

An Obstruction to Delaunay Triangulations in Riemannian Manifolds

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh, Martynchuk Nikolay

► **To cite this version:**

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh, Martynchuk Nikolay. An Obstruction to Delaunay Triangulations in Riemannian Manifolds. Discrete and Computational Geometry, Springer Verlag, 2017, 10.1145/336154.336221 . hal-01583073

HAL Id: hal-01583073

<https://hal.inria.fr/hal-01583073>

Submitted on 6 Sep 2017

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

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

An obstruction to Delaunay triangulations in Riemannian manifolds

Jean-Daniel Boissonnat* Ramsay Dyer† Arijit Ghosh‡ Nikolay Martynchuk§

July 13, 2017

Abstract

Delaunay has shown that the Delaunay complex of a finite set of points P of Euclidean space \mathbb{R}^m triangulates the convex hull of P , provided that P satisfies a mild genericity property. Voronoi diagrams and Delaunay complexes can be defined for arbitrary Riemannian manifolds. However, Delaunay's genericity assumption no longer guarantees that the Delaunay complex will yield a triangulation; stronger assumptions on P are required. A natural one is to assume that P is sufficiently dense. Although results in this direction have been claimed, we show that sample density alone is insufficient to ensure that the Delaunay complex triangulates a manifold of dimension greater than 2.

keywords: Delaunay complex, triangulation, Riemannian manifold

MSC2010: 57R05 (Primary) 54B15, 53B20 (Secondary)

1 Delaunay complex and Delaunay triangulation

Let $(\mathcal{M}, d_{\mathcal{M}})$ be a metric space, and let P be a finite set of points of \mathcal{M} . An *empty ball* is an open ball in the metric $d_{\mathcal{M}}$ that contains no point from P . We say that an empty ball B is *maximal* if no other empty ball with the same centre strictly contains B . A *Delaunay ball* is a maximal empty ball.

A simplex σ is a *Delaunay simplex* if there exists some Delaunay ball B that circumscribes σ , i.e., such that the vertices of σ belong to $\partial B \cap P$. The *Delaunay complex* is the set of Delaunay simplices, and is denoted $\text{Del}_{\mathcal{M}}(P)$. It is an abstract simplicial complex and so defines a topological space, $|\text{Del}_{\mathcal{M}}(P)|$, called its *carrier*. We say that $\text{Del}_{\mathcal{M}}(P)$ *triangulates* \mathcal{M} if $|\text{Del}_{\mathcal{M}}(P)|$ is homeomorphic to \mathcal{M} . A *Delaunay triangulation* of \mathcal{M} is a homeomorphism $H: |\text{Del}_{\mathcal{M}}(P)| \rightarrow \mathcal{M}$.

The *Voronoi cell* associated with $p \in P$ is given by

$$\mathcal{V}_{\mathcal{M}}(p) = \{x \in \mathcal{M} \mid d_{\mathcal{M}}(x, p) \leq d_{\mathcal{M}}(x, q) \text{ for all } q \in P\}.$$

More generally, a *Voronoi face* is the intersection of a set of Voronoi cells: given $\sigma = \{p_0, \dots, p_k\} \subset P$, we define the associated Voronoi face as

$$\mathcal{V}_{\mathcal{M}}(\sigma) = \bigcap_{i=0}^k \mathcal{V}_{\mathcal{M}}(p_i).$$

It follows that σ is a Delaunay simplex if and only if $\mathcal{V}_{\mathcal{M}}(\sigma) \neq \emptyset$. In this case, every point in $\mathcal{V}_{\mathcal{M}}(\sigma)$ is the centre of a Delaunay ball for σ . Thus every Voronoi face corresponds to a Delaunay

*INRIA, DataShape, Sophia-Antipolis, France Jean-Daniel.Boissonnat@inria.fr

†INRIA, DataShape Ramsay.Dyer@inria.fr

‡Indian Statistical Institute, ACM Unit, Kolkata, India arijitiitkgpster@gmail.com

§Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, P.O. Box 407, 9700 AK, Groningen, The Netherlands N.Martynchuk@rug.nl

simplex. The Voronoi cells give a decomposition of \mathcal{M} , denoted $\text{Vor}_{\mathcal{M}}(P)$, called the *Voronoi diagram*. The Delaunay complex of P is the nerve of the Voronoi diagram.

In the case of \mathbb{R}^m equipped with the standard Euclidean metric, Delaunay [Del34] showed that, if P is *generic*, then the natural inclusion $P \hookrightarrow \mathbb{R}^m$ induces a piecewise linear embedding $|\text{Del}_{\mathbb{R}^m}(P)| \hookrightarrow \mathbb{R}^m$ of the Delaunay complex of P into \mathbb{R}^m . This is called the Delaunay triangulation of P (it is a triangulation of the convex hull of P). The point set P is generic if there is no Delaunay ball with more than $m + 1$ points of P on its boundary. Point sets that are not generic are often dismissed in theoretical work, because an arbitrarily small perturbation of the points will almost surely yield a generic point set [Del34]. In the sense of the standard measure in the configuration space $\mathbb{R}^{m \times |P|}$, almost all point sets will yield a Delaunay triangulation.

A similar situation is known for certain standard non-Euclidean geometries, such as Laguerre geometry [Aur87], spaces equipped with a Bregman divergence [BNN10], or Riemannian manifolds of constant sectional curvature¹.

