



HAL
open science

Discretized Riemannian Delaunay Triangulations

Mael Rouxel-Labbé, Mathijs Wintraecken, Jean-Daniel Boissonnat

► **To cite this version:**

Mael Rouxel-Labbé, Mathijs Wintraecken, Jean-Daniel Boissonnat. Discretized Riemannian Delaunay Triangulations. [Research Report] RR-9103, INRIA Sophia Antipolis - Méditerranée. 2017, pp.51. hal-01612924

HAL Id: hal-01612924

<https://inria.hal.science/hal-01612924>

Submitted on 8 Oct 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.



Discretized Riemannian Delaunay triangulations

M. Rouxel-Labbé^{1,2}, M. Wintraecken¹ & J.-D. Boissonat¹

¹INRIA Sophia-Antipolis, France

²GeometryFactory, France

**RESEARCH
REPORT**

N° 9103

Octobre 2017

Project-Teams Datashape



Discretized Riemannian Delaunay triangulations

M. Rouxel-Labbé^{1,2}, M. Wintraecken¹ & J.-D. Boissonnat¹

¹INRIA Sophia-Antipolis, France

²GeometryFactory, France

Project-Teams Datashape

Research Report n° 9103 — Octobre 2017 — 51 pages

Abstract: Anisotropic meshes are desirable for various applications, such as the numerical solving of partial differential equations and graphics. In this report, we introduce an algorithm to compute discrete approximations of Riemannian Voronoi diagrams on 2-manifolds. This is not straightforward because geodesics, shortest paths between points, and therefore distances cannot, in general, be computed exactly.

We give conditions that guarantee that our discrete Riemannian Voronoi diagram is combinatorially equivalent to the exact Riemannian Voronoi diagram. This allows us to build upon recent theoretical results on Riemannian Delaunay triangulations, and guarantee that the dual of our discrete Riemannian Voronoi diagram is an embedded triangulation using both approximate geodesics and straight edges.

Our implementation employs recent developments in the numerical computation of geodesic distances. We observe that, in practice, our discrete Voronoi Diagram is correct in a far wider range of settings than our theoretical bounds imply.

Both the theoretical guarantees on the approximation of the Voronoi diagram and the implementation are new and provides a step towards the practical application of Riemannian Delaunay triangulations.

Key-words: Anisotropic mesh generation; Delaunay triangulation; Voronoi diagram; Geodesic distance;

**RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Triangulations de Delaunay riemanniennes discrètes

Résumé : Les maillages anisotropes sont désirables pour de nombreuses applications, telles que la résolution numérique d'équations aux dérivées partielles ou la visualisation. Dans ce rapport, nous présentons un algorithme qui permet de calculer une approximation discrète d'un diagramme de Voronoi Riemannien sur une 2-variété. Il s'agit d'une tâche complexe car ce diagramme est basé sur la notion de courbe géodésique, qui ne peut en général pas être calculée de manière exacte.

Nous donnons dans ce rapport des conditions qui garantissent que notre diagramme de Voronoi Riemannien discret est combinatoirement équivalent au diagramme de Voronoi Riemannien exact. Ceci nous permet ensuite d'utiliser des résultats récents sur les triangulations de Delaunay Riemanniennes pour garantir le fait que le dual de notre diagramme de Voronoi Riemannien discret est une triangulation plongée, à la fois en utilisant des arêtes géodésiques et des arêtes droites.

Notre implémentation est basée sur de récentes avancées dans le calcul numérique des distances géodésiques. Nous observons en pratique que notre diagramme de Voronoi Riemannien discret est correct dans des conditions beaucoup moins contraignantes que ce que notre théorie implique.

Les garanties théoriques et l'approximation du diagramme de Voronoi sont nouvelles et sont une étape de plus vers une utilisation pratique des triangulations de Delaunay Riemanniennes.

Mots-clés : Génération de maillages anisotropes; Triangulation de Delaunay; Diagramme de Voronoi; Distance géodésique;

1 Introduction

Anisotropic simplicial meshes are triangulations whose elements are elongated along prescribed directions. Anisotropic meshes are known to be well suited when solving PDE's [24, 49, 42]. They can also significantly enhance the accuracy of a surface representation if the anisotropy of the mesh conforms to the curvature of the surface [31].

Many anisotropic mesh generation methods are based on Riemannian manifolds and create meshes whose elements adapt locally to the size and anisotropy prescribed by a metric field. Several authors have considered Voronoi diagrams using anisotropic distances to obtain triangulations adapted to an anisotropic metric field. These authors hoped to build on the well established concepts of the Euclidean Voronoi diagram and its dual structure, the Delaunay triangulation, and their many theoretical and practical results [2]. The Voronoi diagram of a set of points $\mathcal{P} = \{p_i\}$ in a domain Ω is a partition of Ω in a set of regions $\{V_d(p_i)\}$. $V_d(p_i)$ is called the *Voronoi cell* of the *seed* p_i and is defined by

$$V_d(p_i) = \{x \in \Omega \mid d(p_i, x) \leq d(p_j, x), \forall p_j \in \mathcal{P} \setminus p_i\}.$$

Du and Wang [25] defined the anisotropic Voronoi cell of a point p_i in a domain Ω endowed with a metric field g as

$$V_{DW,g}(p_i) = \{x \in \Omega \mid d_{g(x)}(p_i, x) \leq d_{g(x)}(p_j, x), \forall p_j \in \mathcal{P} \setminus p_i\}.$$

Canas and Gortler [14, 15] and Cheng et al. [22] proved that under sufficient sampling conditions, the dual of this Voronoi diagram is an embedded triangulation for planar and surface domains. Labelle and Shewchuk [36] introduced slightly different Voronoi cells:

$$V_{LS,g}(p_i) = \{x \in \Omega \mid d_{g(p_i)}(p_i, x) \leq d_{g(p_j)}(p_j, x), \forall p_j \in \mathcal{P} \setminus p_i\}.$$

While the diagram obtained is easier to study and trace since the bisectors are quadrics, the associated theory only covers the planar setting and its extension to higher dimensions faces inherent difficulties.

Both these diagrams aim to approximate the Riemannian Voronoi diagram (RVD) whose cells are defined by

$$V_g(p_i) = \{x \in \Omega \mid d_g(p_i, x) \leq d_g(p_j, x), \forall p_j \in \mathcal{P} \setminus p_i\}.$$

where $d_g(p, q)$ is the geodesic distance on the Riemannian manifold. Judged too expensive, this Voronoi diagram has not been the subject of any practical work. However, the RVD provides a more favorable framework to develop a theoretically sound method. Additionally, it is devoid of orphan cells, which are regions of the domain disconnected from their seed. These orphans, which usually prevent the dual from being an embedded triangulation, appear in the other anisotropic Voronoi diagrams previously mentioned.

Leibon studied the dual of the Riemannian Voronoi diagram in 2D [38] and gave conditions such that the Riemannian Delaunay triangulation is an embedded triangulation. Together with Letscher [37], he extended that study to Riemannian Delaunay triangulations in any dimension, but their analysis was shown to be incomplete by Boissonnat et al. [4]. Dyer et al. [27] improved the 2D and surface work of Leibon with better bounds.

1.1 Contribution

We approach the Riemannian Voronoi diagram and its dual with a focus on both practicability and theoretical robustness. We introduce an algorithm to compute a discretized Riemannian

Voronoi diagram that is generic with respect to the choice of the method used to compute geodesic distances. The construction of our diagram uses an anisotropic graph structure to greatly reduce the computational time.

By building upon the theoretical results of the Riemannian Delaunay triangulation, we show that, under sufficient conditions, the dual of the Riemannian Voronoi diagram can be embedded as a triangulation (mesh) of the point set, with edges drawn as geodesic arcs, or even as straight edges.

Finally, we present an experimental study of our method in the planar and surface cases.

2 Related work

Due to its wide array of practical applications, anisotropic mesh generation has received a lot of attention and several classes of methods have been proposed.

Delaunay-based methods

The generation of isotropic Delaunay mesh is usually achieved through an iterative refinement algorithm. The common choice to refine a simplex is to choose its circumcenter as Steiner point. Borouchaki et al. [12] introduced the anisotropic Delaunay kernel, iteratively generating anisotropic meshes with a modified Bowyer-Watson algorithm. A different approach to anisotropic Delaunay triangulations was proposed by Boissonnat et al. [10], called locally uniform anisotropic meshes. In their algorithm, the star of each vertex v is composed of simplices that are Delaunay for the metric at v . Each star is built independently and the stars are stitched together in the hope of creating an anisotropic mesh. However, two stars may be combinatorially incompatible, a configuration called an *inconsistency*. Boissonnat et al. proved that these inconsistencies can be resolved by inserting Steiner points, yielding an anisotropic triangulation. The algorithm works in any dimension, can handle complex geometries and provides guarantees on the quality of the simplices of the triangulation. However, the number of vertices required to achieve consistency in the star set can be very large for highly distorted metric fields, making the algorithm unattractive.

Energy-based methods

Optimal Delaunay triangulations (ODT) were introduced by Chen et al. [21, 20] as the triangulation that minimizes the interpolation error $\|u - \hat{u}\|$ of a function u and its piecewise interpolation \hat{u} over the mesh. First applied to isotropic mesh generation using $u(x) = x^2$ [1, 51], ODT was naturally extended to anisotropic metric fields derived from the Hessian of a convex function through the use of weighted Delaunay triangulations [39, 44]. However, as shown by Boissonnat et al. [3], a Riemannian metric field generally cannot always be expressed as the Hessian of a convex function. Chen et al. [19] attempted to remedy this restriction by locally approximating the metric field with local convex functions, but the algorithm may fail to converge or achieve the desired anisotropy. Fu et al. [30] tweaked this approach by including feedback from the neighboring vertices in the expression of the local energy, with better experimental results.

Many other methods have been considered, such as isometric and conformal embeddings [57] and repulsive vertex repositioning [50, 56] techniques, but they possess shortcomings in theoretical or practical aspects.

Geodesic Voronoi diagrams

Peyré and Cohen [46] used isotropic geodesic Voronoi diagrams on surfaces as a tool to achieve isotropic adaptive remeshing. Cao et al. [16] studied the digital Voronoi diagram, a discrete structure diagram obtained by flooding from the seed and proved that they obtain a geometrically and topologically correct dual triangulation. Their work is however limited to the 2D isotropic setting.

3 Background

In this section, we define the core concepts of this work such as nets and power protection as well as some Riemannian geometry.

3.1 Nets

The structures of interest will be built from a finite set of points $\mathcal{P} \subset \mathbb{R}^d$, which we call *samples* or *seeds*. If Ω is a bounded domain, \mathcal{P} is an ϵ -*sample set* for Ω if $d_{\mathbb{E}}(x, \mathcal{P}) < \epsilon$ for all $x \in \Omega$. The set \mathcal{P} is μ -separated if $d_{\mathbb{E}}(p, q) \geq \mu$ for all $p, q \in \mathcal{P}$. We assume that the separation parameter μ is proportional to ϵ , thus there exists a positive λ , with $\lambda \leq 1$, such that $\mu = \lambda\epsilon$. If \mathcal{P} is an ϵ -sample set that is μ -separated, we say that \mathcal{P} is an (ϵ, μ) -*net*.

3.2 Riemannian metric

A metric in \mathbb{R}^d is defined by a symmetric positive definite quadratic form represented by a $d \times d$ matrix G . The distance between two points p and q in Ω , as measured using the metric G , is defined by

$$d_G(x, y) = \sqrt{(x - y)^t G (x - y)}.$$

Given the symmetric positive definite matrix G , we denote by F any matrix such that $\det(F) > 0$ and $F^t F = G$. F is called a *square root* of G . Although not unique, we shall assume that F is fixed. Note that

$$d_G(x, y) = \sqrt{(x - y)^t F^t F (x - y)} = \|F(x - y)\|,$$

where $\|\cdot\|$ stands for the Euclidean norm.

A Riemannian metric field g defined on a domain $\Omega \subset \mathbb{R}^d$ associates a metric $g(p) = G_p = F_p^t F_p$ to any point $p \in \Omega$. When the map $g : p \mapsto G_p$ is constant, we say that the metric field is *uniform*. The metric fields considered will be assumed to be ζ -bilipschitz continuous, meaning that for p, q in Ω ,

$$\frac{1}{\zeta} \|p - q\| \leq \|G_p - G_q\| \leq \zeta \|p - q\|.$$

where $\|\cdot\|$ is the matrix norm operator, *i.e.* for a square matrix A , $\|A\| = \sup_x \frac{\|Ax\|}{\|x\|}$.

Labelle and Shewchuk [36] introduced the concept of *distortion* between two points p and q of Ω :

$$\psi(p, q) = \psi(G_p, G_q) = \max \{ \|F_p F_q^{-1}\|, \|F_q F_p^{-1}\| \},$$

Note that $\psi(G_p, G_q) \geq 1$ with $\psi(G_p, G_q) = 1$ when $G_p = G_q$. A fundamental property of the distortion is to relate the two distances d_{G_p} and d_{G_q} . Specifically, for any pair x, y of points, we have

$$\frac{1}{\psi(p, q)} d_{G_p}(x, y) \leq d_{G_q}(x, y) \leq \psi(p, q) d_{G_p}(x, y).$$

The following lemma shows how the Lipschitz condition imposed on our metric fields allows us to compute the size of a neighborhood, given a distortion bound.

Lemma 3.1 (Distortion and neighborhood size for bilipschitz metric field) *Let g be a ζ -bilipschitz Riemannian metric field. Let p_0 be a point on Ω and $\psi_0 \geq 1$ a distortion bound. If $\psi(g(p), g(p_0)) \leq \psi_0$, then*

$$\|p - p_0\| \leq \zeta(1 + \psi_0^2) \|g(p_0)\|.$$

Proof We have

$$\begin{aligned} \|g(p) - g(p_0)\| &= \left\| g(p)^{\frac{1}{2}} g(p_0)^{-\frac{1}{2}} g(p_0) g(p_0)^{-\frac{1}{2}} g(p)^{\frac{1}{2}} - g(p_0) \right\| \\ &\leq \|g(p_0)\| + \left\| g(p)^{\frac{1}{2}} g(p_0)^{-\frac{1}{2}} g(p_0) g(p_0)^{-\frac{1}{2}} g(p)^{\frac{1}{2}} \right\| \\ &\leq \|g(p_0)\| + \left\| g(p)^{\frac{1}{2}} g(p_0)^{-\frac{1}{2}} \right\| \|g(p_0)\| \left\| g(p_0)^{-\frac{1}{2}} g(p)^{\frac{1}{2}} \right\| \\ &\leq \|g(p_0)\| (1 + \psi_0^2). \end{aligned}$$

Since the metric field is ζ -bilipschitz, we have

$$\frac{1}{\zeta} \|p - p_0\| \leq \|g(p) - g(p_0)\|$$

and we can deduce a condition on the size of the neighborhood.

$$\|p - p_0\| \leq \zeta(1 + \psi_0^2) \|g(p_0)\|.$$

□

3.3 Geodesics

In the following, we assume that our domain has no boundaries and is geodesically complete in order to avoid punctures. A path between two points x, y of Ω is a function $\gamma : [0, 1] \rightarrow \Omega$ such that $\gamma(0) = x$ and $\gamma(1) = y$. If Ω is endowed with the Riemannian metric field g , the Riemannian length of a path γ is defined as

$$l_g(\gamma) = \int_0^1 \sqrt{\gamma'(t)^t g(\gamma(t)) \gamma'(t)} dt.$$

Let $\Gamma(x, y)$ be the set of paths from x to y on Ω : $\Gamma(x, y) = \{\gamma : [0, 1] \rightarrow \Omega \mid \gamma(0) = x \text{ and } \gamma(1) = y\}$. The geodesic distance between two points is given by

$$d_g(x, y) = \min_{\gamma \in \Gamma(x, y)} l_g(\gamma).$$

We say that $\mathcal{P} = \{p_i\}$ is an (ϵ, μ) -net with respect to the metric field g if, for all $x \in \Omega$, there exists $p_i \in \mathcal{P}$ such that $d_g(p, x) < \epsilon$ and for all $p_i, p_j \in \mathcal{P}$, $d_g(p_i, p_j) \geq \mu$.

