigma)^2} \right) \\ &\leq \frac{12kR(\sigma)\xi_0}{\frac{3}{4\sqrt{k}}\Upsilon(\sigma)^3} \quad \text{by Lemma A.3, with } \eta = \frac{1}{2} \\ &\leq \frac{16k^{\frac{3}{2}}R(\sigma)\xi_0}{\Upsilon(\sigma)^3}. \end{aligned}$$

□

A.2 Circumcentres and distortion maps

It is convenient to introduce the affine space $N(\sigma)$, which is the space of centres of circumscribing balls for a simplex $\sigma \in \mathbb{R}^m$. If σ is a non-degenerate k -simplex, then $N(\sigma)$ is an affine space of dimension $m - k$ perpendicular to $\text{aff}(\sigma)$ and containing $C(\sigma)$.

The transition functions introduce a small metric distortion, which motivated our interest in the properties of perturbed simplices. In order to extend the perturbation algorithm [BDG13b] to the setting of curved manifolds, we are interested in quantifying how the test for the hoop property behaves under a perturbation of the interpoint distances. Specifically, if a point p is at a distance $\alpha_0 R$ from the diametric sphere of a simplex σ in one coordinate frame, what can we say about the distance of p from $S^{m-1}(\sigma)$ when measured by the metric of another coordinate frame? To this end, we are interested in the behaviour of the circumcentre under the influence of a mapping that is not distance preserving. As a first step in this direction, we observe another consequence of [BDG13c, Lemma 4.3]:

Lemma A.7 (Circumscribing balls under distortion) Suppose $\phi : \mathbb{R}^m \supset U \rightarrow V \subset \mathbb{R}^m$ is a homeomorphism such that, for some positive ξ_0 ,

$$|d(x, y) - d(\phi(x), \phi(y))| \leq \xi_0 d(x, y) \quad \text{for all } x, y \in U.$$

Suppose also that $\sigma \subset U$ is a k -simplex, and that $B(c, r)$ is a circumscribing ball for σ with $c \in U$. Let $\tilde{\sigma} = \phi(\sigma)$. If

$$\xi_0 \leq \left(\frac{\Upsilon(\sigma)}{4} \right)^2,$$

then there is a circumscribing ball $B(\tilde{c}, \tilde{r})$ for $\tilde{\sigma}$ such that

$$d(\phi(c), \tilde{c}) \leq \frac{3\sqrt{k}r^2\xi_0}{\Upsilon(\sigma)\Delta(\sigma)},$$

and

$$|\tilde{r} - r| \leq \frac{5\sqrt{k}r^2\xi_0}{\Upsilon(\sigma)\Delta(\sigma)}.$$

Proof By the perturbation bounds on ϕ , the distances between $\phi(c)$ and the vertices of $\tilde{\sigma}$ differ by no more than $2\xi_0r$, and these distances are all bounded by $(1 + \xi_0)r$. In this context [BDG13c, Lemma 4.3] says that there exists a $\tilde{c} \in N(\tilde{\sigma})$ such that

$$d(\phi(c), \tilde{c}) \leq \frac{(1 + \xi_0)r2\xi_0r}{\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma})}.$$

We apply Lemma A.3, using $\eta = \frac{1}{2}$, to obtain $\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma}) \geq \frac{3}{4\sqrt{k}}\Upsilon(\sigma)\Delta(\sigma)$. We find

$$d(\phi(c), \tilde{c}) \leq \frac{8\sqrt{k}(1 + \xi_0)r^2\xi_0}{3\Upsilon(\sigma)\Delta(\sigma)}.$$

The announced bound on $d(\phi(c), \tilde{c})$ is obtained by observing that $\xi_0 \leq \frac{1}{16}$.