Leibon and Letscher [LL00] announced sampling density conditions which would ensure that the Delaunay complex defined by the intrinsic metric of an arbitrary compact Riemannian manifold triangulates the manifold. When triangulating submanifolds of dimension 3 and higher in Euclidean space using Delaunay techniques, it was subsequently discovered that near degenerate “sliver” simplices pose problems which cannot be escaped simply by increasing the sampling density. In particular, developing an example on a 3-manifold presented by Cheng et al. [CDR05], Boissonnat et al. [BGO09, Lemma 3.1] show that, using the metric of the ambient Euclidean space restricted to the submanifold, the resulting Delaunay complex (called the *restricted Delaunay complex*) need not triangulate the original submanifold, even with dense well-separated (no two points are too close) sampling.

Here we develop a similar example from the perspective of the intrinsic metric of the manifold. It can be argued that this is an easier way to visualize the problem, since we confine our viewpoint to a three dimensional space and perturb the metric, without referring to deformations into a fourth ambient dimension. This viewpoint also provides an explicit counterexample to the results announced by Leibon and Letscher [LL00]. We construct a fixed compact Riemannian manifold and demonstrate (Theorem 1) that for any sampling density there exists a point set that meets the sampling density and has good separation, but that does not admit a Delaunay triangulation, and furthermore, this property is retained when the point set is subjected to a sufficiently small perturbation. Thus, not only is there no sampling density that is sufficient to ensure the existence of a Delaunay triangulation on all compact Riemannian manifolds, but for any given fixed compact manifold, there may exist no sampling density that can guarantee a Delaunay triangulation. In particular, this latter property will complicate attempts to use Delaunay techniques in the construction and asymptotic analysis of optimal triangulations [Cla06, dL11].

Although we focus on the intrinsic metric, the same qualitative construction applies to the restricted ambient metric, demonstrating that the claim by Cairns [Cai61], that the Voronoi faces of the restricted Voronoi diagram are closed topological balls if the sampling is sufficiently dense, is also incorrect. In Section 3, we demonstrate a stable configuration in three dimensions in which the Voronoi face associated with a tetrahedron is two distinct points (0-balls), and the Voronoi face associated with an edge is two distinct topological 2-balls (see $\mathcal{V}_{\mathcal{M}}(uv)$ in Figure 3(a)). Since the restricted ambient metric is also qualitatively susceptible to this kind of construction, the Voronoi faces in the restricted metric are not necessarily themselves closed balls, regardless of the sampling density. This contradicts the natural interpretation of the statement that the Voronoi diagram “is a subdivision of [the manifold] into the closed cells of a complex” [Cai61]. However, our construction does not eliminate the possibility that the Voronoi faces are each collections of disjoint closed topological balls of the appropriate dimension. In this case, a generous interpretation of Cairns’ result would hold, but it is not clear to us from Cairns’ two-page note that the proof “is easy to complete”.

¹This is standard folklore. We are not aware of a published reference.

The main purpose of this note is to clear up these persistent misconceptions that have appeared in the published literature.

2 A qualitative argument

As we show in this section with a qualitative argument, the problem can be viewed as arising from the fact that in a manifold of dimension $m > 2$, the intersection of two metric spheres is not uniquely specified by m points. We demonstrate the issue in the context of Delaunay balls. The problem is developed quantitatively in terms of the Voronoi diagram in Section 3.

We work exclusively on a three dimensional domain. The problem is a local one, and we are not concerned with “boundary conditions”; we are looking at a coordinate patch on a densely sampled compact 3-manifold.

One core ingredient in Delaunay’s triangulation result [Del34] is that any triangle τ is the face of exactly two tetrahedra. This follows from the observation that a triangle has a unique circumcircle, and that any circumscribing sphere for τ must include this circle. The affine hull of τ cuts space into two components, and if $\tau \in \text{Del}_{\mathbb{R}^m}(P)$, then it will have an empty circumscribing sphere centred at a point c on the line through the circumcentre and orthogonal to $\text{aff}(\tau)$. The point c is contained on an interval on this line which contains all the empty spheres for τ . The endpoints of the interval are the circumcentres of the two tetrahedra that share τ as a face.

The argument hinges on the assumption that the points are in general position, and on the uniqueness of the circumcircle for τ . If there were a fourth vertex lying on that circumcircle, then there would be three tetrahedra that have τ as a face; the third tetrahedron being the degenerate tetrahedron lying in $\text{aff}(\tau)$, with all four vertices on the circumcircle of τ . But this configuration would violate the assumption of general position.

Now if we allow the metric to deviate from the Euclidean one, no matter how slightly, the guarantee of a well defined unique circumcircle for τ is lost. In particular, if three spheres S_1 , S_2 and S_3 all circumscribe τ , their pairwise intersections will be different in general, i.e.,

$$S_1 \cap S_3 \neq S_2 \cap S_3.$$

Although these intersections may be topological circles that are “very close” assuming the deviation of the metric from the Euclidean one is small enough, “very close” is not good enough when the only genericity assumption allows configurations that are arbitrarily bad.