The following lemma shows how the distortion, initially introduced to relate distances with respect to a fixed metric, can be extended to geodesic distances on a Riemannian manifold.

Lemma 3.2 (Relating geodesic distances and the distortion) *Let $U \subset \Omega$ be open and g and g' two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0)$, $\psi(g(p), g'(p)) \leq \psi_0$. Then, for all $x, y \in U$,*

$$\frac{1}{\psi_0} d_g(x, y) \leq d_{g'}(x, y) \leq \psi_0 d_g(x, y),$$

where d_g and $d_{g'}$ indicate the geodesic distances with respect to g and g' respectively.

Proof Recall that for $p \in B_g(p_0, r_0)$ and, for any pair (x, y) of points, we have

$$\frac{1}{\psi_0} d_{g(p)}(x, y) \leq d_{g'(p)}(x, y) \leq \psi_0 d_{g(p)}(x, y).$$

Therefore, for any curve $\gamma(t)$ in U , we have that

$$\frac{1}{\psi_0} \int \sqrt{\langle \dot{\gamma}, \dot{\gamma} \rangle_{g(\gamma(t))}} dt \leq \int \sqrt{\langle \dot{\gamma}, \dot{\gamma} \rangle_{g'(\gamma(t))}} dt \leq \psi_0 \int \sqrt{\langle \dot{\gamma}, \dot{\gamma} \rangle_{g(\gamma(t))}} dt.$$

Considering the infimum over all paths γ that begin at x and end at y , we find that

$$\frac{1}{\psi_0} d_g(x, y) \leq d_{g'}(x, y) \leq \psi_0 d_g(x, y),$$

which relates the two geodesic distances. \square

3.3.1 Riemannian Voronoi diagrams

A cell of the Riemannian Voronoi diagram (RVD) with respect to a metric field g on Ω is defined by

$$V_g(p_i) = \{x \in \Omega \mid d_g(p_i, x) \leq d_g(p_j, x), \forall p_j \in \mathcal{P} \setminus p_i\}.$$

A *Voronoi face* is the intersection of a set of Voronoi cells. If the dimension of the face is 0 or 1, we respectively use the terms *Voronoi vertex* and *Voronoi edge*. The *bisector* of two seeds p and q is the set $\mathcal{BS}(p, q) = \{x \in \Omega, d(p, x) = d(q, x)\}$.

The *nerve* of a finite covering C is $\text{Nrv}(C) = \{D \subseteq C \mid \cap D \neq \emptyset\}$. A geometric realization of the nerve is a (geometric) simplicial complex \mathcal{K} and a bijection φ between $\text{Nrv}(C)$ and \mathcal{K} that preserves the inclusion, that is τ and σ are simplices of \mathcal{K} such that $\tau \subset \sigma$ if, and only if, there exists $D_\sigma = \varphi(\sigma)$ and $D_\tau = \varphi(\tau)$ in $\text{Nrv}(C)$ such that $D_\sigma \subset D_\tau$. In the following, we identify \mathcal{K} and the realization of the nerve. The nerve of the RVD is $\text{Nrv}(\text{RVD}) = \{\{V(p_i)\} \mid \cap_i V(p_i) \neq \emptyset, p_i \in \mathcal{P}\}$. A geometric realization of $\text{Nrv}(\text{RVD})$ whose vertices are \mathcal{P} is called a *dual* of the RVD. A homeomorphism is a continuous bijection whose inverse is also continuous. If there exists an abstract homeomorphism between the nerve $\text{Nrv}(C)$ and the domain, then we say that the realization of the nerve is a *triangulation*. If there exists an embedding (which provides a homeomorphism) of the dual whose edges are geodesic arcs, such an embedding is called the *Riemannian Delaunay triangulation* (RDT) of \mathcal{P} .

3.4 Protection and power-protection

The concept of protection and power-protection of point sets and simplices were introduced by Boissonnat et al. [7] to provide lower bounds on the thickness and quality of simplices.

Let \mathcal{P} be a finite point set in \mathbb{R}^d . Let σ be a Delaunay simplex whose vertices belong to \mathcal{P} , whose (empty) Delaunay ball is noted $B(c, r)$, where c is the circumcenter and r is the circumradius. For $0 \leq \delta \leq r$, we associate to $B(c, r)$ the dilated balls $B^{+\delta} = B(c, r + \delta)$ and $B_p^{+\delta} = B(c, \sqrt{r^2 + \delta^2})$. We say that σ is δ -protected (resp. δ -power protected) if $B^{+\delta}$ (resp. $B_p^{+\delta}$) is empty of any vertex in $\mathcal{P} \setminus \text{Vert}(\sigma)$. If all the simplices of $\text{Del}(\mathcal{P})$ are δ -protected (resp. δ -power protected) then \mathcal{P} is δ -protected (resp. δ -power protected).

4 Properties of power protected nets in Euclidean space

We detail some properties induced on the Voronoi diagram when the point set is a power protected net.

4.1 Voronoi angle bounds

The (ϵ, μ) -net assumption allows us to deduce lower and upper bounds on the Voronoi angles when the domain is \mathbb{R}^2 endowed with the Euclidean metric field.

Lemma 4.1 *Let $\Omega = \mathbb{R}^2$ and \mathcal{P} be an (ϵ, μ) -net with respect to the Euclidean distance on Ω . Let $V(p)$ be the Voronoi cell of $p \in \mathcal{P}$ and θ the angle at a corner of $V(p)$. Then*

$$2 \arcsin\left(\frac{\mu}{2\epsilon}\right) \leq \theta \leq \pi - \arcsin\left(\frac{\mu}{2\epsilon}\right).$$




Figure 1: Illustration of the notations used in Lemma 4.1

Proof Figure 1 illustrates the different notations used in this proof.

Lower bound.

Let m_{pq} and m_{pr} be the projections of the seed p on respectively the bisectors $BS(p, q)$ and $BS(p, r)$. Since \mathcal{P} is an (ϵ, μ) -net, we have $l_q = |p - m_{pq}| \geq \mu/2$, $l_r = |p - m_{pr}| \geq \mu/2$ and $L = |p - c| \leq \epsilon$. Thus

$$\theta = \arcsin\left(\frac{L}{l_q}\right) + \arcsin\left(\frac{L}{l_r}\right) \geq 2 \arcsin\left(\frac{\mu}{2\epsilon}\right).$$

Note that the argument of the arcsin is within the definition interval since $0 < \mu < 2\epsilon$.

Upper bound.

To get an upper bound on θ , we compute a lower bound on the angle α at p , noting that $\theta = \pi - \alpha$. Let $l_{qr} = |q - r|$, and $R = |c - r|$. By the law of sines, we have

$$\frac{l_{qr}}{\sin(\alpha)} = 2R.$$

Since \mathcal{P} is an (ϵ, μ) -net, we have $l_{qr} \geq \mu$ and $R \leq \epsilon$. Finally,

$$\alpha \geq \arcsin\left(\frac{\mu}{2\epsilon}\right) \implies \theta \leq \pi - \arcsin\left(\frac{\mu}{2\epsilon}\right).$$

□

4.2 Separation of the Voronoi vertices

We show that there exists a lower bound on the Euclidean distance between the Voronoi vertices of δ -power protected ϵ -samples. The main result is presented in Lemma 4.6 with Lemmas 4.2, 4.3 and 4.5 providing intermediary steps.

Lemma 4.2 *Let $B = B(c, R)$ and $B' = B(c', R')$ be two d -balls whose bounding spheres ∂B and $\partial B'$ intersect, and let H be the bisecting hyperplane of B and B' , i.e. the hyperplane that contains the $(d-2)$ -sphere $S = \partial B \cap \partial B'$. Let θ be the angle of the cone (c, S) . Writing $\rho = \frac{R'}{R}$ and $\|c - c'\| = \lambda R$, we have*

$$\cos(\theta) = \frac{1 + \lambda^2 - \rho^2}{2\lambda}. \quad (1)$$

If $R \geq R'$, we get $\cos(\theta) \geq \frac{\lambda}{2}$.

Proof Let $q \in S$, applying the cosine rule to the triangle $\triangle cc'q$ gives

$$\lambda^2 R^2 + R^2 - 2\lambda R^2 \cos \theta = R'^2 \quad (2)$$

which proves Equation (1). If $R \geq R'$, $\rho \leq 1$, and $\cos(\theta) \geq \frac{\lambda}{2}$ immediately follows from Equation (1). □

Lemma 4.3 *Let $B = B(c, R)$ and $B' = B(c', R')$ be two d -balls whose bounding spheres ∂B and $\partial B'$ intersect, and let $\tilde{\theta}$ be the angle of the cone (c, \tilde{S}) where $\tilde{S} = \partial B \cap \partial B'^{+\delta}$. Writing $\|c - c'\| = \lambda R$, we have*

$$\cos \tilde{\theta} = \cos \theta - \frac{\delta^2}{2R^2\lambda}.$$

Proof Let $\tilde{q} \in \tilde{S}$, applying the cosine rule to the triangle $\triangle cc'\tilde{q}$ gives

$$\lambda^2 R^2 + R^2 - 2\lambda R^2 \cos \tilde{\theta} = R'^2 + \delta^2.$$

Subtracting Equation (2) from the previous equality yields $\delta^2 = 2\lambda R^2(\cos \theta - \cos \tilde{\theta})$, which proves the lemma. □

Definition 4.4 (Height) *The height of a point p in a simplex $\sigma \ni p$ is defined by*

$$h(p, \sigma) = d(p, \text{aff}(\sigma \setminus p)).$$

Lemma 4.5 *Let $\sigma = p * \tau$ and $\sigma' = p' * \tau$ be two Delaunay simplices sharing a common facet τ . Here $*$ denotes the join operator. Let $B(\sigma) = B(c, R)$ and $B(\sigma') = B(c', R')$ be the circumscribing balls of σ and σ' respectively. Then σ' is δ -power protected with respect to p , that is $p \notin B(\sigma')^{+\delta}$ if and only if $\|c - c'\| \geq \frac{\delta^2}{2h(p, \sigma)}$.*




Figure 2: Illustration of the construction used Lemmas 4.2 and 4.3.

Proof ∂B and $\partial B'^{+\delta}$ intersect in a $(d-2)$ -sphere \tilde{S} which is contained in a hyperplane \tilde{H} parallel to the hyperplane $H = \text{aff}(\tau)$. For any $\tilde{q} \in \tilde{S}$ we have

$$d(\tilde{H}, H) = d(\tilde{q}, H) = R(\cos(\theta) - \cos(\tilde{\theta})) = \frac{\delta^2}{2\|c - c'\|},$$

where the last equality follows from Lemma 4.3 and $d(\tilde{H}, H)$ denotes the distance between the two parallel hyperplanes. See Figure 2 for a sketch. Since $p \in \partial B$, p belongs to $B(\sigma)^{+\delta}$ if and only if p lies in the strip bounded by H and \tilde{H} , which is equivalent to

$$d(p, H) = h(p, \sigma) < \frac{\delta^2}{2\|c - c'\|}.$$

The result now follows. \square

We can now build a bound on the separation of the Voronoi vertices that is independent of the simplices.

Lemma 4.6 (Voronoi vertex separation bound) *For any two adjacent Voronoi vertices c and c' of the RVD, we have*

$$\|c - c'\| \geq \frac{\delta^2}{4\epsilon}.$$

Proof For any simplex σ , we have $h(p, \sigma) \leq 2R_\sigma$ for all $p \in \sigma$, where R_σ denotes the radius of the circumsphere of σ . For any σ in the triangulation of an ϵ -net, we have $R_\sigma \leq \epsilon$. Thus $h(p, \sigma) \leq 2\epsilon$, and Lemma 4.5 yields $\|c - c'\| \geq \frac{\delta^2}{4\epsilon}$. \square

5 Stability of the power protected net with respect to metric perturbations

We show that a δ -power protected (ϵ, μ) -net stays a protected net, under some conditions, when the metric field is perturbed or approximated.

5.1 Stability of the power protected net property with respect to metric transformation

The square root of a metric allows us to switch between the embedding space endowed with an arbitrary uniform metric and the metric space endowed with an Euclidean metric. The power protected net property is shown to be conserved as the point set is transformed between these spaces.

Lemma 5.1 (Stability under metric transformation) *Let \mathcal{P} be a δ -power protected (ϵ, μ) -net with in Ω . Let $g = g_0$ be a uniform metric field and F_0 is a square root of g_0 . If \mathcal{P} is a δ -power protected (ϵ, μ) -net with respect to g_0 then $\mathcal{P}' = \{F_0 p_i, p_i \in \mathcal{P}\}$ is a δ -power protected (ϵ, μ) -net with respect to the Euclidean metric.*

Proof This is a direct consequence from the observation that

$$d_0(x, y)^2 = (x - y)^t g_0(x - y) = \|F_0(x - y)\|^2 = d(F_0 x, F_0 y)^2$$

□

5.2 Stability of the net property under metric perturbation

The following lemma shows that the net property is conserved when the metric field is perturbed.

Lemma 5.2 (Net stability under metric perturbation) *Let U , ψ_0 , g and g' be defined as in Lemma 3.2. Let \mathcal{P}_U be a point set in U . Suppose that \mathcal{P}_U is an (ϵ, μ) -net of U with respect to g . Then \mathcal{P}_U is an $(\epsilon_{g'}, \mu_{g'})$ -net of U with respect to g' with $\epsilon_{g'} = \psi_0 \epsilon$ and $\mu_{g'} = \mu / \psi_0$.*

Proof By Lemma 3.2, we have that

$$\forall x, y \in U, \frac{1}{\psi_0} d_{g'}(x, y) \leq d_g(x, y) \leq \psi_0 d_{g'}(x, y).$$

Therefore

$$\forall x \in U, \exists p \in \mathcal{P}_U, d_g(x, p) \leq \epsilon \Rightarrow \forall x \in U, \exists p \in \mathcal{P}_U, d_{g'}(x, p) \leq \psi_0 \epsilon,$$

and

$$\forall p, q \in \mathcal{P}_U, d_g(p, q) \geq \mu \Rightarrow \forall p, q \in \mathcal{P}_U, d_{g'}(p, q) \geq \frac{\mu}{\psi_0}.$$

Thus \mathcal{P} is an $(\epsilon_{g'}, \mu_{g'})$ -net of U with $\epsilon_{g'} = \psi_0 \epsilon$ and $\mu_{g'} = \mu / \psi_0$. □

Remark 5.3 *We assumed that $\mu = \lambda \epsilon$. By Lemma 5.2, we have $\epsilon' = \psi_0 \epsilon$ and $\mu' = \frac{\mu}{\psi_0}$. Therefore*

$$\mu' = \frac{\lambda \epsilon}{\psi_0} = \frac{\lambda}{\psi_0^2} \epsilon'.$$

5.3 Stability of the protection property under metric perturbation

It is more complex to show that the protection assumption is conserved under metric perturbation. We restrict ourselves to considering a uniform metric field for g' . We shall always compare the metric field g in a neighborhood U with the uniform metric field $g' = g(p_0)$ where $p_0 \in U$. Because g' and the Euclidean metric field differ by a linear transformation, we can simplify matters and assume that g' is the Euclidean metric field.