Choosing a vertex $\tilde{p} = \phi(p) \in \tilde{\sigma}$, the bound on the difference in the radii follows:

$$\begin{aligned} \tilde{r} = d(\tilde{p}, \tilde{c}) &\geq d(\tilde{p}, \phi(c)) - d(\phi(c), \tilde{c}) \\ &\geq r - \xi_0r - \frac{3\sqrt{k}r^2\xi_0}{\Upsilon(\sigma)\Delta(\sigma)} \\ &\geq r - \frac{5\sqrt{k}r^2\xi_0}{\Upsilon(\sigma)\Delta(\sigma)}, \end{aligned}$$

and similarly for the upper bound. \square

We will find it convenient to have a bound on the circumradius of a simplex, relative to its thickness and longest edge length:

Lemma A.8 If σ is a non-degenerate simplex in \mathbb{R}^m , then

$$R(\sigma) \leq \frac{\Delta(\sigma)}{2\Upsilon(\sigma)}.$$

Proof Let $\sigma = \{p_0, \dots, p_k\}$, We work in $\mathbb{R}^k = \text{aff}(\sigma) \subset \mathbb{R}^m$, and let P be the $k \times k$ matrix whose i^{th} column is $p_i - p_0$. Then, by equating $\|C(\sigma) - p_0\|^2$ with $\|C(\sigma) - p_i\|^2$ and expanding, we find a system of equations that may be written in matrix form as

$$P^\top C(\sigma) = b,$$

where the i^{th} component of the vector b is $\frac{1}{2}(\|p_i\|^2 - \|p_0\|^2)$. Choosing p_0 as the origin, we have $\|C(\sigma)\| = R(\sigma)$, and $\|b\| \leq \frac{1}{2}\sqrt{k}\Delta(\sigma)^2$. Since $s_1(P^{-T}) = s_k(P)^{-1}$, the result follows from Lemma A.1, which says $s_k(P) \geq \sqrt{k}\Upsilon(\sigma)\Delta(\sigma)$. \square

Using the bound on $d(\phi(C(\sigma)), N(\tilde{\sigma}))$ given by Lemma A.7, together with the circumradius bound of Lemma A.6, we obtain a bound on $d(\phi(C(\sigma)), C(\tilde{\sigma}))$ by means of the Pythagorean theorem:

Lemma A.9 (Circumcentres under distortion) Suppose $\phi : \mathbb{R}^m \supset U \rightarrow V \subset \mathbb{R}^m$ is a homeomorphism such that

$$|d(x, y) - d(\phi(x), \phi(y))| \leq \xi_0 d(x, y) \quad \text{for all } x, y \in U.$$

Suppose also that $\sigma \subset U$ is a k -simplex, and let $\tilde{\sigma} = \phi(\sigma)$. If

$$\xi_0 \leq \left(\frac{\Upsilon(\sigma)}{4} \right)^2,$$

then

$$d(\phi(C(\sigma)), C(\tilde{\sigma})) \leq \left[\left(\frac{42k^2}{\Upsilon(\sigma)^3} \right) \xi_0 \right]^{\frac{1}{2}} R(\sigma).$$




Figure 1: Diagram for the proof of Lemma A.9.

Proof Let c be the closest point in $N(\tilde{\sigma})$ to $\phi(C(\sigma))$, and let w be the distance from c to $\phi(C(\sigma))$. Setting z as the distance between c and $C(\tilde{\sigma})$, we have that $d(\phi(C(\sigma)), C(\tilde{\sigma}))^2 = z^2 + w^2$; see Figure 1. Let \hat{c} be the orthogonal projection of $\phi(C(\sigma))$ into $\text{aff}(\tilde{\sigma})$. Then, letting $R = R(\sigma)$, and $\tilde{R} = R(\tilde{\sigma})$, and choosing $\tilde{p} = \phi(p) \in \tilde{\sigma}$, we have

$$\begin{aligned} z^2 &= d(\phi(C(\sigma)), \tilde{p})^2 - d(\tilde{p}, \hat{c})^2 \\ &\leq (1 + \xi_0)^2 R^2 - (\tilde{R} - w)^2 \\ &= R^2 - \tilde{R}^2 + 2\tilde{R}w + 2R^2\xi_0 + \xi_0^2 R^2 - w^2. \end{aligned}$$