The problem is illustrated in Figure 1, where $\tau = \{u, v, p\}$. Here, circumspheres S_1 and S_2 of the Delaunay tetrahedra $\{a, u, v, p\}$ and $\{u, v, p, b\}$ would contain any empty sphere S_3 that circumscribes τ if the metric were Euclidean, but any aberration in the metric may leave a part of S_3 exposed to the outside. This means that in principle another sample point w could lie on S_3 , while S_1 and S_2 remain empty. Thus there would be a third Delaunay tetrahedron, $\{u, v, p, w\}$ that shares τ as a face.

In dimension 2 this problem does not arise [Lei99, DZM08]. The essential difference between dimension 2 and the higher dimensions can be observed by examining the topological intersection properties of spheres. Specifically, two $(m-1)$ -spheres intersect transversely in an $(m-2)$ -sphere. For a non-Euclidean metric, even if this property holds for sufficiently small geodesic spheres, it is only when $m = 2$ that the sphere of intersection of the Delaunay spheres of two adjacent m -simplices is uniquely defined by the vertices of the shared $(m-1)$ -simplex.

3 An obstruction to intrinsic Delaunay triangulations

We now explicitly show how density assumptions alone cannot escape topological problems in the Delaunay complex. The configuration considered here may be recognised as similar to the

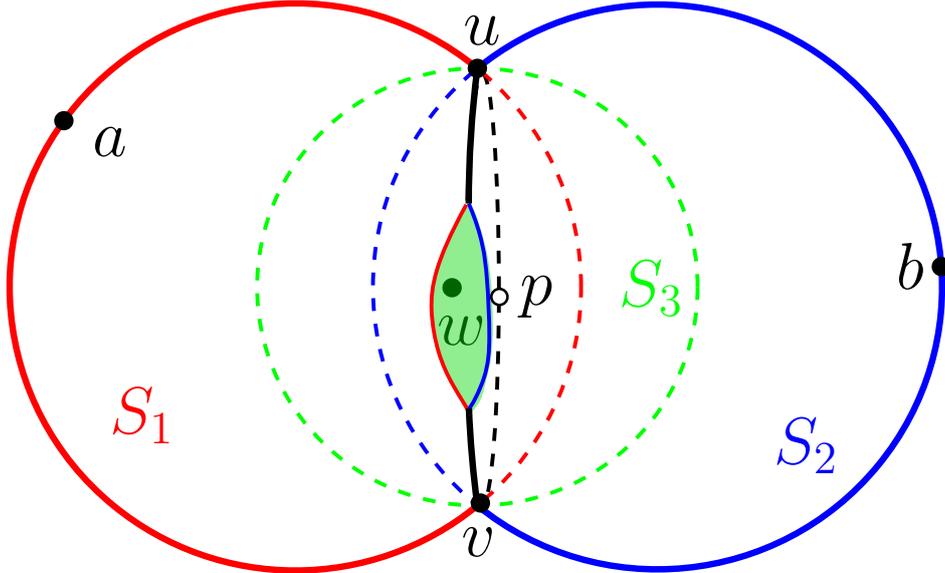


Figure 1: In three dimensions, three closed geodesic balls can all touch three points, u, v, p , on their boundary and yet no one of them is contained in the union of the other two. In Euclidean space $\{u, v, p\}$ defines a unique circumcircle (shown in black) that is the intersection of any two distinct spheres that circumscribe $\{u, v, p\}$, but if the metric deviates from Euclidean, this circle is no longer relevant; we can have three distinct spheres, S_1, S_2 , and S_3 all circumscribing $\{u, v, p\}$ and such that $S_1 \cap S_3 \neq S_2 \cap S_3$. This means that $\{u, v, p\}$ can be a face of three Delaunay tetrahedra, e.g., S_1, S_2 and S_3 could be circumscribing spheres for $\{a, u, v, p\}$, $\{u, v, p, b\}$, and $\{u, v, p, w\}$, respectively.

one described qualitatively in Section 2, but here we consider the Voronoi diagram rather than Delaunay balls. As before we work exclusively in a local coordinate patch on a densely sampled compact 3-manifold.

The idea is to consider four points in the xz -plane in Euclidean space, and show that with a mild perturbation of the metric, this tetrahedron exhibits exactly two distinct circumcentres. We construct a perturbed metric such that for any sampling density, a tetrahedron with two distinct Delaunay balls can exist. This means that the Delaunay complex, defined as the nerve of the Voronoi diagram, will not triangulate the manifold (this Delaunay tetrahedron will have two triangle faces that are not the face of any other Delaunay tetrahedron, as discussed below).

Delaunay's genericity condition [Del34] naturally extends to the setting where \mathcal{M} is a manifold, and Leibon and Letscher [LL00, p. 343] explicitly assume genericity in this sense:

Definition 1 If \mathcal{M} is a Riemannian m -manifold, the set $P \subset \mathcal{M}$, is *generic* if no subset of $m + 2$ points lies on the boundary of a geodesic ball.

Delaunay only imposed the constraint on empty balls, and he showed that any (finite or periodic) point set in Euclidean space can be made generic through an arbitrarily small affine perturbation. That a similar construction of a perturbation can be made for points on a compact Riemannian manifold has not been explicitly demonstrated. However, in light of the construction we now present, it seems that the question is moot when $m > 2$, because an arbitrarily small perturbation from degeneracy will not be sufficient to ensure that a Delaunay triangulation exists.