We now give conditions such that the point set is also protected with respect to g' . The main result appears in Lemma 5.15. Lemmas 5.5 and 5.13 provide intermediate results on the stability of Voronoi cells and vertices.

Recall the following definition of a relaxed Voronoi cell:

Definition 5.4 (Relaxed Voronoi cell) *Let $\omega \in \mathbb{R}$. The relaxed Voronoi cell of the seed p_0 is*

$$V_g^{+\omega}(p_0) = \{x \in U \mid d_g(p_0, x)^2 \leq d_g(p_i, x)^2 + \omega \text{ for all } i \neq 0\}.$$

The following lemma expresses that two Voronoi cells computed in similar metric fields are also similar.

Lemma 5.5 *Let U , ψ_0 , g and g' be defined as in Lemma 3.2. Let $\mathcal{P}_U = \{p_i\}$ be a point set in U . Let $V_{p_0, g}$ denote a Voronoi cell with respect to the Riemannian metric field g . Suppose that the Voronoi cell $V_{g'}^{2\rho^2(\psi_0^2-1)}(p_0)$ lies in a ball of radius ρ with respect to the metric g' , which lies completely in U . Let $\omega_0 = 2\rho^2(\psi_0^2 - 1)$. Then $V_g(p_0)$ lies in $V_{g'}^{+\omega_0}(p_0)$ and contains $V_{g'}^{-\omega_0}(p_0)$.*

Proof Let $\text{BS}_g(p_0, p_i)$ be the bisector between p_0 and p_i with respect to g . Let $y \in \text{BS}_g(p_0, p_i) \cap B_{g'}(p_0, \rho)$, where $B_{g'}(p_0, \rho)$ denotes the ball centered at p_0 of radius ρ with respect to g' . Now $d_g(y, p_0) = d_g(y, p_i)$, and thus

$$\begin{aligned} & |d_{g'}(y, p_0)^2 - d_{g'}(y, p_i)^2| \\ &= |d_{g'}(y, p_0)^2 - d_g(y, p_0)^2 + d_g(y, p_i)^2 - d_{g'}(y, p_i)^2| \\ &\leq |d_{g'}(y, p_0)^2 - d_g(y, p_0)^2| + |d_{g'}(y, p_i)^2 - d_g(y, p_i)^2| \\ &\leq (\psi_0^2 - 1)(d_{g'}(y, p_0)^2 + d_{g'}(y, p_i)^2) \\ &\leq 2\rho^2(\psi_0^2 - 1). \end{aligned}$$

Thus $d_{g'}(y, p_0)^2 \leq d_{g'}(y, p_i)^2 + \omega$ and $d_{g'}(y, p_0)^2 \geq d_{g'}(y, p_i)^2 - \omega$ with $\omega = 2\rho^2(\psi_0^2 - 1)$, which gives us the expected result. \square

Remark 5.6 *Recall that we set $\rho = 2\epsilon_0$ and therefore $\omega_0 = 8\epsilon_0^2(\psi_0^2 - 1)$. We show that this choice is reasonable in Remark 5.14.*

Lemma 5.7 *Let g' be a uniform metric field. Let $p_0 \in \mathcal{P}$. The bisectors of $V_{g'}^{\pm\omega_0}(p_0)$ are hyperplanes.*

Proof The bisectors of $V_{g'}^{\pm\omega_0}(p_0)$ are given by

$$\text{BS}_{g'}^{\pm\omega}(p_0, p_i) = \{x \in \Omega \mid d_{g'}(p_0, x)^2 = d_{g'}(p_i, x)^2 \pm \omega\}.$$

For $x \in \text{BS}_{g'}^{\pm\omega}(p_0, p_i)$, we have by definition that

$$\begin{aligned} & \|x - p_0\|_{g'}^2 = \|x - p_i\|_{g'}^2 \pm \omega \\ \iff & (x - p_0)^t g'(x - p_0) = (x - p_i)^t g'(x - p_i) \pm \omega \\ \iff & 2x^t g'(p_i - p_0) = p_i^t g' p_i - p_0^t g' p_0 \pm \omega. \end{aligned}$$

which is the equation of an hyperplane since g' is uniform. \square

The cells $V_{g'}^{\pm\omega_0}(p_0)$ are unfortunately unpractical to manipulate as we do not have an explicit distance between the boundaries $\partial V_{g'}(p_0)$ and $\partial V_{g'}^{\pm\omega_0}(p_0)$. However that distance can be bounded, as shown in the following lemma.

Lemma 5.8 *Let U , ψ_0 , g and g' be defined as in Lemma 3.2. Assume furthermore that g' is the Euclidean metric field. Let $p_0 \in \mathcal{P}$. We have*

$$d_{g'}(\partial V_{g'}(p_0), \partial V_{g'}^{\pm\omega_0}(p_0)) := \min_{\substack{x \in \partial V_{g'}(p_0) \\ y \in \partial V_{g'}^{\pm\omega_0}(p_0)}} d_{g'}(x, y) \leq \frac{\rho^2(\psi_0^2 - 1)}{\mu_0}.$$

Proof Let m_{ω_0} be the intersection of the segment $[p_0, p_i]$ and the bisector $\text{BS}_{g'}^{-\omega_0}(p_0, p_i)$, for $i \neq 0$. Let m be the intersection of the segment $[p_0, p_i]$ and $\partial V_{g'}(p_0)$, for $i \neq 0$. We have

$$\begin{cases} m_{\omega_0} \in \text{BS}_{g'}^{-\omega_0}(p_0, p_i) & \iff 2m_{\omega_0}^T(p_i - p_0) = p_i^T p_i - p_0^T p_0 - \omega_0 \\ m \in \partial V_{g'}(p_0) & \iff 2m^T(p_i - p_0) = p_i^T p_i - p_0^T p_0 \end{cases}$$

Therefore

$$2(m - m_{\omega_0})^T(p_i - p_0) = \omega_0.$$

Since $(m - m_{\omega_0})$ and $(p_i - p_0)$ are linearly dependent,

$$\omega_0 = 2|m - m_{\omega_0}| |p_i - p_0|.$$

\mathcal{P} is μ_0 -separated, which implies that

$$\omega_0 \geq 2|m - m_{\omega_0}| \mu_0,$$

and

$$|m - m_{\omega_0}| \leq \frac{\omega_0}{2\mu_0} = \frac{\rho^2(\psi_0^2 - 1)}{\mu_0}.$$

□

Definition 5.9 *We denote $\eta = \frac{\rho^2(\psi_0^2 - 1)}{\mu_0}$.*

Definition 5.10 (Eroded Voronoi cell) *Let $\eta \in \mathbb{R}$. The eroded Voronoi cell of p_0 is the morphological erosion of $V_g(p_0)$ by a ball of radius η is*

$$\text{EV}_g^{-\eta}(p_0) = \{x \in V_g(p_0) \mid d_g(x, \partial V_g(p_0)) > \eta\}.$$

Definition 5.11 (Dilated Voronoi cell) *Let $\eta \in \mathbb{R}$. The dilated Voronoi cell of p_0 is*

$$\text{DV}_g^{+\eta}(p_0) = \bigcap_{i \neq 0} H^\eta(p_0, p_i),$$

where $H^\eta(p_0, p_i)$ is the half-space containing p_0 and delimited by the bisector $\text{BS}(p_0, p_i)$ translated away from p_0 by η (see Figure 3).

Lemma 5.12 (Encompassing) *We have*

$$\text{EV}_{g'}^{-\eta}(p_0) \subseteq V_{g'}^{-\omega_0}(p_0) \subseteq V_{g'}(p_0) \subseteq V_{g'}^{+\omega_0}(p_0) \subseteq \text{DV}_{g'}^{+\eta}(p_0).$$




Figure 3: Illustration of the different encompassing cells around p_0 . The RVDs with respect to g and g' are respectively traced in red and black. $V_{g'}(p_0)$ is colored in yellow. The cells $DV_{g'}^{+\eta}$ and $EV_{g'}^{-\eta}$ are colored in green, and the $V_{g'}^{\pm\omega_0}(p_0)$ cells are colored in purple.

Proof Using the notations and the result of Lemma 5.8, we have

$$|m - m_{\omega_0}| \leq \eta.$$

Since the bisectors $BS_{g'}^{\omega}(p_0, p_i)$ are hyperplanes, the result follows. \square

On Figures 3, 4 and 11, $DV_{g'}^{+\eta}$ and $EV_{g'}^{-\eta}$ are shown in green and $V_{g'}$ in yellow. The following lemma shows that Voronoi vertices are close if the distortion between the metric fields g and g' is small.

Lemma 5.13 (Voronoi vertex stability) *Let U , g and g' be defined as in Lemma 5.15. Let \mathcal{P}_U be a point set in U . Let $p_0 \in \mathcal{P}_U$. Let c and c' be the same Voronoi vertex in respectively $V_g(p_0)$ and $V_{g'}(p_0)$. Then $d_{g'}(c, c') \leq \chi_0$ with*

$$\chi_0 = \frac{2\eta}{\sqrt{1 + \frac{\mu_0}{2\epsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\epsilon_0}}}.$$

Proof We use Lemma 5.5. $V_g(p_0)$ lies in $DV_{g'}^{+\eta}$ and contains $EV_{g'}^{-\eta}$. The circumcenters c and c' lie in a parallelogrammatic region centered on c' , itself included in the ball centered on c' and with radius χ . The radius χ is given by half the length of the longest diagonal of the parallelogram (see Figure 4). By Lemma 5.2, \mathcal{P} is an (ϵ_0, μ_0) -net with respect to $g' = g_0$. Let θ

be the angle of the Voronoi corner of $V_{g'}(p_0)$ at c' . By Lemma 4.1, that angle can be bounded:

$$\theta_m = 2 \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right) \leq \theta \leq \pi - \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right) = \theta_M.$$

Since $\pi - \theta_M < \theta_m$, χ is maximal when $\theta > \pi/2$. We thus assume $\theta > \pi/2$, and compute a bound on χ as follows:

$$\begin{aligned} \sin\left(\frac{\pi - \theta}{2}\right) &= \frac{\eta}{\chi} \implies \chi = \frac{\eta}{\sin\left(\frac{\pi - \theta}{2}\right)} \\ &\leq \frac{\eta}{\sin\left(\frac{1}{2} \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right)\right)} \\ &\leq \frac{2\eta}{\sqrt{1 + \frac{\mu_0}{2\epsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\epsilon_0}}} =: \chi_0, \end{aligned}$$

using $\sin\left(\frac{1}{2} \arcsin(\alpha)\right) = \frac{1}{2} (\sqrt{1 + \alpha} - \sqrt{1 - \alpha})$.




Figure 4: Black lines trace the $\text{DRVD}_{g'}$ and brown lines trace the DRVD_g . The $V_{g'}(p_0)$ is colored in yellow. The green region corresponds to $DV_{g'}^{+\eta}(p_0)$ and $EV_{g'}^{-\eta}(p_0)$.

□

Remark 5.14 *Lemmas 5.12 and 5.13 allow us to characterize the parameter ρ more precisely. Indeed, an assumption of Lemma 5.5 is that $V_{g'}^{\omega_0}(p_0)$ is included in a ball $B_g(p_0, \rho)$. We show here that if the sampling of \mathcal{P} is sufficiently dense, such an assumption is reasonable.*

We know that $V_{g'}^{\omega_0}(p_0) \subset DV_{g'}^{+\eta}(p_0)$. Since for $x \in V_{g'}(p_0)$, $d_{g'}(p_0, x) \leq \epsilon_0$ and using Lemma 5.13, we have for $x \in DV_{g'}^{+\eta}(p_0)$

$$d_{g'}(p_0, x) \leq \epsilon_0 + \chi_0.$$

We require $V_{g'}^{\omega_0}(p_0) \subset B_{g'}(p_0, \rho)$, which is verified if $DV_{g'}^{+\eta}(p_0) \subset B_{g'}(p_0, \rho)$, that is if

$$\begin{aligned} & \epsilon_0 + \frac{\eta}{\sin\left(\frac{1}{2} \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right)\right)} \leq \rho \\ \Leftrightarrow & \frac{2\rho^2(\psi_0^2 - 1)}{\mu_0\left(\sqrt{1 + \frac{\mu_0}{2\epsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\epsilon_0}}\right)} \leq \rho - \epsilon_0 \\ \Leftrightarrow & \frac{2(\psi_0^2 - 1)}{\sqrt{1 + \frac{\mu_0}{2\epsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\epsilon_0}}} \leq \frac{\mu_0(\rho - \epsilon_0)}{\rho^2} \\ \Leftrightarrow & \frac{2\psi_0(\psi_0^2 - 1)}{\sqrt{1 + \frac{\mu_0}{2\epsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\epsilon_0}}} \leq \frac{\lambda\epsilon(\rho - \epsilon_0)}{\rho^2} \text{ by Remark 5.3.} \end{aligned}$$

The parameter ρ can be chosen arbitrarily as long as it is greater than ϵ_0 . We take $\rho = 2\epsilon_0$, which imposes

$$\begin{aligned} & \frac{\psi_0^2(\psi_0^2 - 1)}{\sqrt{1 + \frac{\mu_0}{2\epsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\epsilon_0}}} \leq \frac{\lambda}{8} \\ \Leftrightarrow & \frac{\psi_0^2(\psi_0^2 - 1)}{\sqrt{1 + \frac{\lambda_0}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda_0}{2\psi_0^2}}} \leq \frac{\lambda}{8}. \end{aligned}$$

Recall that the parameter λ is fixed and $\lambda < 1$. By continuity of the metric field, $\lim_{\epsilon \rightarrow 0} \psi_0 = 1$, therefore the left hand side goes to 0 and this condition is eventually satisfied as the sampling is made denser.

We are now ready to show that the power protection property is conserved when the metric field is perturbed.

Lemma 5.15 (Protection stability under metric perturbation) *Let U , g , g' and ψ_0 be defined as in Lemma 3.2. Assume additionally that g' is the Euclidean metric field. Let \mathcal{P}_U be a δ -power protected (ϵ, μ) -net in U with respect to g . If δ is well chosen, then \mathcal{P}_U is a δ' -power protected net with respect to g' , with*

$$\delta' = \left(\frac{1}{\psi_0^2} - \psi_0^2\right) (\epsilon + \chi_0)^2 - \frac{4\epsilon\chi_0}{\psi_0^2} + \frac{\delta^2}{\psi_0^2}.$$

Proof By Lemma 5.2, we know that \mathcal{P}_U is (ϵ', μ') -net with respect to g' . Let $q \in \mathcal{P}_U$, with q not a vertex in the dual of c . Since \mathcal{P} is a δ -power protected net with respect to g , we have $d_g(c, q) > \sqrt{r^2 + \delta^2}$. On one hand, we have

$$\begin{aligned} d_{g'}(c', q) & \geq d_{g'}(q, c) - d_{g'}(c, c') \\ & \geq \frac{1}{\psi_0} d_g(q, c) - \chi_0 \\ & \geq \frac{1}{\psi_0} \sqrt{r^2 + \delta^2} - \chi_0. \end{aligned}$$

On the other hand, for any $p \in \mathcal{P}_U$ such that p is a vertex of the dual of c , we have

$$\begin{aligned} r' = d_{g'}(c', p) & \leq d_{g'}(c, p) + d_{g'}(c', c) \\ & \leq \psi_0 d_g(c, p) + \chi_0 \\ & \leq \psi_0 r + \chi_0. \end{aligned}$$

Thus δ' -power protection of \mathcal{P}_U with respect to g' requires

$$\begin{aligned} \frac{1}{\psi_0} \sqrt{r^2 + \delta^2} - \chi_0 &> \chi_0 + \psi_0 r \\ \iff \sqrt{r^2 + \delta^2} &> \psi_0(2\chi_0 + \psi_0 r). \end{aligned}$$

This is verified if

$$\begin{aligned} \delta^2 &> (\psi_0(2\chi_0 + \psi_0 r))^2 - r^2 = 4\chi_0^2\psi_0^2 + 4\chi_0\psi_0^3r + \psi_0^4r^2 - r^2 \\ &= 4\chi_0^2\psi_0^2 + 4\chi_0\psi_0^3r + (\psi_0^4 - 1)r^2 \end{aligned}$$

for all $r \in [\mu/2, \epsilon]$. This gives us

$$\delta^2 > 4\chi_0^2\psi_0^2 + 4\chi_0\psi_0^3\epsilon + (\psi_0^4 - 1)\epsilon^2. \quad (3)$$

This condition on δ is only reasonable if the right hand side is not too large. Indeed, since \mathcal{P} is an ϵ -sample, we must have $d_g(c, q) < 2\epsilon$. However, we have that $d_g(c, q)^2 > d_g(c, p)^2 + \delta^2$ by δ -power protection of \mathcal{P} with respect to g . Because $d_g(c, p) < \epsilon$, it suffices that $\delta < \epsilon$. We will now show that this is reasonable by examining the limit of the right hand side of Inequality (3).