Using Lemma A.6, we write \tilde{R} in terms of R , as $|R - \tilde{R}| \leq sR$, where

$$s = \frac{16k^{\frac{3}{2}} \xi_0}{\Upsilon(\sigma)^3}.$$

Then using Lemma A.7 to bound w , and writing Δ , and Υ , instead of $\Delta(\sigma)$ and

$\Upsilon(\sigma)$, we find

$$\begin{aligned}
d(\phi(C(\sigma)), C(\tilde{\sigma}))^2 &\leq R^2 - (1-s)^2 R^2 + 2w(1+s)R + 2\xi_0 R^2 + \xi_0^2 R^2 \\
&\leq 2(sR + w + ws)R + (2 + \xi_0)\xi_0 R^2 \\
&\leq \left[2 \left(\frac{16k^{\frac{3}{2}}}{\Upsilon^3} + \frac{3\sqrt{k}R}{\Upsilon\Delta} + \frac{54k^2 R \xi_0}{\Upsilon^4 \Delta} \right) + (2 + \xi_0) \right] R^2 \xi_0 \\
&\leq \left[2 \left(\frac{16k^{\frac{3}{2}}}{\Upsilon^3} + \frac{3\sqrt{k}}{2\Upsilon^2} + \frac{27k^2}{16\Upsilon^3} \right) + 3 \right] R^2 \xi_0 \quad \text{using Lemma A.8} \\
&\leq \left[\frac{42k^2}{\Upsilon^3} \right] R^2 \xi_0.
\end{aligned}$$

□

A.3 The size of the domains

The domains U_{ij} on which the transition functions are defined need to be large enough to accommodate two distinct requirements. First, the domain of the transition function φ_{ji} must contain a large enough neighbourhood of p'_i that we can apply the metric stability result of [BDG13c] to ensure that $\text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$ will be the same as $\text{star}(p'_i; \text{Del}(\mathbf{P}'_j))$ whenever $p'_j \in \text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$. The second requirement is that any potential forbidden configuration in the region of interest must lie entirely within the domain of the transition function associated with each of its vertices.

We recall the stability result [BDG13c, Theorem 4.17] that we will use:

Lemma A.10 (Delaunay stability under metric perturbation) Suppose \mathbf{Q}'_i is a (μ'_0, ϵ'_i) -net and $\text{conv}(\mathbf{Q}'_i) \subseteq U \subset \mathbb{R}^m$ and $d_j : U \times U \rightarrow \mathbb{R}$ is such that $|d_i(x, y) - d_j(x, y)| \leq \xi$ for all $x, y \in U$. Suppose also that $\mathbf{S} \subseteq \mathbf{Q}'_i$ is a set of interior points such that every m -simplex $\sigma \in \text{star}(\mathbf{S})$ is Γ_0^m -thick and δ -protected and satisfies $d_i(p, \partial U) \geq 2\epsilon'_i$ for every vertex $p \in \sigma$. If

$$\xi \leq \frac{\Gamma_0^m \mu'_0}{36} \delta,$$

then

$$\text{star}(\mathbf{S}; \text{Del}_{d_j}(\mathbf{Q}'_i)) = \text{star}(\mathbf{S}; \text{Del}_{d_i}(\mathbf{Q}'_i)).$$

The notation $\text{Del}_{d_j}(\mathbf{Q}'_i)$ in Lemma A.10 means that the metric d_j is used to compute the Delaunay complex of \mathbf{Q}'_i . For our purposes, d_j is the pullback by φ_{ji} of the Euclidean metric on U_j . Thus we have the identification

$$\text{star}(\mathbf{S}; \text{Del}_{d_j}(\mathbf{Q}'_i)) \cong \text{star}(\varphi_{ji}(\mathbf{S}); \text{Del}(\varphi_{ji}(\mathbf{Q}'_i))).$$

We will use $\mathbf{S} = \{p'_i\}$. Some argument is required to ensure that Lemma A.10 provides a route to the desired equivalence

$$\text{star}(p'_i; \text{Del}(\mathbf{P}'_i)) \cong \text{star}(p'_i; \text{Del}(\mathbf{P}'_j)), \quad \text{when } p'_j \in \text{star}(p'_i; \text{Del}(\mathbf{P}'_i)). \quad (12)$$