For a point x in a compact Riemannian manifold \mathcal{M} (without boundary), the injectivity radius at x is the supremum of radii r such that a geodesic ball with radius r , i.e., $B_{\mathcal{M}}(x, r) = \{y \in \mathcal{M} \mid d_{\mathcal{M}}(x, y) < r\}$, has the property that each $y \in B_{\mathcal{M}}(x, r)$ is connected to x by a unique minimising geodesic. The *injectivity radius* of \mathcal{M} , denoted $\text{inj}(\mathcal{M})$, is the infimum of

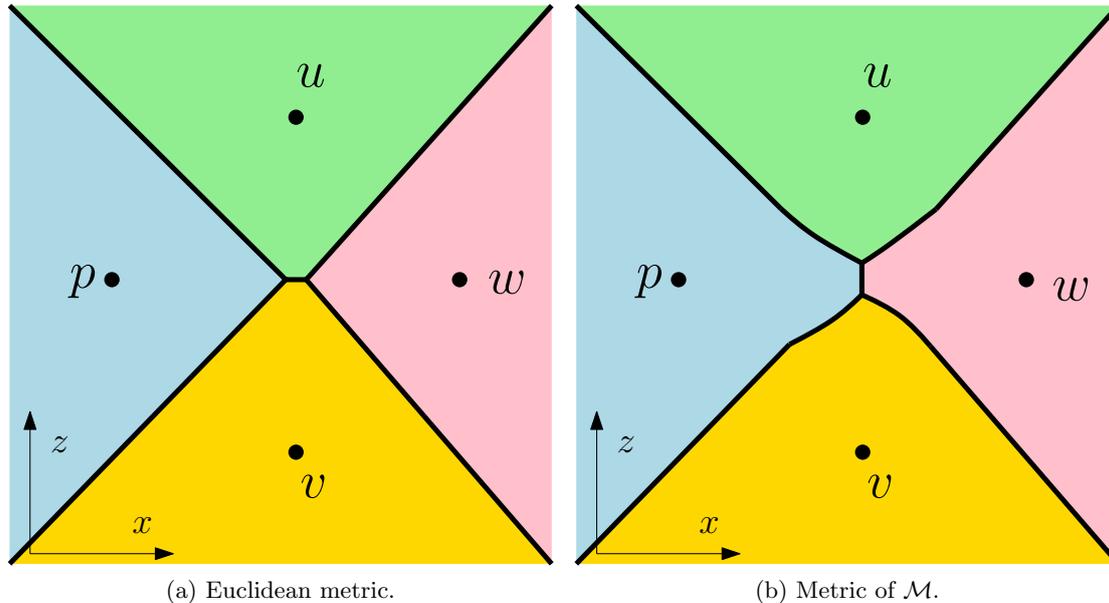


Figure 2: A qualitative schematic of a vertical slice (the $y = 0$ plane) of the Voronoi diagram, seen from the negative y axis (the axes are indicated in the lower left of the figures). In this plane the Euclidean diagram (a) is qualitatively different from the diagram (b) defined by the metric $d_{\mathcal{M}}$ of \mathcal{M} .

the injectivity radii at the points of \mathcal{M} . When \mathcal{M} is a compact manifold, $\text{inj}(\mathcal{M}) > 0$; see Chavel [Cha06, §III.2].

We say $P \subset \mathcal{M}$ is ϵ -dense if $d_{\mathcal{M}}(x, P) < \epsilon$ for any $x \in \mathcal{M}$. If $d_{\mathcal{M}}(p, q) \geq \tilde{\epsilon}$ for all $p, q \in P$, then P is $\tilde{\epsilon}$ -separated. The set P is an ϵ -net if it is ϵ -dense and ϵ -separated.

Theorem 1 For every integer $m > 2$ there exists a compact Riemannian m -manifold \mathcal{M} such that for any $\epsilon > 0$, with $\epsilon < \text{inj}(\mathcal{M})/2$, there exists a generic ϵ -net $P \subset \mathcal{M}$ such that $\text{Del}_{\mathcal{M}}(P)$ does not triangulate \mathcal{M} . Furthermore, this holds true for all perturbed point sets $P' \subset \mathcal{M}$ that are within a small positive Hausdorff distance of P .

Proof. Consider 3-dimensional Euclidean space, parameterised by x , y , and z , to be the parameter domain of a coordinate chart of a compact Riemannian 3-manifold \mathcal{M} . Place points u and v at $\pm\epsilon a$ on the z axis, and points w and p at $\pm\epsilon a(1 + \xi)$ on the x -axis. Here $\epsilon > 0$ is the sampling parameter that parameterises the scale of the example, $a = 1/\sqrt{2}$, and $\xi > 0$ is a small parameter that will be constrained by considerations below. We will fix a metric $d_{\mathcal{M}}$ on \mathcal{M} , and the value of ξ will depend on this metric, as well as the sampling parameter ϵ .

With this configuration, u and v will share a Voronoi face in the Euclidean metric which will extend indefinitely in the y -direction, but have a very short extent in the x -direction, assuming ξ is very small. Figure 2(a) depicts the slice of this Voronoi diagram defined by the xz -plane, as viewed from somewhere on the negative y -axis.