We note, see Lemma 5.13, that

$$\chi_0 = \frac{8\epsilon(\psi_0^2 - 1)}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right)},$$

where $\epsilon_0 = \psi_0\epsilon$ and $\mu_0 = \mu/\psi_0 = \lambda\epsilon/\psi_0$. So that

$$\begin{aligned} 4\chi_0^2\psi_0^2 + 4\chi_0\psi_0^3\epsilon + (\psi_0^4 - 1)\epsilon^2 &= \frac{256(\psi_0^2 - 1)\psi_0^2}{\lambda^2 \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right)^2} (\psi_0^2 - 1)\epsilon^2 \\ &+ \frac{32\psi_0^3}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right)} (\psi_0^2 - 1)\epsilon^2 \\ &+ (\psi_0^2 + 1)(\psi_0^2 - 1)\epsilon^2. \end{aligned} \quad (4)$$

This means that the right hand side of Inequality (3) is of the form $f(\psi_0)(\psi_0^2 - 1)\epsilon^2$, where $f(\psi_0)$ is a function that tends to a constant as ψ_0 tends to 1. So the bound given in Inequality (3) may be easily satisfied if the metric distortion is sufficiently small.

We now provide an explicit value for δ' in terms of δ . Let $\xi = d_g(c, q)$ and $\xi' = d_{g'}(c', q')$. We have the following bounds on r' and ξ' :

$$\begin{aligned} \frac{1}{\psi_0}(r - \xi) &\leq r' \leq \psi_0(r + \xi) \\ \frac{1}{\psi_0} \sqrt{(r - \chi_0)^2 + \delta^2} &\leq \xi' \leq \psi_0 \sqrt{(r + \chi_0)^2 + \delta^2}. \end{aligned}$$

If we had $\tilde{\delta}$ -power protection, we would have

$$\begin{aligned} r'^2 + \tilde{\delta}^2 \leq \xi'^2 &\iff \tilde{\delta}^2 \leq \xi'^2 - r'^2 \\ &\iff \tilde{\delta}^2 \leq \frac{1}{\psi_0^2} \left((r - \chi_0)^2 + \delta^2 \right) - \psi_0^2 (r + \chi_0)^2 \\ &\iff \tilde{\delta}^2 \leq \frac{1}{\psi_0^2} (r + \chi_0)^2 - \frac{4r\chi_0}{\psi_0^2} + \frac{\delta^2}{\psi_0^2} - \psi_0^2 (r + \chi_0)^2 \\ &\implies \tilde{\delta}^2 \leq \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\epsilon + \chi_0)^2 - \frac{4\epsilon\chi_0}{\psi_0^2} + \frac{\delta^2}{\psi_0^2}. \end{aligned}$$

Therefore we can take $\delta'^2 = \left(\frac{1}{\psi_0^2} - \psi_0^2\right) (\epsilon + \chi_0)^2 - \frac{4\epsilon\chi_0}{\psi_0^2} + \frac{\delta^2}{\psi_0^2}$. Note that with this definition, δ' goes to δ as ψ_0 goes to 1, which proves that our value of δ' is legitimate. \square

6 The discretized Riemannian Voronoi diagram

We now introduce our algorithm that approximates the exact Riemannian Voronoi diagram.

6.1 Computing geodesics

The computation of geodesic distances and paths on discrete domains is in itself a difficult task to which much work has been dedicated in the last decades. Historically, geodesic distances were first evaluated with the Euclidean metric naturally derived from the embedding of the domain in \mathbb{R}^d . Many approaches exist depending on the expectations on the geodesic distance, exact or approximate, heuristic or with theoretical guarantees [47].

Algorithms have more recently been introduced to compute geodesic distances when the domain is endowed with a prescribed metric field g that exhibits anisotropy. We detail some methods that we empirically evaluated. Konukoglu et al. [35] proposed a heuristic method to extend the fast marching method to the anisotropic setting by adding a recursive correction scheme. However, this implies loss of theoretical convergence of the fast marching method and the recursivity induces a heavy computational cost. Mirebeau [43] introduced a grid-based modification of the fast marching method with theoretical guarantees, but the underlying grid structure does not extend easily to manifolds. Campen et al. [13] proposed a heuristic approach based on a modified Dijkstra algorithm. While their algorithm does not offer guarantees, it is nevertheless fast and precise with little requirement on the quality of the discrete domain. Lastly, a very different approach was considered by Crane et al. [23], who solve linear elliptic equations to compute a global geodesic distance map, obtaining an approximate distance between any two points of the domain. The algorithm is extremely fast compared to front propagation-based methods, but its parameters require delicate tuning and the approximation is generally much worse, making it unattractive for our work.

6.2 Canvas generation

We refer to the discretization of the domain Ω upon which geodesic distances are computed as the *canvas*, and denote it \mathcal{C} . Notions related to the canvas explicitly carry *canvas* in the name (for example, an edge of \mathcal{C} is a *canvas edge*). Except in Section 7, an isotropic triangulation is used as canvas. The generation of this triangulation is performed through an isotropic refinement algorithm driven by a *sizing field* constraint s : if a simplex σ of \mathcal{C} has a circumradius R_σ larger than $s(c_\sigma)$, where c_σ is the circumcenter of σ , then it is refined by inserting c_σ . We will show how the sizing field is chosen in Section 8.

6.3 Generation of the diagram

The procedure used to generate our approximate diagram is similar to the classic Dijkstra algorithm.

Each seed is associated a color (an integer) and each canvas vertex is assigned three fields:

- a status: KNOWN, TRIAL or FAR (initialized at FAR),
- a distance: the current smallest geodesic distance to a seed (initialized at ∞),

- a color: the color of the closest Voronoi seed.

The initialization of the canvas with the seed set is done by locating each seed on the canvas. The vertices of the simplex containing the seed are then attributed TRIAL as status, their color is set to the color of the seed and their distance is set to the distance to the seed.

The vertices of the canvas are colored through simultaneous Dijkstra-like front propagations from all the seeds: the vertices are sorted in a priority queue \mathcal{Q} that is kept ordered by increasing distance to the vertex's closest seed. At each iteration, the point with lowest distance to a seed becomes KNOWN and is removed from the queue. The distances at its neighbors that are not KNOWN are evaluated again, and the queue is updated.

When the queue is empty, the result is a partition of the vertices of the canvas. To each seed p_i , we associate a *discrete cell* $V^d(p_i)$ made of all canvas simplices who have at least one vertex of the color of p_i . We call the set of these cells the *discretized Riemannian Voronoi diagram* (DRVD).

Refinement of the DRVD.

New seeds can easily be inserted in an existing DRVD: the Voronoi cell of a new point is obtained by spreading from this new seed with a single front propagation algorithm. When the sampling of the point set is already dense on the domain, the insertion is a local operation as only a relatively small amount of canvas simplices carry the color of any seed.

6.4 Extracting the nerve of the DRVD

To construct a geometric realization of the nerve \mathcal{N} , we must first extract the connectivity of the DRVD. Using a triangulation as canvas allows us to obtain \mathcal{N} easily: each canvas k -simplex σ of \mathcal{C} has $k + 1$ vertices $\{v_0, \dots, v_k\}$ with respective colors $\{c_0, \dots, c_k\}$, which correspond to seeds $\{p_{c_0}, \dots, p_{c_k}\} \in \mathcal{P}$. If all colors c_i are different, then the intersection $D = \bigcap_{i=0 \dots k} V^d(p_{c_i})$ is non-empty and D belongs to the nerve (see Figure 7). We say that the simplex σ *witnesses* (or is a witness of) D . If for all $D \in \mathcal{N}$ there exists $\sigma \in \mathcal{C}$ such that σ witnesses D , and if for all $\sigma \in \mathcal{C}$ the simplex witnessed by σ belongs to \mathcal{N} , then we say that \mathcal{C} *captures* \mathcal{N} .

7 Fast DRVD computations

The computation of geodesic distances constitutes the expected bottleneck in the construction of the DRVD. This cost can be reduced in two independent ways: improving the speed of geodesic computations and decreasing the amount of geodesic distance estimations. We rely on previous work for the estimation of the geodesic distances and improving the current methods goes beyond the scope of this report. We can however dramatically decrease the number of geodesic computations required by changing the nature of the canvas, as we shall explain now.

Assume that a number of consecutive canvas edges form a line in the domain. If the metric were constant along this line, the accuracy of the distance computation would not be improved by adding edges. Extrapolating from this simple example, one can intuit that it is preferable for the canvas not to be an isotropic discretization of the domain: directions in which the metric field does not vary too much can be discretized only sparsely, as long as other constraints are honored (see the theoretical requirements in Section 8). Unfortunately, an anisotropic mesh is exactly what we want to construct. However, for geodesic distance computations, a triangulation is not required and a graph suffices. The locally Delaunay algorithm of Boissonnat et al. [10] gives such a graph based on the connectivity of its stars. We recall in the next section the main principles of this algorithm.

7.1 The local Delaunay algorithm

The local Delaunay algorithm, introduced by Boissonnat et al. [10], is an attempt to extend the concept of Delaunay triangulation to non Euclidean settings by computing independently at each vertex its connectivity to other vertices. The algorithm constructs for every vertex $v \in \mathcal{P}$ the Delaunay triangulation in the metric M_v of v , which is the image through the stretching transformation F_v^{-1} of the Euclidean Delaunay triangulation $\text{Del}(F_v(\mathcal{P}))$. The *star* S_v of v in $\text{Del}_v(\mathcal{P})$ is the sub-complex of $\text{Del}_v(\mathcal{P})$ formed by the simplices that are incident to v . The collection of all the stars is called the (anisotropic) star set of \mathcal{P} , noted $S(\mathcal{P})$.

Since the connectivity of each star is set according to the metric at the center of the star, *inconsistencies* may exist amongst the stars of the sites: a simplex σ that appears in $\text{Del}_v(\mathcal{P})$ might not appear in $\text{Del}_w(\mathcal{P})$ with $v, w \in \sigma$ (see Figure 5, left). The algorithm refines the set of sites \mathcal{P} until each simplex conforms to a set of a criteria (size, quality, etc.) and there are no more inconsistencies amongst the stars, which is proven to happen.




Figure 5: Two stars in an inconsistent configuration, [9].

Unfortunately, the algorithm requires in practice a large number of vertices to solve inconsistencies when dealing with highly-distorted metric fields, making it unattractive to use. Figure 6 shows an example of metric field that is difficult for the star set. The algorithm required 4430 vertices to achieve consistency, a number considerably higher than our method as can be seen on the results produced in Section 14.

7.2 Using star sets as canvasses

While the local Delaunay algorithm is not competitive, its star set structure provides, regardless of its consistency, a graph that can be used to compute geodesic distances. Since the stars conform to the metric field, this anisotropic canvas is much sparser than an equivalent isotropic triangulation. This difference in density is accentuated as the anisotropy ratio increases. Indeed, the Voronoi cells become very thin for extreme anisotropy ratios and require a very dense sampling when using isotropic canvasses. On the other hand, the amount of vertices in a star set-based canvas can be kept at around two orders of magnitude that of the point set regardless of the anisotropy. In other words, it suffices in practice that every Riemannian Voronoi cell contains around 100 canvas vertices to capture the nerve. Finally, it is fast to compute and can be trivially parallelized.




Figure 6: Example of a difficult metric field for the local Delaunay algorithm: the hyperbolic shock in a rectangular region of the plane. 4430 vertices are required to obtain a consistent star set. Comparison can be made with our method in Figure 21.

The use of the star set as canvas therefore greatly decreases the computational time as the amount of vertices that make up the canvas is drastically reduced without any change in the extracted nerve (see Figure 7 and Section 14). This approach is not unlike the stencil-based geodesic distance computation method of Mirebeau [43] previously described. Indeed, our stars can be considered as non-uniform stencils. Although they offer less guarantees than Mirebeau's stencils, Delaunay stars allow to handle easily the case of (curved) manifolds.

The theoretical results presented in Section 8 use an isotropic triangulation as canvas for simplicity, but these results can be extended, under some conditions, to star set canvasses that are not necessarily consistent.

8 Overview of the theoretical sections

Our theoretical work consists of two parts. First, we expose sufficient conditions such that the nerves of the DRVD and the RVD coincide. We then give requirements such that the geodesic and straight edge duals of the RVD are embedded triangulations. These conditions will be expressed through sampling requirements on the point set \mathcal{P} with respect to the input metric and through density requirements on the sizing field of the canvas. To simplify computations, we will consider that the sizing field is constant over the domain.

In the following, we assume that our domain has no boundaries and is geodesically complete in order to avoid punctures. Until Section 9.4, we assume that geodesic distances are computed




Figure 7: Discretized Riemannian Voronoi diagram of the same seed set on two different canvas. On the left, an isotropic triangulation (1M vertices). On the right, an inconsistent star set (40k vertices). The nerves of both DRVD are identical. The canvas simplices colored in red are witnesses of Voronoi vertices. The canvas simplices colored in grey are witnesses of Voronoi edges.

exactly.

9 Equating the nerves of the DRVD and the RVD

We give conditions on the seed set \mathcal{P} and the canvas \mathcal{C} that guarantee that the DRVD has the same nerve as the exact RVD. Our analysis consists of four steps. We firstly show that the Voronoi cells are well formed in Euclidean space and deduce conditions on \mathcal{P} and \mathcal{C} such that the nerves coincide in the Euclidean setting. These can be then be translated to conditions in the cases of uniform and arbitrary Riemannian metric fields in \mathbb{R}^2 . Finally, we solve the case of a general surface.

9.1 Equating the nerves of the DRVD and the RVD: Euclidean metric field

We prove in details that if the sampling and the canvas are dense enough, the nerve of our DRVD coincide with the nerve of the exact Riemannian Voronoi diagram, in the context of the Euclidean space $\Omega = \mathbb{R}^2$. The theorem to prove is the following:

Theorem 9.1 *Let \mathcal{P} be an (ϵ, μ) -net of \mathbb{R}^2 that is δ -power protected. Let \mathcal{C} be the canvas, a triangulation of \mathbb{R}^2 . Let s_0 be the uniform sizing field of \mathcal{C} . If s_0 is small enough, then $\text{Nrv}(\text{DRVD}) = \text{Nrv}(\text{RVD})$.*

More explicit requirements and sufficient conditions are given in Lemmas 9.2 and 9.3.

Lemma 9.2 *Let $p_0 \in \mathcal{P}$ be a seed of the RVD. Let $V(p_0)$ be its corresponding cell. If*

$$s_0 < \frac{\delta^2 \mu}{8\epsilon^2} \sqrt{1 - \left(\frac{\mu}{2\epsilon}\right)^2}$$

then no edge of \mathcal{C} intersects two non-adjacent Voronoi faces of V_0 .