We first establish our “region of interest”. We demand, for all $i \in \mathcal{N}$, that \mathbf{P}_i be a (μ_0, ϵ_i) -net for $B_i(p_i, 8\epsilon_i)$, and we define $\mathbf{Q}'_i = \mathbf{P}'_i \cap B_i(p_i, 6\epsilon_i)$. Since \mathbf{P}'_i changes as the algorithm progresses, points may come and go from \mathbf{Q}'_i , but we will ensure that when the algorithm terminates, \mathbf{Q}'_i will contain no forbidden configurations.

Lemma A.11 For all $i \in \mathcal{N}$ we have

$$\text{star}(p'_i; \text{Del}(\mathbf{Q}'_i)) = \text{star}(p'_i; \text{Del}(\mathbf{P}'_i)).$$

and if $p \in \text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$, then $d_i(p, \partial B_i(p_i, 6\epsilon_i)) > 2\epsilon'_i$.

If $B_i(p_i, 6\epsilon_i) \subseteq U_{ij}$ whenever $p'_j \in \text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$, then

$$\text{star}(p'_i; \text{Del}(\varphi_{ji}(\mathbf{Q}'_i))) = \text{star}(p'_i; \text{Del}(\mathbf{P}'_j)).$$

Proof We have $d_i(p'_i, \partial B_i(p_i, 6\epsilon_i)) \geq \frac{24}{5}\epsilon'_i - \frac{1}{4}\epsilon'_i > 4\epsilon'_i$. The density assumption guarantees that if $\sigma^m \in \text{star}(p'_i; \text{Del}(\mathbf{Q}'_i))$, then $R(\sigma^m) < \epsilon'_i$, and the observation that $B_i(C(\sigma^m), R(\sigma^m)) \subset B_i(p_i, 4\epsilon_i)$, leads to the first equality, and the bound on the distance from p to $\partial B_i(p_i, 6\epsilon_i)$.

The second equality follows from two observations. First we show that if $\sigma^m \in \text{star}(p'_i; \text{Del}(\varphi_{ji}(\mathbf{Q}'_i)))$, then $R(\sigma) < \epsilon'_j$. Since $\varphi_{ji}(\mathbf{Q}'_i) \subset \mathbf{P}'_j$, and \mathbf{P}'_j is ϵ'_j -dense for $B = B_j(p_j, 8\epsilon_j)$, it is sufficient to show that $d_j(p'_i, \partial B) \geq 2\epsilon'_j$. Since $p'_j \in \text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$, we have $d_j(p'_i, p'_j) \leq (1 + \xi_0)d_i(p'_i, p'_j) \leq (1 + \xi_0)2\epsilon'_i \leq 2(1 + \xi_0)(1 + \epsilon_0)\frac{5}{4}\epsilon_j \leq 5\epsilon_j$. Thus since $d_j(p_j, p'_j) \leq \frac{1}{4}\epsilon_j$, we have $d_j(p'_j, \partial B) \geq 8\epsilon_j - \frac{21}{4}\epsilon_j = \frac{11}{4}\epsilon_j \geq \frac{11}{5}\epsilon'_j$. This establishes that the Delaunay ball for σ^m must remain empty when points outside of B are considered.

The second observation required to establish the second equality is that if $q \in \mathbf{P}'_j$ is such that $d_j(p'_i, q) < \epsilon'_j$, then $q \in \varphi_{ji}(B_i(p_i, 6\epsilon_i))$. Indeed, we have $d_i(p'_i, q) \leq 2(1 + \xi_0)\epsilon'_j \leq 2(1 + \xi_0)(1 + \epsilon_0)\frac{5}{4}\epsilon_i \leq 5\epsilon_i$. The result follows since $d_i(p_i, p'_i) \leq \frac{1}{4}\epsilon_i$. \square

If $p'_j \in \text{star}(p'_i; \text{Del}(\mathbf{P}'_i))$, then $d_i(p_i, p_j) < \frac{1}{2}\epsilon_i + 2\epsilon'_i \leq 3\epsilon_i$. Thus Lemma A.11 establishes the first requirement on U_{ij} , namely

$$B_i(p_i, 6\epsilon_i) \subset U_{ij} \quad \text{if } d_i(p_i, p_j) < 3\epsilon_i. \quad (13)$$

The second requirement arises from the fact that we wish to ensure that there are no forbidden configurations in \mathbf{Q}'_i . This will be sufficient for us to apply Lemma A.10.