We now construct a metric on \mathcal{M} such that the geodesic distance between u and v will be greater than the distance between p and w , thus qualitatively changing the structure of the Voronoi diagram in the xz -plane, as shown in Figure 2(b). To do this we employ a C^∞ non-negative bump function $f: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ that is symmetric: $f(-y) = f(y)$, and satisfies $f'(y) \leq 0$ for $y > 0$ within our region of interest, and a strict inequality $f'(y) < 0$ for $y > 0$ in some open neighbourhood of the origin. Thus there is a unique maximum at the origin. As an explicit example, we may assume that the underlying set of \mathcal{M} is the torus $\mathbb{R}^3/(2\mathbb{Z})^3$, and that $f(y) = A(1 + \cos(\pi y))$, where the amplitude is constrained by $A \leq 3/8$, as described below. We

define the metric tensor of \mathcal{M} in our coordinate system by

$$g(q) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 + f(y(q)) \end{pmatrix},$$

where $y(q)$ denotes the y -coordinate of the point q .

By integrating along the z -axis we obtain an upper bound on the geodesic distance between v and u :

$$d_{\mathcal{M}}(v, u) \leq 2 \int_0^{a\epsilon} (1 + f(0))^{1/2} dz = 2a\epsilon(1 + f(0))^{1/2}. \quad (1)$$

In fact, if the right hand side of (1) is less than the injectivity radius of \mathcal{M} , then equality with $d_{\mathcal{M}}(v, u)$ is attained. Indeed, in this case the curve represented by the z -axis is the minimising geodesic between v and u . This follows from the symmetry of the metric: the reflection through the xz -plane $(x, y, z) \mapsto (x, -y, z)$ is an isometric involution that leaves points in the xz -plane fixed. Since an isometry maps geodesics to geodesics, and the minimising geodesic between v and u is unique, it must lie in the xz -plane, and, again by symmetry, the minimising curve is the straight line in the parameter domain.

We require that all four points lie within a ball centred at the origin, and with diameter less than $\text{inj}(\mathcal{M})$. The final constraint that we impose on f is that $f(0) \leq 3/4$. Another reason for this constraint is described below, but in particular it implies that $a(1 + f(0))^{1/2} < 1$. Then we see that, provided that $\epsilon < \text{inj}(\mathcal{M})/2$,

$$d_{\mathcal{M}}(v, u) = 2a\epsilon(1 + f(0))^{1/2} < \text{inj}(\mathcal{M}). \quad (2)$$

We also require that $d_{\mathcal{M}}(p, w) < d_{\mathcal{M}}(u, v)$, so that we have the qualitative Voronoi diagram in the xz -plane depicted by Figure 2(b). Observing that

$$d_{\mathcal{M}}(p, w) = d_{\mathbb{R}^3}(p, w) = 2a\epsilon(1 + \xi),$$

and equating with (2), we see that $d_{\mathcal{M}}(p, w) = d_{\mathcal{M}}(u, v)$ if ξ is equal to the critical value defined by

$$\xi_0 = (1 + f(0))^{1/2} - 1,$$

and so we require $\xi < \xi_0$.

The goal is to show that the tetrahedron $\sigma = \{u, v, w, p\}$ can be a Delaunay tetrahedron admitting two distinct Delaunay balls in an ϵ -net P . To that end, we require that it has a circumcentre c on the positive y -axis, at a distance less than ϵ from the vertices ($-c$ will also be a circumcentre). Let the y -coordinate of c be $b\epsilon$. We argue that for any $b > 0$, there is a $\xi < \xi_0$ such that $c = (0, b\epsilon, 0)$ and $-c$ are circumcentres of σ .

The geodesic distance between u and a given c is no greater than the geodesic length of the straight line between them in the parameter domain:

$$\begin{aligned} d_{\mathcal{M}}(u, c) &\leq \int_0^{b\epsilon} \left((1 + f(y)) \frac{a^2}{b^2} + 1 \right)^{\frac{1}{2}} dy \\ &< \int_0^{b\epsilon} \left((1 + f(0)) \frac{a^2}{b^2} + 1 \right)^{\frac{1}{2}} dy = b\epsilon \left((1 + \xi_0)^2 \frac{a^2}{b^2} + 1 \right)^{\frac{1}{2}}. \end{aligned} \quad (3)$$

Also, the geodesic distance between u and c is greater than the Euclidean distance:

$$d_{\mathcal{M}}(u, c) > d_{\mathbb{R}^3}(u, c) = b\epsilon \left(\frac{a^2}{b^2} + 1 \right)^{\frac{1}{2}}.$$

It follows then that for any $b > 0$, there exists a positive $\tilde{\xi}(b) < \xi_0$ such that

$$d_{\mathcal{M}}(u, c) = b\epsilon \left((1 + \tilde{\xi}(b))^2 \frac{a^2}{b^2} + 1 \right)^{\frac{1}{2}}. \quad (4)$$

Since $d_{\mathcal{M}}(u, c)$ varies smoothly with b , $\tilde{\xi}$ must be a smooth function.