Lemma 9.3 *Let $p_0 \in \mathcal{P}$ be a seed of the RVD. Let $V(p_0)$ be its corresponding cell. If*

$$s_0 < \frac{\delta^2 \mu}{8\epsilon^2} \sqrt{1 - \left(\frac{\mu}{2\epsilon}\right)^2}$$

then every Voronoi vertex of $V(p_0)$ is witnessed by a canvas simplex.

Remark 9.4 Lemma 9.2 ensures that $\text{Nrv}(\text{DRVD}) \subseteq \text{Nrv}(\text{RVD})$ and Lemma 9.3 ensures that $\text{Nrv}(\text{RVD}) \subseteq \text{Nrv}(\text{DRVD})$.

The rest of this section is dedicated to proving Lemmas 9.2 and 9.3. Theorem 9.1 is a straightforward consequence of these lemmas. We start by proving Lemma 9.2.

Assume that $V(p_0)$ has at least 3 vertices, otherwise we are done as any Voronoi edge of $V(p_0)$ is adjacent to the other Voronoi edges. Let $\{F_i\}$ be the set of faces (edges) of the Voronoi cell $V(p_0)$. One seeks a lower bound on the minimum of

$$d(F_j, F_k) = \min_{x \in F_j, y \in F_k} d(x, y)$$

where F_j and F_k are not adjacent. Note that by convexity of the Voronoi cell, the minimum of this distance will be attained at a vertex of a face.

Lemma 9.5 Assume that \mathcal{P} is a δ -power protected (ϵ, μ) -net. Let $l = \min_{i,j} d(v_i, F_j)$ where F_j is a Voronoi edge whose extremities are not v_i . The length l can be bounded from below:

$$l \geq s_m = \frac{\delta^2 \mu}{8\epsilon^2} \sqrt{1 - \left(\frac{\mu}{2\epsilon}\right)^2}.$$




Figure 8: Corner of a Voronoi cell. l is the minimal distance between two non-adjacent edges.

Proof Figure 8 shows an example of a Voronoi cell, with 6 Voronoi vertices. Assume without loss of generality that v_1 and $F = [v_0v_5]$ are the arguments that realize the minimum of the distance and let θ be the angle at v_1 . Since we have δ -power protection, the Voronoi vertices of $V(p_0)$ are separated and $|v_0 - v_1| \geq \frac{\delta^2}{4\epsilon}$ by Lemma 4.6. If $\theta \geq \pi/2$, we immediately have $l \geq \frac{\delta^2}{4\epsilon}$. Assume now that $\theta \leq \pi/2$. Let $v = |v_0 - v_1|$. From Lemma 4.1, we have

$$\theta \geq 2 \arcsin \left(\frac{\mu}{2\epsilon} \right)$$

Therefore,

$$l = \sin(\theta)v \geq \sin \left(2 \arcsin \left(\frac{\mu}{2\epsilon} \right) \right) \frac{\delta^2}{4\epsilon} = \frac{\delta^2 \mu}{4\epsilon^2} \sqrt{1 - \left(\frac{\mu}{2\epsilon} \right)^2},$$

using $\sin(2\alpha) = 2\sin(\alpha)\sqrt{1 - \sin^2(\alpha)}$. Since

$$\frac{\mu}{\epsilon} \sqrt{1 - \left(\frac{\mu}{2\epsilon}\right)^2} \leq 1,$$

the minimum is attained when $\theta \leq \pi/2$ and we have

$$s_m = \frac{\delta^2 \mu}{4\epsilon^2} \sqrt{1 - \left(\frac{\mu}{2\epsilon}\right)^2}.$$

□

If the canvas \mathcal{C} honors a uniform sizing field s , all the edges e in \mathcal{C} have length $|e| < 2s$. Thus if $s_0 < s_m/2$, then no edge crosses two non-adjacent Voronoi edges. This ensures that for all $D \in \text{Nrv}(\text{DRVD})$, then $D \in \text{Nrv}(\text{RVD})$, which proves Lemma 9.2.

We now prove Lemma 9.3; that is, if s_0 is small enough, then every Voronoi vertex is witnessed by a at least one canvas simplex.

Since \mathcal{C} is a partition of \mathbb{R}^2 , every Voronoi vertex is located in a maximal canvas simplex $\sigma_i \in \mathcal{C}$.

Remark 9.6 *A Voronoi vertex v_i might be located on a canvas edge (or even a canvas vertex) and belong to more than one canvas simplex. However this does not create any issue in the following proofs, and one can select randomly any canvas simplex that witnesses the Voronoi vertex.*

Two issues potentially prevent the canvas from capturing the nerve of the RVD:

- multiple Voronoi vertices v_i can belong to the same canvas simplex,
- the canvas simplex σ_i in which v_i is located might not witness v_i (see Figure 9 for a problematic case). This happens when a canvas edge crosses over two (adjacent) Voronoi edges, which is not prevented by Lemma 9.2.

We prove that these issues can be remedied when the canvas is dense enough with the following lemmas.

Lemma 9.7 *If $s_0 < \frac{\delta^2}{8\epsilon}$, then no two Voronoi vertices are located in the same maximal canvas simplex.*

Proof If s_0 is the sizing field of \mathcal{C} , then any edge of \mathcal{C} has its length bounded by $2s_0$. Since we have assumed δ -power protection of the Delaunay simplices, the Voronoi vertices are separated by $\frac{\delta^2}{4\epsilon}$ by Lemma 4.6. Therefore, if $s_0 < \frac{\delta^2}{8\epsilon}$ then every Voronoi vertex falls in a different canvas simplex. □

Recall that for a given v_i , σ_i is the associated canvas simplex that contains v_i . Assume that not every σ_i is also a witness of v_i (otherwise we are done since we have a bijection between $\{v_i\}$ and $\{\sigma_i\}$ by Lemma 9.7). Without loss of generality, we can assume that v_0 is such vertex (thus we assume that σ_0 is not a witness of v_0). We will show that if $s_0 \leq s_m/2$, there exists a canvas simplex that witnesses of v_0 .

Lemma 9.8 *The vertices of σ_0 are colored with two different colors.*




Figure 9: Illustration of a canvas (in green) and three Voronoi cells (boundaries in red). On the left, the canvas simplex containing the Voronoi vertex v does not witness the dual of that Voronoi vertex.




Figure 10: the RVD is traced in black, the canvas is shown with green edges. The algorithm starts in the triangle $\Delta q_0 q_1 q_2$ and eventually finds the witness $\Delta q_4 q_5 q_6$ of v_0 .

Proof A canvas simplex is colored with at most three different colors. Since we have assumed σ_0 does not witness v_0 and that s_0 is small enough such that the Voronoi vertices fall in different canvas simplices, σ_0 is not colored with three different colors. The vertices cannot either have all the same colors because v_0 is located in σ_0 and Voronoi cells are convex. Therefore, the vertices of σ_0 are painted with two different colors. \square

Assume that the conditions of Lemma 9.2 are fulfilled. No canvas edge can intersect two non-adjacent Voronoi faces. Let q_0, q_1 and q_2 be the vertices of σ_0 . Without loss of generality, we can assume that q_0 is the vertex that has a different color (see Figure 10 for an illustration). Let c_0 and c_1 be the colors at respectively q_0 and q_1 . Let c_2 be the missing color to capture v_0 .

Let p_0 , p_1 and p_2 be the seeds of the cells V_0 , V_1 and V_2 (whose respective colors are c_0 , c_1 and c_2).

Begin by picking the edge e_0 of σ_0 that is closest to the seed p_2 . This is an edge $e_0 = [p_0, p_i]$ with $i = 1$ or 2 . Again without loss of generality, assume $i = 1$. Consider the canvas simplex σ_1 that shares the edge e with σ_0 . Let q_3 and c_3 be the vertex of σ_1 that does not belong to e .

There are three possibilities:

- $c_3 = c_2$. We are done: σ_1 witnesses v_0 .
- $c_3 = c_0$ or $c_3 = c_1$. Repeat the same process: pick the edge of σ_1 that is closest to p_2 and has two different colors; find the adjacent σ_2 and look at the color of its third vertex. This algorithm finishes since we move closer to p_2 with every step and the mesh is finite.
- $c_3 = c_4$ with $c_4 \neq c_0, \dots, c_2$. This is impossible because it means one of the edges of σ_1 crosses two non-adjacent Voronoi edges, which contradicts our density assumption.

Therefore, we eventually reach a σ_i that has three different colors, and witnesses v_0 .

This proves Lemma 9.3 and ensures that $\forall D \in \text{Nrv}(\text{RVD}), D \in \text{Nrv}(\text{DRVD})$. And thus if

$$s_0 < \min\left(\frac{s_m}{2}, \frac{\delta^2}{8\epsilon}\right) = \frac{s_m}{2} = \frac{\delta^2 \mu}{8\epsilon^2} \sqrt{1 - \left(\frac{\mu}{2\epsilon}\right)^2},$$

then $\text{Nrv}(\text{RVD}) = \text{Nrv}(\text{DRVD})$.

9.2 Equating the nerves of the DRVD and the RVD: uniform metric field

We prove here that in the case of a uniform metric field, density conditions on the sampling of the point set and on the canvas can be exhibited such that the nerves of the DRVD and the RVD coincide.

By Lemma 5.1, if a point set \mathcal{P} in \mathbb{R}^2 is a δ -power protected (ϵ, μ) -net with respect to g_0 then the point set $\mathcal{P}' = \{F_0 p_i, p_i \in \mathcal{P}\}$, where F_0 is the square root of g_0 , is a δ -power protected (ϵ, μ) -net with respect to the Euclidean metric. We can deduce an upper bound on the sizing field of \mathcal{C} in the case of a uniform metric field using the results of section 9.1 for the Euclidean setting.

Theorem 9.9 *Let \mathcal{P} be a point set in Ω . Let g be a uniform metric field on Ω (for all $x \in \Omega$, $g(x) = g_0$). Let \mathcal{C} be the canvas, and s_0 its sizing field. Let $\{\lambda_i\}$ and $\{v_i\}$ be respectively the eigenvalues and the normalized eigenvectors of the matrix g_0 . If*

$$s_0 < \left(\min_i \sqrt{\lambda_i}\right) \left(\frac{\delta^2 \mu}{8\epsilon^2} \sqrt{1 - \left(\frac{\mu}{2\epsilon}\right)^2}\right),$$

then the canvas captures the nerve of the RVD.

Proof Let \mathcal{D} be the DRVD of \mathcal{P} over \mathcal{C} . Let F_0 be a square root of g_0 . Let $\mathcal{P}_0 = \{F_0 p, p \in \mathcal{P}\}$ be the transformed point set. Let \mathcal{C}_0 be a canvas of the transformed space. Let \mathcal{D}_0 be the DRVD of \mathcal{P}_0 in the Euclidean metric, over \mathcal{C}_0 . Let s_0 be the upper bound on the sizing field of \mathcal{C}_0 provided by Theorem 9.1 such that \mathcal{D}_0 is captured by \mathcal{C}_0 .

Given an eigenvector v_i , the corresponding eigenvector has length $\frac{1}{\sqrt{\lambda_i}}$ in the Euclidean space. Let \mathcal{C}'_0 be the image of \mathcal{C} by F_0 . Note that \mathcal{C}'_0 and \mathcal{C}_0 are two different triangulations of the same space. If the sizing field s_0 of \mathcal{C} is smaller than αs_0 , with

$$\alpha = \frac{1}{\max_i \left(\frac{1}{\sqrt{\lambda_i}} \right)} = \min_i \sqrt{\lambda_i},$$

then every edge of \mathcal{C}'_0 is smaller than $2s_0$. This implies that \mathcal{C}'_0 captures \mathcal{P}_0 and therefore \mathcal{C} captures \mathcal{P} . \square

This settles the case of a uniform metric field.

9.3 Equating the nerves of the DRVD and the RVD: arbitrary metric field.

We consider an arbitrary metric field g over \mathbb{R}^2 . In a sufficiently small neighborhood of \mathbb{R}^2 , the metric field is almost constant, and can be well approximated with a uniform metric field, a configuration that we have dealt with in the previous section. We shall always compare the metric field g in a neighborhood U with the uniform metric field $g' = g(p_0)$ where $p_0 \in U$. Because g' and the Euclidean metric field differ by a linear transformation, we can simplify matters and assume that g' is the Euclidean metric field.

The core hypothesis to obtain a bound on the sizing field in the Euclidean and uniform metric field settings was that the point set \mathcal{P} formed a protected net. We have shown in Lemmas 5.2 and 5.15 that if this assumption is made on \mathcal{P} with respect to the Riemannian metric field g , it can be carried over as we locally approximate the metric field with a uniform metric field g' .

Theorem 9.10 *Let U , g and g' be defined as in Lemma 3.2. Let \mathcal{P}_U be a point set in U and let $p_0 \in \mathcal{P}_U$. Suppose that \mathcal{P}_U is a δ -power protected (ϵ, μ) -net of with respect to g . Let $V_g(p_0)$ be the Voronoi cell of p_0 in the DRVD of \mathcal{P}_U with respect to g . If the sampling criterion ϵ is small enough, there exists a sizing field s_0 small enough such that the cell $V_g(p_0)$ is captured.*

Lemma 5.5 gives us that $V_g(p_0)$ lies in $DV_{g'}^{+\eta}(p_0)$ and contains $EV_{g'}^{-\eta}(p_0)$. Since $V_g(p_0)$ contains $EV_{g'}^{-\eta}(p_0)$, if s_0 is small enough such that $EV_{g'}^{-\eta}(p_0)$ is captured, then $V_g(p_0)$ is also captured. Proving that $EV_{g'}^{-\eta}(p_0)$ is captured is done similarly to the Euclidean case. While we do not explicitly have a protected net when dealing with relaxed Voronoi cells, we can still extract the core property that the Voronoi vertices are separated, as shown by the next lemma.

Lemma 9.11 (Separation of the vertices of the eroded Voronoi cell) *Assume U , g and $g' = g_0$ as above. Assume that the point set \mathcal{P}_U is a δ -power protected (ϵ, μ) -net with respect to the Riemannian metric field g . Then the Voronoi vertices of $EV_{g'}^{-\eta}(p_0)$ are separated.*

Proof The notations used below are illustrated on Figure 11. By Lemmas 5.2 and 5.15, \mathcal{P} is δ_0 -power protected (ϵ_0, μ_0) -net with respect to g_0 . Let $L_0 = \frac{\delta_0^2}{4\epsilon_0}$ be the separation bound induced by the δ_0 -power protection property of \mathcal{P}_U (see Lemma 4.6). Let l be the distance between any two adjacent Voronoi vertices of $EV_{g'}^{-\eta}(p_0)$. Let $\theta_0 = 2 \arcsin \left(\frac{\mu_0}{2\epsilon_0} \right)$ be the lower bound on the angle of a corner of $V_{g'}(p_0)$ (see Lemma 9.5). Remark that the lower bound on the smallest angle θ_0 in a Voronoi cell is unchanged by the erosion of the Voronoi cell. We have

$$l = L - 2l, \text{ with } l = \frac{\eta}{\tan \left(\frac{\theta}{2} \right)}.$$




Figure 11: Eroded Voronoi cell and a zoom on one of the corners.