Lemma A.12 (Protected stars) If there are no forbidden configurations in \mathbf{Q}'_i , then all the m -simplices in $\text{star}(p'_i; \text{Del}(\mathbf{Q}'_i))$ are Γ_0 -good and δ -protected, with $\delta = \delta_0\mu'_0\epsilon'_i$.

Proof Since \mathbf{P}'_i is a (μ'_0, ϵ'_i) -net for $B_i(p_i, 8\epsilon_i)$, it follows that \mathbf{Q}'_i is a (μ'_0, ϵ'_i) -net. Thus if there are no forbidden configurations in \mathbf{Q}'_i , then by [BDG13b, Lemma 3.5], all the m -simplices in $\text{Del}_1(\mathbf{Q}'_i)$ will be Γ_0 -good and δ -protected, with $\delta = \delta_0\mu'_0\epsilon'_i$.

The sampling criteria ensure that every point on $\partial \text{conv}(\mathbf{Q}'_i)$ must be at a distance of less than $2\epsilon'_i$ from $\partial B_i(p_i, 6\epsilon_i)$. Thus $d_i(p_i, \partial \text{conv}(\mathbf{Q}'_i)) > 6\epsilon_i - 2\epsilon'_i \geq \frac{14}{5}\epsilon'_i$. Also, $d_i(p_i, p'_i) \leq \frac{\epsilon'_i}{4}$, and we find that $d_i(p'_i, \partial \text{conv}(\mathbf{Q}'_i)) \geq \frac{51}{20}\epsilon'_i$. Thus, since $d_i(p'_i, C(\sigma)) < \epsilon'_i$ if σ is in $\text{star}(p'_i; \text{Del}(\mathbf{Q}'_i))$, we have $\text{star}(p'_i; \text{Del}(\mathbf{Q}'_i)) \subseteq \text{Del}_1(\mathbf{Q}'_i)$, and hence the result. \square

According to Lemma 2.2 $\mathcal{P}3$, if τ is a forbidden configuration in \mathbf{Q}'_i , then $\Delta(\tau) < \frac{15}{4}\epsilon_i$, and it follows that if $p'_j \in \tau$, then $\tau \subset B_i(p_j, 4\epsilon_i)$. We will require that each potential forbidden configuration in \mathbf{Q}'_i lies within the domain of any transition function associated one of its vertices. Thus we demand that

$$B_i(p_j, 4\epsilon_i) \cap B_i(p_i, 6\epsilon_i) \subset U_{ij} \quad \text{if } p_j \in B_i(p_i, 6\epsilon_i). \quad (14)$$

For simplicity we accommodate Equations (13) and (14) by demanding that

$$B_i(p_j, 9\epsilon_i) \cap B_i(p_i, 6\epsilon_i) \subset U_{ij} \quad \text{if } p_j \in B_i(p_i, 6\epsilon_i). \quad (15)$$

In summary, Lemmas A.10, A.11, and A.12 combine to yield the desired equivalence of stars (12), under the assumption that \mathcal{Q}'_i has no forbidden configurations. We take $U = B_i(p_i, 6\epsilon_i)$ in Lemma A.10, and Equation (3) yields $\xi \leq \xi_0 12\epsilon_i$. Using $\delta = \Gamma_0^{m+1} \mu'_0 \epsilon'_i \geq \frac{1}{2} \Gamma_0^{m+1} \mu_0 \epsilon_i$ we have:

Lemma A.13 (Stable stars) If

$$\xi_0 \leq \frac{\Gamma_0^{2m+1} \mu_0^2}{2^{12}},$$

and there are no forbidden configurations in \mathcal{Q}'_i , then for all $p'_j \in \text{star}(p'_i; \text{Del}(\mathcal{P}'_i))$, we have

$$\text{star}(p'_i; \text{Del}(\mathcal{P}'_i)) \cong \text{star}(p'_i; \text{Del}(\mathcal{P}'_j)).$$

We have established minimal requirements on the size of the domains U_{ij} , but these requirements may implicitly demand more. Although $\varphi_{ji} : U_{ij} \rightarrow U_{ji}$ is close to an isometry, ϵ_i may be almost twice as large as ϵ_j . Thus the requirement on U_{ij} may imply that $U_{ji} = \varphi_{ji}(U_{ij})$ is significantly larger than Equation (15) demands.

Clearly we must have

$$\bigcup_{j \in \mathcal{N}_i} U_{ij} \subset U_i.$$

We have also explicitly demanded that \mathcal{P}_i be a (μ_0, ϵ_i) -net for $B_i(p_i, 8\epsilon_i)$. We will assume that $B_i(p_i, 8\epsilon_i) \subset U_i$.

A.4 Hoop distortion

We will rely primarily on Properties $\mathcal{P}1$ and $\mathcal{P}4$ of forbidden configurations (Lemma 2.2), and the stability of the circumcentres exhibited by Lemma A.9. We have the following observation about the properties of forbidden configurations under the influence of the transition functions:

Lemma A.14 Assume $\xi_0 \leq \left(\frac{\Gamma_0^k}{4}\right)^2$. If $\tau = p'_i * \sigma \subset \mathcal{Q}'_j \subset U_j$ is a forbidden configuration, where σ is a k -simplex, then $\tilde{\sigma} = \varphi_{ij}(\sigma) \subset \mathcal{P}'_i$ is $\tilde{\Gamma}_0^k$ -thick, with

$$\tilde{\Gamma}_0^k = \frac{2}{5\sqrt{k}} \Gamma_0^k,$$

has a radius satisfying

$$R(\tilde{\sigma}) \leq 2 \left(1 + \frac{16k^{\frac{3}{2}} \xi_0}{\Gamma_0^{3k}} \right) (1 + \epsilon_0) \epsilon_i,$$

and $d_i(p'_i, S^{m-1}(\tilde{\sigma})) \leq 2\tilde{\alpha}_0 \epsilon_i$, where

$$\tilde{\alpha}_0 = \left(\alpha_0 (1 + \xi_0) + \left(\frac{12k^{\frac{3}{2}}}{\Gamma_0^{2k}} \right) \xi_0^{\frac{1}{2}} \right) (1 + \epsilon_0).$$

Proof The bound for $\tilde{\Gamma}_0^k$ follows immediately from Lemma A.3, and the fact that σ is Γ_0^k -thick (Lemma 2.2 $\mathcal{P}4$). Likewise, the radius bound is a direct consequence of Lemma A.6 and Lemma 2.2 $\mathcal{P}2$.

The bound on $\tilde{\alpha}_0$ is obtained from Property $\mathcal{P}1$ with the aid of Lemmas A.6 and A.9. We have $d_i(p'_i, S^{m-1}(\tilde{\sigma})) = |d_i(p'_i, C(\tilde{\sigma})) - R(\tilde{\sigma})|$, and we are able to get a tighter upper bound on $R(\tilde{\sigma}) - d_i(p'_i, C(\tilde{\sigma}))$, than we can for $d_i(p'_i, C(\tilde{\sigma})) - R(\tilde{\sigma})$. Thus