The symmetry of the metric implies that the unique minimizing geodesic between p and c must lie in the xy -plane, where the metric coincides with the Euclidean one. It follows that, for $p = (-(1 + \xi)a\epsilon, 0, 0)$, we have

$$d_{\mathcal{M}}(p, c) = d_{\mathbb{R}^3}(p, c) = b\epsilon \left((1 + \xi)^2 \frac{a^2}{b^2} + 1 \right)^{\frac{1}{2}}. \quad (5)$$

Thus when $\tilde{\xi}(b) = \xi$ we have $d_{\mathcal{M}}(u, c) = d_{\mathcal{M}}(p, c)$, and it follows (by symmetry) that c is a circumcentre for σ .

We need to keep b small enough that $d_{\mathcal{M}}(u, c) < \epsilon$, so that σ can be a Delaunay tetrahedron in an ϵ -net. It is sufficient to demand that $b \leq b_{\max} = a/2$; then the fact that $d_{\mathcal{M}}(u, c) < \epsilon$ follows from the constraint $f(0) \leq 3/4$ and the definitions $a^2 = 1/2$, and $b^2 \leq b_{\max}^2 = 1/8$:

$$d_{\mathcal{M}}(u, c) < b\epsilon \left((1 + f(0)) \frac{a^2}{b^2} + 1 \right)^{\frac{1}{2}} \leq \epsilon.$$

Also, it follows from the construction that the vertices of σ meet the separation criterion of an ϵ -net: geodesic distances between the points are bounded below by the Euclidean distances in the parameter domain, e.g., $d_{\mathcal{M}}(u, w) > \|u - w\| = a\epsilon(1 + (1 + \xi)^2)^{1/2} > \epsilon$. Thus, σ can be realised as a Delaunay simplex with exactly two distinct Delaunay balls in an ϵ -net.

The configuration can be realised as part of an ϵ -net so that the two circumcentres of σ remain Voronoi vertices. Furthermore, this is not a degenerate configuration in the sense of Definition 1, because the requirements of an ϵ -net are met without placing any further vertices on the Delaunay spheres associated to σ ; i.e., σ is not the proper face of any Delaunay simplex.

When we deviate from a Euclidean metric we can encounter degeneracies different from those described by Definition 1. For example, if we were to choose $\xi = \xi_0$ in our example, then σ would have a unique circumcentre (at the origin), and it could be a Delaunay simplex that is not the proper face of any other. But an arbitrarily small perturbation, sending u or v towards the origin, for example, would mean that σ has no circumcentre at all.

However, at least for some positive $b \leq b_{\max}$, the presented example is not degenerate in any strict sense: there is some number $\rho > 0$ such that the vertices of σ can be each independently displaced by a distance ρ without disturbing the combinatorial structure of the Voronoi diagram. In other words, this bad configuration is represented by a set of positive measure in the configuration space.

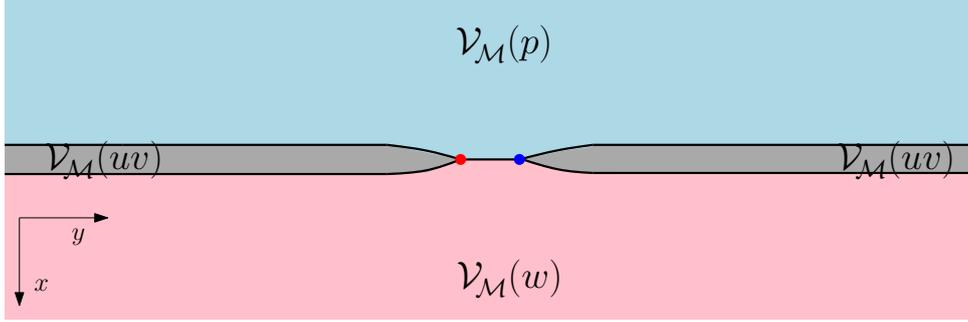
In order to demonstrate this it is sufficient to show that the circumcentres of σ depend continuously on the position of the vertices, at least in a neighbourhood of $s_0 = (p, v, w, u)$. This can be shown by demonstrating that the origin is a regular value of the map

$$h: \mathbb{R}^3 \rightarrow \mathbb{R}^3 \quad (6)$$

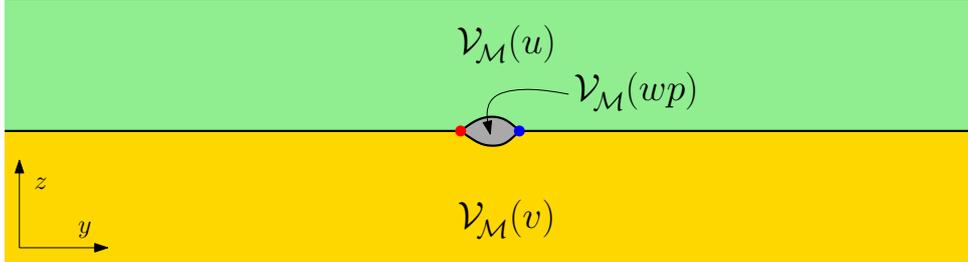
$$q \mapsto (d_{\mathcal{M}}(q, u) - d_{\mathcal{M}}(q, p), d_{\mathcal{M}}(q, v) - d_{\mathcal{M}}(q, p), d_{\mathcal{M}}(q, w) - d_{\mathcal{M}}(q, p)).$$

Indeed, if this is the case, then the implicit function theorem applied to the function $F: \mathbb{R}^{12} \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$ defined by (6), such that $F(s_0, q) = h(q)$, ensures that for all vertex positions s in a neighbourhood of s_0 , there will be circumcentres (solutions to $F(s, q) = 0$) close to c and $-c$.