Noting that

$$\tan(\arcsin(x)) = \frac{x}{\sqrt{1-x^2}} = \frac{1}{\sqrt{\frac{1}{x^2}-1}},$$

we obtain

$$\begin{aligned} l &= L - 2 \frac{\eta}{\tan\left(\frac{\theta}{2}\right)} \\ &\geq L - 2 \frac{\eta}{\tan\left(\arcsin\left(\frac{\mu_0}{2\epsilon_0}\right)\right)} \\ &= L - 2\eta \sqrt{\left(\frac{2\epsilon_0}{\mu_0}\right)^2 - 1} \\ &\geq \frac{\delta_0^2}{4\epsilon_0} - 2 \frac{\rho_0^2(\psi_0^2 - 1)}{\mu_0} \sqrt{\left(\frac{2\epsilon_0}{\mu_0}\right)^2 - 1}. \end{aligned}$$

We need l to be positive. This can be achieved by enforcing that the lower bound is positive:

$$\frac{\delta_0^2}{4\epsilon_0} > 2 \frac{\rho_0^2(\psi_0^2 - 1)}{\mu_0} \sqrt{\left(\frac{2\epsilon_0}{\mu_0}\right)^2 - 1}$$

Recall that $\epsilon_0 = \psi_0\epsilon$, $\mu_0 = \lambda\epsilon/\psi_0$ and $\rho_0 = 2\psi_0\epsilon$. Using this notation we see that $l > 0$ if

$$\delta_0^2 > \frac{32}{\lambda} (\psi_0^2 - 1) \epsilon^2 \psi_0^4 \sqrt{\left(\frac{2\psi_0^2}{\lambda}\right)^2 - 1}. \quad (5)$$

This condition is easy to satisfy when ψ_0 goes to 1, because the right hand side of Inequality (5) is proportional to $(\psi_0^2 - 1)\epsilon^2$. We however would like to give this condition in terms of δ so that it may be compared with Inequality (3). In Lemma 5.15, we have seen that

$$\delta_0^2 = \frac{\delta^2}{\psi_0^2} + \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\epsilon + \chi_2)^2 - \frac{4\epsilon\chi_2}{\psi_0^2},$$

with, by Lemma 5.13,

$$\chi_2 = \frac{8\psi_0^3\epsilon(\psi_0^2 - 1)}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right)}.$$

Thus

$$\frac{\delta^2}{\psi_0^2} + \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\epsilon + \chi_2)^2 - \frac{4\epsilon\chi_2}{\psi_0^2} > \frac{32}{\lambda} (\psi_0^2 - 1) \epsilon^2 \psi_0^4 \sqrt{\left(\frac{2\psi_0^2}{\lambda} \right)^2 - 1}$$

which is equivalent to

$$\delta^2 > \frac{32}{\lambda} (\psi_0^2 - 1) \epsilon^2 \psi_0^6 \sqrt{\left(\frac{2\psi_0^2}{\lambda} \right)^2 - 1} - (1 - \psi_0^4) (\epsilon + \chi_2)^2 + 4\epsilon\chi_2.$$

After simplifying the $(\epsilon + \chi_2)$ -term, we find

$$\delta^2 > \frac{32}{\lambda} (\psi_0^2 - 1) \epsilon^2 \psi_0^6 \sqrt{\left(\frac{2\psi_0^2}{\lambda} \right)^2 - 1} + \frac{32\psi_0^3(\psi_0^2 - 1)\epsilon^2}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right)} + (\psi_0^2 - 1) \epsilon^2 (\psi_0^2 + 1). \quad (6)$$

If we compare this to the bound given by Inequality (3), made explicit in Inequality (4), namely

$$\begin{aligned} \delta^2 &> \frac{256(\psi_0^2 - 1)\psi_0^2}{\lambda^2 \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right)^2} (\psi_0^2 - 1) \epsilon^2 \\ &+ \frac{32\psi_0^3}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right)} (\psi_0^2 - 1) \epsilon^2 \\ &+ (\psi_0^2 + 1)(\psi_0^2 - 1) \epsilon^2, \end{aligned} \quad (7)$$

we see that the two final terms in the sum in Inequalities (6) and (7) are of equal magnitude. The final terms would in fact agree if we would not have simplified the equation, because χ_2 is of order $\epsilon(\psi_0^2 - 1)$. The first term disappears for Inequality (7) in the limit where ψ_0 tends to 1, while this is not the case for the first term in Inequality (6).

This means that at least when ψ_0 goes to 1, Equation 6 provides the tougher bound. However both inequalities have to be verified explicitly in a practical setting. \square

We can now provide a bound on the sizing field so that it captures $EV_{g'}^{-\eta}(p_0)$. Similarly to the Euclidean case (see Lemma 9.5), the bound is given by

$$s_m = \sin(\theta_0) \left(L_0 - \frac{2\eta}{\tan\left(\frac{\theta_0}{2}\right)} \right).$$

Following the same logic than in Lemmas 9.2 and 9.3, the final bound s_0 is $\min(\frac{s_m}{2}, \frac{l}{2})$. From the relaxed Euclidean Voronoi cell, we can deduce a bound on the sizing field needed to capture the relaxed Voronoi cell in the case of a uniform metric field, as was done in Lemma 9.9. Therefore, for a point set and a canvas sufficiently dense, $EV_g^{-\eta}(p_0)$ as well as $V_g^{-\eta}(p_0)$ are captured, and thus $V_g(p_0)$ is captured.

Taking the minimum of all the sizing field values over all $p \in \mathcal{P}$, we get a sizing field such that all the Voronoi cells are captured. Finally, this proves that $\text{nrv}(\text{DRVD}) = \text{nrv}(\text{RVD})$ in the general setting, when the geodesics are exactly computed.

9.4 Approximate geodesic computations

We have so far assumed that geodesics are computed exactly. This is however not the case in practice. Nevertheless, once the error in the approximation of the geodesic distances is small enough, the computation of the DRVD with approximate geodesic distances can be equivalently seen as the computation of a second DRVD with a slightly different metric field, using exact geodesic distances. Another layer is again added to the proof, handled similarly to the case of an arbitrary metric field.

10 Riemannian Delaunay triangulation

We have now given sufficient conditions for the DRVD and the AVGD to have the same nerve. The next step is to prove that that the Delaunay complex is homeomorphic to the manifold.

Several works have been dedicated to conditions that guarantee this, see for example Edelsbrunner and Shah [29] and Boissonnat, Dyer and Ghosh [5, 6]. We shall be relying on the latter ones for many results. In two dimensions, the Delaunay complex is homeomorphic to the surface under very light conditions, see [38] and [28]. We prove the same results using a more geometrical approach where we “paint” the Delaunay simplex on the manifold in an intrinsic manner and, along the way, fix a flaw in [28].

10.1 Riemannian simplices

For each simplex $\sigma_{\mathcal{D}}$ in the Delaunay complex we can consider its Riemannian realization. The definition of the Riemannian realization is based on the Riemannian center of mass [33, 26], also known as a Karcher mean [34].

Let us denote by ψ^n the n -dimensional standard simplex with coordinates $\lambda = (\lambda_i)$, that is

$$\psi^n = \{(\lambda_i) \in \mathbb{R}^{n+1} \mid \lambda_i \geq 0, \sum_i \lambda_i = 1\}.$$

If the dimension is obvious, we shall drop it from the notation. For any $n + 1$ points p_i on the manifold and $(\lambda) \in \psi^n$ we can define

$$\mathcal{E}_\lambda(x) = \frac{1}{2} \sum_i \lambda_i d_M(x, p_i)^2.$$

We shall call the points p_i vertices.

Theorem 10.1 (Karcher) *Let M be a manifold whose sectional curvature K is bounded, that is $\Lambda_- \leq K \leq \Lambda_+$. Let us consider the function \mathcal{E}_λ on B_ρ , a geodesic ball of radius ρ that contains the set $\{p_i\}$. If*

- ρ is less than half the injectivity radius,
- if $\Lambda_+ > 0$ then

$$\rho < \frac{\pi}{4\sqrt{\Lambda_+}},$$

then \mathcal{E}_λ has a unique minimum point in B_ρ , which is called the center of mass.

The Riemannian realization of a n -dimensional Delaunay simplex $\sigma_{\mathcal{D}}$ of M whose vertices lie in a sufficiently small ball, so that the conditions of Theorem 10.1 are satisfied, can now be defined as:

Definition 10.2 *The image of ψ under the map*

$$\begin{aligned} \mathcal{B}_{\sigma_{\mathcal{D}}} : \quad \psi &\rightarrow M \\ \lambda &\mapsto \operatorname{argmin}_{x \in \bar{B}_\rho} \mathcal{E}_\lambda(x) \end{aligned}$$

is called the Riemannian realization of the Delaunay simplex $\sigma_{\mathcal{D}}$ and is denoted by $\sigma_{\mathcal{D},M}$.

As in [26], we can define non-degeneracy of a geometric simplex.

Definition 10.3 *The Riemannian realization of a Delaunay simplex, $\sigma_{\mathcal{D},M}$, is said to be non-degenerate if, and only if, it is homeomorphic to ψ .*

10.2 2-manifolds

The question of non-degeneracy of Riemannian simplices is non-trivial in arbitrary dimension, see [26]. In two dimension it is sufficient for three points not to lie on a geodesic, see Rustamov [48] and Dyer, Vegter and Wintraecken [26] or Wintraecken [54].

We are interested in conditions that guarantee that the Riemannian realization of the Delaunay complex gives a homeomorphism between the Delaunay complex \mathcal{D} and the manifold M . The homeomorphism is in fact a piecewise smooth homeomorphism where the homeomorphism per simplex $\sigma_{\mathcal{D},M}$ is given by the maps $\mathcal{B}_{\sigma_{\mathcal{D}}}$. It is of course necessary that all simplices $\sigma_{\mathcal{D},M}$ in the complex are non-degenerate.

We start with some definitions:

Definition 10.4 *A mapping f of an oriented simplicial complex K that is a n -manifold (with boundary), into Euclidean space \mathbb{E}^n endowed with an orientation, is called simplex-wise positive if for each top dimensional simplex σ_i^n , f is smooth and bijective in σ_i^n (the partial derivatives being continuous on the boundary), as well as orientation preserving. For any simplicial complex L , let L^k , or $(L)^k$, denote the subcomplex containing all cells of L of dimension $\leq k$. With f in K as above, any point q of $\mathbb{E}^n \setminus f(K^{n-1})$ is in the image of a certain number of h of n -simplices of K ; we say q is covered h times.*

Lemma 10.5 (Whitney's lemma [53]) *Let f be simplex-wise positive in K . Then for any connected open subset R of $\mathbb{E}^n \setminus f(\partial K)$, any two points of R not in $f(K^{n-1})$ are covered the same number of times. If this number is 1, then f , considered in the open subset $R' = f^{-1}(R)$ of K only, is bijective onto R .*

We recall some definitions and a lemma from [38] and [28].

Definition 10.6 (Strong convexity radius) A geodesic disk $B_r(x)$ centered at x on a Riemannian surface is called strongly convex if for any two point in the closed disk $\bar{B}_r(x)$ there is exists a unique minimizing geodesic connecting these points which lies in the interior of $B_r(x)$. The strong convexity radius is the largest radius for which the disk is strongly convex.

We have the following bound on the convexity radius [18, Theorem IX.6.1]:

Lemma 10.7 Suppose the sectional curvatures of M are bounded by $K \leq \Lambda_+$, and i_M is the injectivity radius. If

$$r < \min \left\{ \frac{i_M}{2}, \frac{\pi}{2\sqrt{\Lambda_+}} \right\},$$

then $\bar{B}_M(x; r)$ is convex. (If $\Lambda_+ \leq 0$, we take $1/\sqrt{\Lambda_+}$ to be infinite.)

Remark 10.8 If the vertices of an intrinsic simplex lie inside a convex disk, then the simplex is wholly contained in the convex disk.

Definition 10.9 (Pseudo-disks (Boissonnat and Oudot [8])) A family of topological disks on a surface are Pseudo-Disks if for any two distinct disks, their boundary either do not intersect, or they intersect tangentially at a single point, or they intersect transversely at exactly two points.

Lemma 10.10 (Small circle intersection ([28, 38])) All geodesic disks whose radius is strictly smaller than the strong convexity radius are pseudo-disks.

We assume that there are at least two intersection points of the two circles, and want to prove that there are not more than 2. For the other cases, we refer to [28]. This is precisely where the proof of [28] has a hole. Indeed, it is claimed that given two convex bodies C_1 and C_2 in dimension 2, the boundary of $\partial C_1 \cap \partial C_2$ consists of 2 connected arcs, one on ∂C_1 and one on ∂C_2 . This is not true as can be seen by taking two ellipses and rotating one by π .

Proof Let us assume that two geodesic circles of radius r and 1 intersect 2 times on one side of the extended geodesic between the two centers c_1 and c_2 . Denote the intersection points A and B . Two of the geodesics between A and B and c_1 and c_2 respectively must intersect in the point p , because of the uniqueness of geodesics. Let us say that that Ac_2 and Bc_1 intersect. Denote the lengths of the geodesics Ap : a_1 the length pc_2 : a_2 , the length Bp : b_1 and the length pc_1 : b_2 . By definition $a_1 + a_2 = 1$ and $b_1 + b_2 = r$, so $a_1 + a_2 + b_1 + b_2 = 1 + r$. Because p does not lie on the geodesic Bc_2 or Ac_1 , we have $a_1 + b_2 > 1$ and $a_2 + b_1 > r$ and $a_1 + a_2 + b_1 + b_2 > 1 + r$. The latter inequality yields a contradiction. \square




Figure 12: Illustrations for the proof of Lemma 10.10

We want to prove that the Riemannian realization of the Delaunay complex induces a homeomorphism from the complex to the surface.

Corollary 10.11 *Two Riemannian Delaunay triangles t_1 and t_2 lying within a disk of radius less than the convexity radius on a Riemannian surface, of which no three vertices lie on a geodesic, that share an edge intersect only in this edge.*

Proof Let v_1 and v_2 be the two vertices that are shared by $t_1 = \{v_1, v_2, v_3\}$ and $t_2 = \{v_1, v_2, v_4\}$ and denote by $\gamma_{v_1 v_2 E}$ the extended geodesic between v_1 and v_2 . If v_3 and v_4 lie on either side of $\gamma_{v_1 v_2 E}$, there is nothing to prove. If v_3 and v_4 lie on the same side of $\gamma_{v_1 v_2 E}$ (lets say the left), we consider geodesic circumscribing circles of t_1 and t_2 , the boundaries of the disks D_{v_3} and D_{v_4} . Lemma 10.10 implies that these circles only intersect in v_1 and v_2 . Moreover the arc given by the part of ∂D_{v_3} on the left side of $\gamma_{v_1 v_2 E}$ is contained in D_{v_4} or vice versa the arc given by the part of ∂D_{v_4} on the left side of $\gamma_{v_1 v_2 E}$ is contained in D_{v_3} . This contradicts the empty ball property of Delaunay ball, so this configuration can not occur. \square

Corollary 10.12 *Let t_1 and t_2 be two Riemannian Delaunay triangles lying within a disk of radius less than the strong convexity radius on a Riemannian surface. Suppose that no three vertices lie on a geodesic, and that share t_1 and t_2 a vertex. Then intersection of the interiors of t_1 and t_2 is empty.*

Proof Suppose not, then let us consider the Delaunay disks for both triangles. The Delaunay disks contain the triangles because of convexity. The Delaunay circles intersect in the common vertex. Because they must contain no vertices they must intersect in at least two more points, contradicting Lemma 10.10. \square




Figure 13: The Riemannian Delaunay triangles t_1 and t_2 and their respective Delaunay balls.