$$\begin{aligned}
d_i(p'_i, S^{m-1}(\tilde{\sigma})) &= |d_i(p'_i, C(\tilde{\sigma})) - R(\tilde{\sigma})| \\
&\leq (1 + \xi_0)d_j(p'_i, C(\sigma)) + d_i(\varphi_{ij}(C(\sigma)), C(\tilde{\sigma})) - (R(\sigma) - |R(\tilde{\sigma}) - R(\sigma)|) \\
&\leq (1 + \xi_0)(\alpha_0 R(\sigma) + R(\sigma)) - R(\sigma) + d_i(\varphi_{ij}(C(\sigma)), C(\tilde{\sigma})) + |R(\tilde{\sigma}) - R(\sigma)| \\
&\leq \left(\alpha_0(1 + \xi_0) + \xi_0 + \left[\left(\frac{42k^2}{\Upsilon(\sigma)^3} \right) \xi_0 \right]^{\frac{1}{2}} + \frac{16k^{\frac{3}{2}}\xi_0}{\Upsilon(\sigma)^3} \right) R(\sigma) \\
&\leq 2 \left(\alpha_0(1 + \xi_0) + \xi_0 + \left[\left(\frac{42k^2}{\Gamma_0^{3k}} \right) \xi_0 \right]^{\frac{1}{2}} + \frac{16k^{\frac{3}{2}}\xi_0}{\Gamma_0^{3k}} \right) \epsilon_j \\
&\leq 2 \left(\alpha_0(1 + \xi_0) + \left(\xi_0^{\frac{1}{2}} + \frac{7k}{\Gamma_0^{\frac{3k}{2}}} + \frac{16k^{\frac{3}{2}}\xi_0^{\frac{1}{2}}}{\Gamma_0^{3k}} \right) \xi_0^{\frac{1}{2}} \right) (1 + \epsilon_0)\epsilon_i \\
&\leq 2 \left(\alpha_0(1 + \xi_0) + \left(\frac{12k^{\frac{3}{2}}}{\Gamma_0^{2k}} \right) \xi_0^{\frac{1}{2}} \right) (1 + \epsilon_0)\epsilon_i.
\end{aligned}$$

□

We have abused the notation slightly because $\tilde{\tau} = \varphi_{ij}(\tau)$ need not actually satisfy the $\tilde{\alpha}_0$ -hoop property definition $d_i(p, S(\tilde{\tau}_p)) \leq \tilde{\alpha}_0 R(\tilde{\tau}_p)$, because $R(\tilde{\tau})$ may be larger than $2\epsilon_i$. However we are not concerned with the $\tilde{\alpha}_0$ -property for $\tilde{\tau}$; instead we desire a condition that will permit the extended algorithm to emulate the original Euclidean perturbation algorithm [BDG13b], and guarantee that forbidden configurations such as τ cannot exist in any of the sets \mathcal{Q}'_j .

The bounds in Lemma A.14 can be further simplified. We have announced them in this intermediate state in order to elucidate the roles played by ξ_0 and ϵ_0 . In particular, there is no need to significantly constrain ϵ_0 . The original perturbation algorithm for points in Euclidean space [BDG13b] extends to the case of a non-constant sampling radius simply by replacing α_0 by $\tilde{\alpha}_0 \leq (1 + \epsilon_0)\alpha_0 \leq 2\alpha_0$, as can be seen by setting $\xi_0 = 0$ in the expression for $\tilde{\alpha}_0$.

In the general case of interest here, we see from the expression for $\tilde{\alpha}_0$ presented in Lemma A.14, that ξ_0 must be considerably more constrained with respect to Γ_0 if we are to obtain an expression for $\tilde{\alpha}_0$ that goes to zero as Γ_0 goes to zero. For the purposes of the algorithm, we do not require the bounds on the radius or the thickness.

Lemma A.15 If

$$\xi_0 \leq \left(\frac{\Gamma_0^{2m+1}}{4} \right)^2,$$

then for any forbidden configuration $\tau = p'_j * \sigma \in \mathcal{Q}'_i$, there is a simplex $\tilde{\sigma} = \varphi_{ji}(\sigma) \subset \mathcal{P}'_j$ such that $d_j(p'_j, S^{m-1}(\tilde{\sigma})) \leq 2\tilde{\alpha}_0\epsilon_j$, where

$$\tilde{\alpha}_0 = \frac{2^{16}m^{\frac{3}{2}}\Gamma_0}{\mu_0^3}.$$

Proof By the properties of a forbidden configuration, σ is a k -simplex with $k \leq m$. From Lemma A.14,