In order to estimate the Jacobian determinant of h at $q = c$, let $H(q) = H(x, y, z) = d_{\mathcal{M}}(q, u) - d_{\mathcal{M}}(q, p)$ be the first component of h , and observe that, by symmetry, $H(x, y, -z) =$



(a) xy -plane seen from positive z -axis



(b) yz -plane seen from positive x -axis

Figure 3: Looking at cross-sections; the positive y -direction is to the right (the axes are indicated in the lower left of the figures). The tetrahedron $\sigma = \{p, u, v, w\}$, admits exactly two small circumballs with distinct centres (the red and blue points).

$d_{\mathcal{M}}(q, v) - d_{\mathcal{M}}(q, p)$ is the second component. Let $W(q) = d_{\mathcal{M}}(q, w) - d_{\mathcal{M}}(q, p)$ be the third component. We observe that $W(q) = 0$ if $x(q) = 0$. Therefore, denoting partial derivatives with subscripts, we have $W_y = W_z = 0$ at $q = c$. Also, we observe directly that W_x is negative. Let $W_x(c) = t < 0$. Therefore the Jacobian matrix has the form

$$Dh(c) = \begin{pmatrix} H_x & H_y & H_z \\ H_x & H_y & -H_z \\ t & 0 & 0 \end{pmatrix},$$

and to show that its determinant is nonzero, we need to verify that H_y and H_z are nonzero at c .

By inspection $H_z(c) < 0$ since c is a critical point of $d_{\mathcal{M}}(q, p)$ with respect to variation in the z direction, and $d_{\mathcal{M}}(q, u)$ is strictly decreasing at c as q moves up a vertical line. For $H_y(c)$, we exploit the fact that we have explicit expressions, (4) and (5), that describe $d_{\mathcal{M}}(u, c)$ and $d_{\mathcal{M}}(p, c)$. When differentiating with respect to b at the point where $\tilde{\xi}(b) = \xi$, we find

$$H_y(c) = \partial_b d_{\mathcal{M}}(c, u) - \partial_b d_{\mathcal{M}}(c, p) = \frac{a^2(1 + \xi)\tilde{\xi}'\epsilon}{((1 + \xi)^2 a^2 + b^2)^{\frac{1}{2}}}. \quad (7)$$

Equation (3) shows that for any $b > 0$, we have $\tilde{\xi}(b) < \xi_0 = \tilde{\xi}(0)$, and in particular, this is true for $b = b_{\max}$. So it follows from the mean value theorem that there is a $b \in (0, b_{\max}]$ such that $\tilde{\xi}'(b) < 0$, and (7) implies that when ξ is chosen so that this b represents a circumcentre, we have $H_y(c) < 0$. Thus this configuration is stable with respect to small perturbations of the vertices.

This constructed Delaunay tetrahedron σ with two distinct circumcentres implies a defect in the Voronoi diagram that prevents $\text{Del}_{\mathcal{M}}(P)$ from triangulating \mathcal{M} . A schematic view of the Voronoi diagram associated with σ is depicted in Figure 3. The circumcentres of σ are the Voronoi vertices shown in the diagram. A Voronoi edge corresponds to a Delaunay triangle, and the cofaces of the triangle are the Delaunay tetrahedra that correspond to the Voronoi vertices

at the ends of the Voronoi edge. So for triangle $t = \{u, p, w\}$, since the Voronoi vertex at each end of $\mathcal{V}_{\mathcal{M}}(t)$ represents the same simplex, σ , there will be only one coface of t (likewise for triangle $\{v, p, w\}$). Since our manifold \mathcal{M} has no boundary, $\text{Del}_{\mathcal{M}}(P)$ cannot triangulate \mathcal{M} .

Finally, observe that this kind of counter-example to Delaunay triangulations will also exist in higher dimensional manifolds. For example, the same basic construction when $m > 3$ can be obtained by starting with a regular simplex σ of dimension $(m - 2)$ such that its circumcentre is at the origin, and its circumradius is $\epsilon a(1 + \xi)$, with $a = 1/\sqrt{2}$, and such that

$$\text{aff}(\sigma) = \{x \in \mathbb{R}^m \mid x_{m-1} = x_m = 0\}.$$

Then place two additional vertices at $\pm a\epsilon$ on the x_{m-1} -axis to obtain an m -simplex. The metric g is then the diagonal matrix with $g_{ii}(q) = (1 + f(x_{m-1}(q)))$ if $i = m$ and $g_{ii}(q) = 1$ otherwise, and f defined exactly as before. \square

The qualitative construction of the previous section (Figure 1) is similar to the one we have considered here except that it conceptually employs a different perturbation function, whose support need not be symmetric with respect to the xz -plane, so that the simplex $\{u, v, p, w\}$ can be dual to a single stable Voronoi vertex. The Delaunay perspective on the example represented in Figure 3 would be similar to Figure 1, except that there would be two distinct green Delaunay balls for σ .