Remark 10.13 *Let us assume that p_i is a Delaunay vertex of a Delaunay complex, then p_i does not lie in the interior or on an edge of a Delaunay triangle, otherwise it would violate the empty ball property.*

Corollary 10.14 *Let p be a Delaunay vertex. And denote by e_j the Delaunay edges emanating from p . Then there is a neighborhood $p \in U$ such that: for every point q in $U \setminus (\cup_i e_i \cup p)$ lies in a single Delaunay triangle.*

Proof Because of Remark 10.13, there exists a neighborhood U of p such that all Delaunay triangles with non-empty intersection with U have p as a vertex. Corollary 10.12 implies that the pairwise intersections of the interiors of the simplices that have p as a vertex is empty. The statement now follows. \square

Remark 10.15 *Suppose that all Voronoi cells in a Voronoi diagram $V_{\{p_i\}}$ are contained in a geodesic disk of radius less than a sixth of the convexity radius. Then for any p_i the Delaunay simplices that have p_i as a vertex are contained in a geodesic disk of radius less than a third of the convexity radius.*

Lemma 10.16 *Voronoi cells in a convex disk are star shaped, in particular they are topological disks.*

Proof Let p be the center of the Voronoi cell, and b a point on the boundary. Then the geodesic between b and p is completely contained in the Voronoi cell. Suppose that it is not, then there is a point on the geodesic that is closer to another Voronoi center, say q , but then it follows that b is closer to q too. \square

Lemma 10.17 *Let $\mathcal{V}_{\{p_i\}}$ be the Riemannian Voronoi diagram of the point set $\{p_i\}$ on an oriented Riemannian surface M . Suppose that no three vertices lie on a geodesic. Assume that all Voronoi cells V_{p_i} lie within balls of radius less than a sixth of the convexity radius. Then every Delaunay edge is shared by at least two Delaunay triangles.*

Proof Let us suppose that V_p and V_q are the Voronoi cells that share a Voronoi face. Call $\gamma_{pq,E}$ the extended geodesic between p and q in a disk centered on p of radius less than the convexity radius. The geodesic $\gamma_{pq,E}$ can be subdivided into three parts, namely the geodesic γ_{pq} between p and q and the two connected components of $\gamma_{pq,E} \setminus \gamma_{pq} = \gamma_{pq,Ep} \cup \gamma_{pq,Eq}$, with $p \in \overline{\gamma_{pq,Ep}}$ and $q \in \overline{\gamma_{pq,Eq}}$. It is clear that every point in $\gamma_{pq,Ep}$ is closer to p than q . The limit we have placed on the size of the Voronoi cell now implies that the intersection of V_p and V_q consists of segments curves. Moreover every Delaunay edge lies in at least one Delaunay triangle. Now let us consider one of the segments of the curves. Suppose that the end points of this segment consist of intersection points between the same three Voronoi cells. We can now consider the Delaunay balls centered at the end points and we find two Delaunay balls intersecting in three points. So we find a contradiction with the pseudo-ball property. This means that every Delaunay edge is shared by at least two Delaunay triangles. \square

Lemma 10.18 *Under the same conditions as Lemma 10.17, the Delaunay complex is a topological two-manifold at every vertex.*

Proof Let us consider an arbitrary vertex p_i . The vertex p_i can not be disconnected, because by compactness there exists at least one vertex p_j for which the distance between p_i and the other vertices is minimized. Therefore p_i and p_j share a Delaunay edge. By Lemmas 10.17 and 10.11, this edge lies in two triangles whose interiors do not intersect. These triangles share no other edge, otherwise they would coincide. By induction over a finite (due to compactness) number of triangles, we find the star of the vertex. No overlaps can take place because of Lemma 10.12. \square

We now have the following:

Corollary 10.19 *Under the same conditions as Lemma 10.17, the Delaunay complex is a topological two manifold.*

With this result, we are now able to state the following:

Theorem 10.20 *Let $\mathcal{V}_{\{p_i\}}$ be the Riemannian Voronoi diagram of the point set $\{p_i\}$ on an oriented Riemannian surface M . Suppose that no three vertices lie on a geodesic. Assume that all Voronoi cells V_{p_i} lie within balls of radius less than a sixth of the convexity radius. Then the Riemannian realization of the Delaunay complex of the Riemannian Voronoi diagram induces a piecewise smooth homeomorphism from the Delaunay complex to the surface M .*

Proof We are going to apply Whitney's lemma (see Lemma 10.5). Because we assume that no three vertices lie on a geodesic, the geometric realizations of two dimensional simplices are non-degenerate, that is diffeomorphisms. Due to Corollary 10.11, we can assume that two triangles that share an edge have compatible orientation. This means that Riemannian realization of the Delaunay complex \mathcal{D} induces a simplex-wise positive map, which we shall call f .

Now let us intersect the surface at an arbitrary point with a geodesic disk of radius $1 - \epsilon$ times the convexity radius and remove all simplices that do not completely lie in the intersection. The resulting complex is a subcomplex of the Delaunay complex and has a boundary (the boundary may have more connected components and singular points). Remark 10.15 implies that the boundary on the surface lies at least a distance of $2/3 - \epsilon$ times the convexity radius from the center of the geodesic disk. By Lemma 10.5, we now have that with exception of the skeleton the number of covers is constant. By Corollary 10.14, it is one. This means that f is locally a piecewise smooth homeomorphism. Due to Remark 10.15 and because the point has been chosen arbitrarily, it follows that f is a global piecewise smooth homeomorphism. \square

11 Straight edge realization of the RVD

Starting from an embedded Riemannian Delaunay triangulation, we give conditions such that "straightening" the geodesic edges of the RDT preserves the property of being an embedded triangulation.

We first prove that, under some sampling conditions, the geodesic path and the straight edge between two seeds are close from one another. This result is then used to prove that straightening edges creates no inversion of simplices or self-intersections.

Lemma 11.1 (Distance bound between a geodesic path and a straight segment) *Let U , ψ_0 , g and g' be defined as in Lemma 3.2, but g' is now assumed to be constant. Let us further denote the geodesics with respect to g and g' connecting x and y as γ_g and $\gamma_{g'}$. If $z \in \gamma_g$ then $d_{g'}(z, \gamma_{g'}) \leq a$, where*

$$a = \frac{1}{2} \sqrt{\psi_0^2 - \frac{1}{\psi_0^2} d_g(x, y)}$$

Proof Denote by $\pi(z, \gamma_{g'})$ the point on $\gamma_{g'}$ that is closest to z . Note that since g' is uniform, this point is unique. Moreover, if $\pi(z, \gamma_{g'})$ does not equal x or y , one can use Pythagoras' theorem:

$$\begin{aligned} d_{g'}(x, z)^2 &= d_{g'}(x, \pi(z, \gamma_{g'}))^2 + d_{g'}(\pi(z, \gamma_{g'}), z)^2 \\ d_{g'}(z, y)^2 &= d_{g'}(z, \pi(z, \gamma_{g'}))^2 + d_{g'}(\pi(z, \gamma_{g'}), y)^2 \end{aligned}$$

On the other hand, we have by definition of the distance that the length $L_{g'}(x, z, y)$ of any curve from x to y via z is at least $d_{g'}(x, z) + d_{g'}(z, y)$.

Let $c \in \mathbb{R}$ such that $d_{g'}(\pi(z, \gamma_{g'}), x) = \frac{1}{2}d_{g'}(x, y) - c$ and $d_{g'}(\pi(z, \gamma_{g'}), y) = \frac{1}{2}d_{g'}(x, y) + c$. If $\pi(z, \gamma_{g'})$ does not equal x or y ,

$$L_{g'}(x, z, y) > \sqrt{a^2 + \left(\frac{1}{2}d_{g'}(x, y) - c\right)^2} + \sqrt{a^2 + \left(\frac{1}{2}d_{g'}(x, y) + c\right)^2}. \quad (8)$$

It is apparent from symmetry, but it can also easily be verified by hand or Mathematica, that the right hand side of Inequality (8) is minimized with respect to c when $c = 0$, thus

$$L_{g'}(x, z, y) > 2\sqrt{a^2 + \left(\frac{1}{2}d_{g'}(x, y)\right)^2}.$$

Using the metric distortion as in Lemma 3.2, we see that

$$L_g(x, z, y) > \frac{2}{\psi_0} \sqrt{a^2 + \left(\frac{1}{2}d_{g'}(x, y)\right)^2},$$

for all curves that go through z . In particular, we assumed that $z \in \gamma_g$, therefore

$$\begin{aligned} d_g(x, y) &> \frac{2}{\psi_0} \sqrt{a^2 + \left(\frac{1}{2}d_{g'}(x, y)\right)^2} \\ \iff \frac{\psi_0^2}{4} d_g(x, y)^2 &> a^2 + \frac{1}{4} d_{g'}(x, y)^2 \geq a^2 + \frac{1}{4\psi_0^2} d_g(x, y)^2. \end{aligned}$$

Finally, bringing the metric terms to one side, we obtain

$$\frac{1}{4} \left(\psi_0^2 - \frac{1}{\psi_0^2} \right) d_g(x, y)^2 > a^2.$$

We know that if $\pi(z, \gamma_{g'})$ is x or y , then $L_{g'}(x, z, y) \geq a + \sqrt{d_{g'}(x, y)^2 + a^2}$. This can be easily seen as we are in (rescaled) Euclidean space, and $\pi(z, \gamma_{g'})$ is x or y , z lies on one of the two semi-circles around x and y . It now follows that $L_{g'}(x, z, y) \geq L_{g'}(x, \tilde{z}, y)$ where \tilde{z} is one of the end points of the semi circle. By Pythagoras' theorem, $L_{g'}(x, \tilde{z}, y) = a + \sqrt{d_{g'}(x, y)^2 + a^2}$. \square

Lemma 11.2 (Embedding of the SRDT) *If the RDT is an embedded triangulation and if the sampling criterion ϵ is small enough, then the SRDT is also an embedded triangulation.*

Proof Consider a Riemannian triangle σ of the RDT with vertices P, Q and R . By Lemma 4.1, we have bounds on the value θ_v of the angle at a Voronoi corner:

$$2 \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right) \leq \theta_v \leq \pi - \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right).$$

We can thus deduce bounds on the triangle angles θ_t :

$$\theta_m = \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right) \leq \theta_t \leq \pi - 2 \arcsin\left(\frac{\mu_0}{2\epsilon_0}\right) = \theta_M.$$

Using the definition of a given in Lemma 11.1, we have at a corner of the triangle Δpqr

$$\sin(\theta) = \frac{a}{l} \iff l = \frac{a}{\sin(\theta)}$$

We need to find a lower bound on the length of the straight edge PQ eroded by the protecting tubes of the edges PQ and PR (see Figure 14).

$$\begin{aligned}
l = L - (l_R + l_Q) &\geq \mu_0 - \left(\frac{a_R}{\sin(\theta_R)} + \frac{a_Q}{\sin(\theta_Q)} \right) \\
&\geq \mu_0 - \frac{\sqrt{\psi_0^2 - \frac{1}{\psi_0^2}} d_{g'}(p, q)}{2 \sin(\theta_R)} - \frac{\sqrt{\psi_0^2 - \frac{1}{\psi_0^2}} d_{g'}(p, r)}{2 \sin(\theta_Q)} \\
&\geq \mu_0 - \frac{\epsilon_0 \sqrt{\psi_0^2 - \frac{1}{\psi_0^2}}}{\sin\left(\arcsin\left(\frac{\mu_0}{2\epsilon_0}\right)\right)} \\
&\geq \mu_0 - \frac{2\epsilon_0^2 \sqrt{\psi_0^2 - \frac{1}{\psi_0^2}}}{\mu_0}.
\end{aligned}$$




Figure 14: The Riemannian Delaunay triangulation is shown in red and the AGDT in black. The geodesic edges between seeds P-R and P-Q lie in their respective protection tubes (colored in beige) around the straight edges.

It is now sufficient to show that the distance l on Figure 14 is strictly positive. Indeed, if we let the metric tend to the constant metric the geodesic triangle tends to a straight triangle. This is a homotopy and the orientation of the triangles remains constant if l is strictly positive.

We require $l > 0$, that is

$$\begin{aligned} l \geq \mu_0 - \frac{2\epsilon_0^2 \sqrt{\psi_0^2 - \frac{1}{\psi_0^2}}}{\mu_0} > 0 &\iff \mu_0^2 > 2\epsilon_0^2 \sqrt{\psi_0^2 - \frac{1}{\psi_0^2}} \\ &\iff \frac{\lambda^2}{\psi_0^2} \epsilon^2 > 2\psi_0^2 \epsilon^2 \sqrt{\psi_0^2 - \frac{1}{\psi_0^2}} \\ &\iff \lambda^2 > 2\psi_0^4 \sqrt{\psi_0^2 - \frac{1}{\psi_0^2}}, \end{aligned}$$

using $\epsilon_0 = \psi_0 \epsilon$, $\mu_0 = \mu / \psi_0$ and $\mu = \lambda \epsilon$. The left hand side is fixed, the right hand side goes to 0 as ϵ goes to 0 (and ψ_0 goes to 1).

This proves that no issue can appear between the two edges PQ and PR when straightened. This reasoning can also be applied to the other pairs of edges in ΔPQR , and finally to all the simplices of the RDT, proving that the SRDT is also an embedded triangulation of the domain. \square

12 Extension to surfaces

The previous results may be generalized to surfaces embedded in \mathbb{R}^3 . We shall assume that, apart from the metric induced by the embedding of the domain in Euclidean space, there is a second metric g defined on M . Let $\pi_p : M \rightarrow T_p M$ be the orthogonal projection of points of M on the tangent space $T_p M$ at p . π_p is a local diffeomorphism for a sufficiently small neighborhood $U_p \subset T_p M$, see [45]. The pullback of the metric g with the inverse projection π_p^{-1} defines a metric on U_p . We have a metric on a subset of a two dimensional space, in this case the tangent space, giving us a setting that we have already solved. Boissonnat et al. [4, Lemma 3.7] give bounds on the metric distortion of the projection on the tangent space. This result allows to translate the δ -power protected (ϵ, μ) -net properties between the tangent space and the manifold.

13 Constructing the point set \mathcal{P}

Our theory introduces requirements on the canvas \mathcal{C} (through its sizing field) and on the point set \mathcal{P} (through its power protected net nature) such that the two duals of the DRVD are embedded triangulations. The former is easily achieved through refinement. We now explain how a δ -power protected ϵ -sample can be built iteratively through a refinement algorithm that is similar to that of the locally uniform Delaunay method proposed by Boissonnat et al. [10].

In the local Delaunay algorithm, the obstacle that potentially prevents the star set from being a triangulation is the presence of inconsistencies: a configuration where adjacent stars whose connectivity is incompatible. It is shown that, once the distortion is small enough, inconsistencies are only created by *quasi-cosphericities*, the existence $n + 1$ points almost living on a sphere. Steiner points can then be chosen in a small zone around the circumcenter called the *picking region* such that no quasi-cosphericities are created (see Figure 15) as the thickened Delaunay spheres encroaching on the picking region do not cover it.

Power protection and the absence of quasi-cosphericities are two aspects of the same concept. Indeed, both can be thought of as thickened circumscribing balls that must be empty of any point other than the vertices of the circumscribed simplex. In Boissonnat et al.'s setting, this thickening depends on the metric distortion and goes to 0 as the distortion becomes small, which




Figure 15: The picking region of the simplex τ and a forbidden zone created by three neighboring vertices, [9].

makes the core of their proof of the (eventual) existence of valid Steiner points. In our context, the thickening is explicitly given by the δ coefficient of the power protection and must be bigger than a given value δ_0 . However, this bound δ_0 also goes to 0 as the distortion becomes small, which happens as the sampling gets denser. The same smart refinement algorithm that prevents quasi-cosphericity in their method will therefore create here a power protected point set.

The initialization of the point set \mathcal{P} is made with a few random points on the domain. The refinement algorithm is made of two criteria:

- A sizing criterion to obtain an ϵ -sample. Simplices are refined by the insertion of their circumcenter.
- A protection criterion to obtain a δ -power protected point set. The Steiner point is chosen with a `pick_valid` routine that selects a good refinement point (see [11] for details). If no good refinement point exists (yet) in the picking region, the circumcenter is instead inserted, which lowers the distortion and increases the probability of finding a valid Steiner point at a further step.