$$\begin{aligned}\tilde{\alpha}_0 &= \left(\alpha_0(1 + \xi_0) + \left(\frac{12k^{\frac{3}{2}}}{\Gamma_0^{2k}} \right) \xi_0^{\frac{1}{2}} \right) (1 + \epsilon_0) \\ &\leq 2 \left(\frac{2^{13}\Gamma_0}{\mu_0^3} (1 + \xi_0) + \left(\frac{12m^{\frac{3}{2}}}{\Gamma_0^{2m}} \right) \frac{\Gamma_0^{2m+1}}{4} \right) \\ &< \left((1 + 2^{-4}) + m^{\frac{3}{2}} \right) \frac{2^{14}\Gamma_0}{\mu_0^3} \\ &< \frac{2^{16}m^{\frac{3}{2}}\Gamma_0}{\mu_0^3}.\end{aligned}$$

□

References

- [AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22(4):481–504, 1999. 10
- [BDG13a] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Constructing intrinsic Delaunay triangulations of submanifolds. Research Report RR-8273, INRIA, 2013. arXiv:1303.6493. 4, 10, 15
- [BDG13b] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Constructing stable Delaunay triangulations. Research Report RR-8275, INRIA, 2013. 4, 5, 7, 11, 12, 13, 20, 24, 26
- [BDG13c] J.-D. Boissonnat, R. Dyer, and A. Ghosh. The stability of Delaunay triangulations. Research Report RR-8276, INRIA, 2013. (to appear in *Int. J. Comp. Geom. & Appl.* special issue for SoCG 2012). 4, 5, 6, 12, 15, 16, 17, 19, 20, 21, 23
- [BG11] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. Technical Report RR-7142 v.3, INRIA, 2011. (To appear in *DCG*). 5, 15, 16
- [Boo86] W. M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry*. Academic Press, Orlando, Florida, second edition, 1986. 9
- [BWY11] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic Delaunay mesh generation. Technical Report RR-7712, INRIA, 2011. (To appear in *SIAM J. of Computing*). 5
- [CDR05] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA*, pages 1018–1027, 2005. 5
- [Cha06] I. Chavel. *Riemannian Geometry, A modern introduction*. Cambridge, 2nd edition, 2006. 15
- [Dye10] R. Dyer. *Self-Delaunay meshes for surfaces*. PhD thesis, Simon Fraser University, Burnaby, Canada, 2010. 16

- [ELM⁺00] H. Edelsbrunner, X.Y. Li, G. Miller, A. Stathopoulos, D. Talmor, S.H. Teng, A. Üngör, and N. Walkington. Smoothing and cleaning up slivers. In *STOC*, pages 273–277. ACM, 2000. 4
- [Gal11] J. Gallier. Ch. 4: Construction of manifolds from gluing data. In *Notes on differential geometry and Lie groups*. 2011. Unpublished manuscript accessed June 2013: <http://www.cis.upenn.edu/~jean/gbooks/manif.html>. 5
- [GKS00] M. Gopi, S. Krishnan, and C. T. Silva. Surface reconstruction based on lower dimensional localized Delaunay triangulation. volume 19, pages 467–478, 2000. 4
- [Gli05] D. Glickenstein. Geometric triangulations and discrete Laplacians on manifolds, 2005. [arxiv:math.MG/0508188v1](http://arxiv.org/abs/math.MG/0508188v1). 16
- [GW04] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. *Discrete & Comp. Geom.*, 32:245–267, 2004. 5
- [HKV12] A. N. Hirani, K. Kalyanaraman, and E. B. VanderZee. Delaunay Hodge star. *Computer-Aided Design*, 45(2):540–544, 2012. 16
- [JP13] D. Jiménez and G. Petrova. On matching point configurations. Preprint accessed 2013.05.17: <http://www.math.tamu.edu/~gpetrova/JP.pdf>, 2013. 18
- [LL00] G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *SoCG*, pages 341–349, 2000. 4