4 Discussion

We have shown that for constructing a Delaunay triangulation of an arbitrary compact Riemannian manifold, a sampling density requirement is not sufficient in general. One approach to this problem might be to constrain the kind of metrics that are considered. However, the example shown here exhibits a problem even with a mild deviation from homogeneity, so the admissible metrics will be severely limited. We expect this kind of problem to arise any time one attempts to construct an anisotropic triangulation using the empty sphere property in dimension $m > 2$, unless the $(m-1)$ -spheres circumscribing m points are constrained to all intersect in a unique $(m-2)$ -sphere defined by those points.

For example, in the anisotropic framework studied by Canas and Gortler [CG12] in the 2-dimensional setting, we can in three dimensions construct a problematic configuration as discussed in Section 2. In this setting, the “metric” is no longer symmetric, but the Delaunay triangulation is still described by the empty ball property; the balls are ellipsoids defined by the quadratic form associated with the centre. We can construct three ellipsoids that each have $\tau = \{u, v, p\}$ on their boundary, and such that the problem of three Delaunay tetrahedra that have τ as a face occurs in qualitatively the same way as in Figure 1. The Delaunay complex in this setting is still defined as the nerve of the Voronoi diagram, so a quantitative example analogous to that presented in Section 3 could also be constructed. In our obstruction the full dimensional Voronoi cells are connected (“orphan free”), as is always the case in the Riemannian metric. This is a *topological* obstruction to an embedding of this anisotropic Delaunay complex in \mathbb{R}^3 . (Even when the topology of the complex is correct, the natural piecewise linear map into Euclidean space defined by the coordinates of the vertices may not be an embedding, i.e., some simplices may be “inverted”; see [CG15, Figure 16, Section 7]. Our construction in this note does not shed light on this kind of geometric problem.)

An alternate approach to avoiding the problematic topological obstructions to triangulation is to constrain the kinds of point sets considered. Pursuing this direction, we developed the notion of *protected* Delaunay balls in Euclidean space [BDG13], which provides a parameter that allows us to quantify how close a point set is to degeneracy. We showed that point sets with a positive bound on this parameter can be produced, and used this to construct an algorithm which, when given a sufficiently dense point set P on an arbitrary compact Riemannian manifold

\mathcal{M} , will produce a Delaunay complex $\text{Del}(P')$ which triangulates \mathcal{M} and whose vertices P' are a perturbation of the original point set P ; see [BDG17].

Acknowledgements

We thank Gert Vegter for suggesting that we look for an obstruction that can exist at all scales. We also thank David Cohen-Steiner and Mathijs Wintraecken for illuminating discussions.

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). Partial support has also been provided by the Advanced Grant of the European Research Council GUDHI (Geometric Understanding in Higher Dimensions).

Arijit Ghosh is supported by Ramanujan Fellowship number SB/S2/RJN-064/2015. Part of this work was done when he was a researcher at the Max-Planck-Institute for Informatics, Germany supported by the IndoGerman Max Planck Center for Computer Science (IMPECS). Part of this work was also done while he was a visiting scientist at the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata, India.

References

- [Aur87] F. Aurenhammer. Power Diagrams: Properties, Algorithms and Applications. *SIAM J. Comput.*, 16(1):78–96, 1987. 2
- [BDG13] J.-D. Boissonnat, R. Dyer, and A. Ghosh. The stability of Delaunay triangulations. *Int. J. Comp. Geom. & Appl.*, 23(04n05):303–333, 2013. eprint: arXiv:1304.2947. 9
- [BDG17] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay triangulation of manifolds. *Foundations of Computational Mathematics*, 2017. (Preprint: arXiv:1311.0117). 10
- [BGO09] J.-D. Boissonnat, L. J. Guibas, and S. Y. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. *Discrete & Comp. Geom.*, 42:37–70, 2009. 2
- [BNN10] J.-D. Boissonnat, F. Nielsen, and R. Nock. Bregman Voronoi diagrams. *Discrete & Comp. Geom.*, 44(2):281–307, 2010. 2
- [Cai61] Stewart S. Cairns. A simple triangulation method for smooth manifolds. *Bulletin of the American Mathematical Society*, 67(4):389–390, 07 1961. 2
- [CDR05] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *Proceedings of SODA*, pages 1018–1027, 2005. 2
- [CG12] G. D. Canas and S. J. Gortler. Duals of orphan-free anisotropic Voronoi diagrams are embedded meshes. In *Proceedings of SoCG*, pages 219–228, New York, NY, USA, 2012. ACM. 9
- [CG15] Guillermo D. Canas and Steven J. Gortler. On the embeddability of Delaunay triangulations in anisotropic, normed, and Bregman spaces. *CoRR*, abs/1512.03589, 2015. 9
- [Cha06] I. Chavel. *Riemannian Geometry, A modern introduction*. Cambridge, 2nd edition, 2006. 5
- [Cla06] K. L. Clarkson. Building triangulations using epsilon-nets. In *Proceedings of STOC*, 2006. 2

- [Del34] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934. 2, 3, 4
- [dL11] David de Laat. Approximating manifolds by meshes: Asymptotic bounds in higher codimension. Master’s thesis, University of Groningen, 2011. 2
- [DZM08] 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. 3
- [Lei99] G. Leibon. *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization*. PhD thesis, UCSD, 1999. arXiv:math/0011016v1. 3
- [LL00] G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *Proceedings of SoCG*, pages 341–349, 2000. 2, 4