The μ -separation of the point set is implicitly obtained with this choice of Steiner points and can be computed explicitly through the insertion radius, as done in [11]. In similar fashion to the local Delaunay algorithm, the refinement algorithm can be shown to terminate and to produce a δ -power protected (ϵ, μ) -net.

In practice, neither the sampling nor the canvas need to be as dense as our theory requires them to be and a farthest point refinement can be used to generate the point set.

14 Results

We now present some results for 2D and planar domains. The different geodesic distance computation algorithms described in Section 6.1 were implemented and tested and while the method proposed by Campen et al.'s does not offer any theoretical guarantee, its precision and convergence were in practice satisfactory and its speed a significant advantage. CGAL [17] was used to generate isotropic Delaunay canvases and our newer implementation of the star set algorithm for star sets.

We consider two types of metric fields: artificial ones, chosen for their difficulty, and curvature-based metric fields as more realistic examples.




Figure 16: On the left, the DRVD of 4010 seeds in a square endowed with a swirly metric field. On the right, the straight edge dual triangulation of the DRVD.

Computation time

The extraction of the nerve is performed on the fly during the coloring the canvas. This operation and the construction of the duals are negligible in cost compared to computing geodesic distances.

Figure 16 uses an isotropic canvas of a 4 by 4 square and a swirly metric field. The canvas is composed of 1,286,862 vertices and took 29 seconds to generate. The computation of the RDT and SRDT took almost 6 hours. Comparatively, computing the RDT and SRDT using a star set as canvas only required 121,264 stars (vertices) and took 6 minutes, including 3 minutes to generate the star set. This improvement becomes even more significant as the anisotropy increases.

Duals

The swirly and hyperbolic shock metric field (see Figures 16 and 21) are good examples of metric fields that are typically difficult for the non-geodesic Voronoi diagrams of Du and Wang and Labelle and Shewchuk. Indeed, they are characterized by both noticeable anisotropy ratios and a highly rotational vector field (vectors given at a point by the eigenvectors of the metric at this point). These regions are easily populated by orphan cells and an improper dual for the VS or LS diagram if the point set is not sufficiently dense. On the other hand, this issue cannot appear in the RVD and both the RDT and the SRDT quickly become triangulations.

Preprocessing the metric field, typically with a Laplacian smoothing, is a technique also often used (see for example [30]) that lowers the fidelity to the input metric field and can even nullify the anisotropy such as in the highly rotational region of the hyperbolic shock, resulting in isotropic elements. On the other hand, our algorithm produces a DRVD that honors the metric field closely and produces triangulations whose anisotropic elements conform to the input metric field with even very few points for both the RDT and the SRDT (see Figure 21).

Geodesic edges The geodesic paths are traced by backtracking along the gradient of the geodesic distance map, as described by Yoo et al. [55]. In the case of planar and surface domains, tracing the geodesic paths between the seeds provides the complete triangulation as the Riemannian simplices are simply the interior of the region formed by the three geodesic edges on the surface. Figure 17 shows the DRVD and the RDT for the “chair” surface endowed with a curvature-induced anisotropic metric field. The anisotropy of the regions near parabolic lines is well captured with few Riemannian simplices.



Figure 17: the DRVD of 1020 seeds on the “chair” surface, with a curvature induced metric field. The RDT is traced in black.

Straight edges By definition, the RVD captures the metric field more accurately than other methods that typically only consider the metric at the vertices. This allows us to construct Riemannian triangulations, but also has a positive influence on the straight edge realization of the DRVD. Figure 18 and 19 shows the SRDT of two surfaces endowed with the hyperbolic shock metric field and a curvature metric field. Both SRDT rapidly become embedded triangulations during the refinement process, as illustrated on the leftmost example.

Edge lengths and angles The quality of a mesh and its conforming to the metric field can be evaluated through the measure, in the metric, of the lengths of its edges and of its angles. We evaluate the edge lengths and the angles of the point set shown on Figure 21 (the SRDT is shown) for both the RDT and the SRDT. Note that the point sets are generated with a sizing criterion intended to create unit meshes, that is RDTs whose geodesic edges are ideally of length 1. Results for the edge lengths of the RDT and the angles of the SRDT are shown in Figure 20 and show good control on the edge lengths and the angles. We also observed that while the results are naturally better for the RDT, since the point set is generated with respect to the geodesic distance, the quality of the elements of the SRDT stays relatively close to that of the RDT.




Figure 18: The unit sphere endowed with the hyperbolic metric field. DRVD (left) and SRDT (right).

Optimization

Optimization is often used to improve the quality of a triangulation. Centroidal Voronoi tessellations are Voronoi diagrams whose generators are also the centroids (centers of mass) of their respective cells [40]. The famous Lloyd algorithm [41] iteratively moves the seeds to the center of mass of their respective cell and recomputes the Voronoi diagram of this new seed set.

We approximate the Riemannian Voronoi center of mass of a cell V with the following formula:

$$c_g = \frac{\sum_i c_i |t_i|_{g_i}}{\sum_i |t_i|_{g_i}}. \quad (9)$$

where $|t_i|_{g_i}$ is the area of the triangle t_i in the metric g_i and the t_i make a partition of V . The canvas conveniently provides this decomposition of a geodesic Voronoi cell in small triangles, making the approximation of c_g accurate. The formula in equation (9) does not extend to surfaces as the result of the weighted sum might not lie on the domain. In that setting, we use a process similar to Wang et al. [52]. As its Euclidean counterpart, this algorithm comes with no guarantees (not even for termination) but works well in practice. In Figure 21, the initial SRDT of 320 seeds has been optimized with 100 iterations. The metric field is well captured, especially in the rotational region, with few elements.

15 Metric fields

We detail below the different metric fields used in our experiments.

15.1 Curvature

Let S be an implicit surface defined by $f(x, y, z) = 0$. At each point p of S , we denote by N the normal vector $-\nabla f / \|\nabla f\|$ and by H the Hessian of f . The principal curvatures c_{\max} , c_{\min} are the eigenvalues of the matrix

$$C = P \cdot H' \cdot P,$$

$$\text{where } H' = H \|\nabla f\| \text{ and } P = I_3 - N \cdot N^t,$$

see e.g. [32]. By convention, we define c_{\max} to be the eigenvalue with the maximal absolute value. Let U_{\max} , U_{\min} and N be the normalized eigenvectors of C , and let $U = [U_{\max}, U_{\min}, N]$. The




Figure 19: the “Fertility” surface endowed with a curvature induced metric field, SRDT with 220 seeds (left), DRVD (center) and SRDT (right) with 6020 seeds.




Figure 20: Histograms for the optimized mesh shown on Figure 21. On the left, RDT edge lengths; on the right, SRDT angles values.

metric at the point p is defined by $M_p = F_p^t F_p$, where:

$$\begin{aligned}
 F_p &= U \cdot \Delta \cdot U^t, \\
 \text{with } \Delta &= \text{diag}\{e_{\max}, e_{\min}, e_n\}, \\
 e_{\max} &= \sqrt{|c_{\max}|}, \\
 e_{\min} &= \sqrt{|c_{\min}|}.
 \end{aligned}$$

If one of the curvature values vanishes, we replace it with a small strictly positive value ϵ such that the metric stays positive definite.

15.2 Hyperbolic shock

If E denotes such a scalar field defined in \mathbb{R}^3 and S a surface, we may want to mesh S to provide the best approximation of E on S . The triangles are elongated orthogonally to the direction




Figure 21: The optimized SRDT of 4000 seeds in a planar domain endowed with a hyperbolic shock induced metric field (left). On the right, a zoom on a rotational region of the metric field shows the difference between pre- (above) and post- (bottom) optimization.

of the gradient ∇E and the density depends both on the norm and the changing rate of the gradient.

We define the metric field as follows :

$$F_M = U \cdot \text{diag}\{1/(1 + \varphi \|\nabla E\|), 1, 1\} \cdot U^t$$

where $U = [\frac{\nabla E}{\|\nabla E\|}, U_1, U_2]$, U_1 and U_2 being two arbitrary unit vectors that form an orthogonal basis of the orthogonal complement of ∇E . If the gradient is zero, the metric field is isotropic. Parameter φ controls how the facets are stretched with respect to the norm of the gradient.

The hyperbolic shock is an arbitrary sine-shaped shock function that can be described by the following equation:

$$E(x, y, z) = \tanh\left(\frac{1}{\lambda}(2x - \sin(5y))\right) + x^3 + xy^2,$$

with $\lambda = 0.6$. The metric field is derived from E as described above.

15.3 Swirl

The Swirl metric field is based on a uniform vector field that creates a swirl to which an anisotropy ratio that depends on the distance to the center is attached. The metric field has a parameter ν that controls the circular motion; we use $\nu = 3$. The metric M at a point $P(x, y)$ is constructed as follows.

First, construct the eigenvectors v_1 and v_2 of the matrix:

$$\begin{aligned} v_{1,x} &= \frac{-(xs + yc) + (xc - ys)(1 - x^2 - y^2)^2}{r^2} \\ v_{1,y} &= \frac{(xc - ys) + (xs + yc)(1 - x^2 - y^2)^2}{r^2} \\ n &= \sqrt{v_{1,x}^2 + v_{1,y}^2} \\ v1 &= \left(\frac{v_{1,x}}{n}, \frac{v_{1,y}}{n} \right) \quad v2 = \left(\frac{-v_{1,y}}{n}, \frac{v_{1,x}}{n} \right) \end{aligned}$$

where $c = \cos(\nu)$ and $s = \sin(\nu)$.

Then construct the eigenvalues λ_1 and λ_2 (responsible for the anisotropy ratio):

$$\begin{aligned} f &= \frac{\sqrt[3]{r} + \sqrt{r}}{2} \\ \lambda_1 &= \frac{1}{|50(1 - (f - 1)^2)|} \quad \lambda_2 = 1 \end{aligned}$$

where $r^2 = x^2 + y^2$.

We can now assemble the matrix :

$$\begin{aligned} U &= (v1, v2) \\ \Delta &= \text{Diag}(\lambda_1, \lambda_2) \\ F &= U^t \Delta U \end{aligned}$$

And finally $M = F^t F$.

16 Conclusion

We have introduced a practical approach to computing Riemannian Voronoi diagrams and exposed theoretical conditions on the point set and the canvas such that geodesic and straight edge realizations are embedded triangulations of the domain in planar and surface cases. The theoretical requirements are demanding, but in practice both the geodesic and straight edge duals require relatively few points to become embedded triangulations, even with non-trivial, highly distorted metric fields. The RVD and its duals are shown to be particularly well suited to capture the metric field in regions where it is both anisotropic and curved. No preprocessing or smoothing of the metric field is required, a technique that is often used and results in noticeable loss of anisotropy in this type of region.

Future work

The use of the star set greatly increases the speed of the algorithm, and additional techniques are still to be explored. A first technique to improve the computation time is to observe that during most of the computation of the DRVD, there is very little interaction between the different fronts. Parallelization of the algorithm should therefore yield significant gains. If the metric we consider were almost constant in significant parts of our domain, a hybrid diagram that would only use geodesics in zones where the metric varies a lot would decrease the running time significantly.

Another interesting topic to explore would be the applied uses of Riemannian Delaunay triangulations. Indeed, much work has been committed to the theoretical study of these triangulations, but it still stays untouched in practice. Finite elements can be for example defined on the RDT in our context as the canvas naturally offers a piecewise linear transformation between a simplex of the RDT and the traditional reference element of the finite element method. One can also intuit that Riemannian simplices adapted to a metric field could be used to further decrease numerical error.

Finally, the work presented here focused on planar and surface domains. A future paper will provide results for Riemannian manifolds of arbitrary dimension.

Acknowledgments

The second and third authors have received funding from the European Research Council under the European Union's ERC Grant Agreement number 339025 GUDHI (Algorithmic Foundations of Geometric Understanding in Higher Dimensions).

References

- [1] Pierre Alliez, David Cohen-Steiner, Mariette Yvinec, and Mathieu Desbrun. Variational tetrahedral meshing. *ACM Trans. Graph.*, 24(3):617–625, 2005.
- [2] Franz Aurenhammer and Rolf Klein. Voronoi diagrams. In J. Sack and G. Urrutia, editors, *Handbook of Computational Geometry*, pages 201–290. Elsevier Science Publishing, 2000.
- [3] Jean-Daniel Boissonnat, David Cohen-Steiner, and Mariette Yvinec. Comparison of algorithms for anisotropic meshing and adaptive refinement. Research Report ACS-TR-362603, INRIA, 2008.
- [4] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Constructing intrinsic Delaunay triangulations of submanifolds. Research Report RR-8273, INRIA, 2013. arXiv:1303.6493.
- [5] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. The stability of Delaunay triangulations. *Int. J. Comp. Geom. & Appl.*, 23(04n05):303–333, 2013. eprint: arXiv:1304.2947.
- [6] Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh. Delaunay stability via perturbations. *Int. J. Comp. Geom. & Appl.*, to appear, 2014. (Preprint: arXiv:1310.7696).
- [7] Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh, and Steve Y. Oudot. Only distances are required to reconstruct submanifolds. Research report, INRIA Sophia Antipolis, 2014.
- [8] Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, 2005.
- [9] Jean-Daniel Boissonnat, Kan-Le Shi, Jane Tournois, and Mariette Yvinec. Anisotropic delaunay meshes of surfaces. *ACM Trans. Graph.*, 34(2):14:1–14:11, 2015.
- [10] Jean-Daniel Boissonnat, Camille Wormser, and Mariette Yvinec. Locally uniform anisotropic meshing. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 270–277. ACM, 2008.
- [11] Jean-Daniel Boissonnat, Camille Wormser, and Mariette Yvinec. Anisotropic Delaunay mesh generation. *SIAM Journal on Computing*, 44(2):467–512, 2015.

-
- [12] Houman Borouchaki, Paul Louis George, Frédéric Hecht, Patrick Laug, and Eric Saltel. Delaunay mesh generation governed by metric specifications. part i. algorithms. *Finite elements in analysis and design*, 25(1):61–83, 1997.
- [13] Marcel Campen, Martin Heistermann, and Leif Kobbelt. Practical anisotropic geodesy. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing*, SGP '13, pages 63–71. Eurographics Association, 2013.
- [14] Guillermo D. Canas and Steven J. Gortler. Orphan-free anisotropic Voronoi diagrams. *Discrete and Computational Geometry*, 46(3), 2011.
- [15] Guillermo D. Cañas and Steven J. Gortler. Duals of orphan-free anisotropic Voronoi diagrams are embedded meshes. In *SoCG*, pages 219–228, New York, NY, USA, 2012. ACM.
- [16] Thanh-tung Cao, Herbert Edelsbrunner, and Tiow-seng Tan. Proof of correctness of the digital delaunay triangulation algorithm. *Computational Geometry: Theory and Applications*, 48, 2015.
- [17] CGAL. *User and Reference Manual*. CGAL Editorial Board, 4.7 edition, 2015.
- [18] Isaac Chavel. *Riemannian Geometry, A modern introduction*. Cambridge, 2nd edition, 2006.
- [19] Long Chen. Mesh smoothing schemes based on optimal Delaunay triangulations. In *Proceedings of the 13th International Meshing Roundtable*, pages 109–120. Sandia National Laboratories, 2004.
- [20] Long Chen, Pengtao Sun, and Jinchao Xu. Optimal anisotropic meshes for minimizing interpolation errors in L^p -norm. *Mathematics of Computation*, 76:179–204, January 2007.
- [21] Long Chen and Jinchao Xu. Optimal Delaunay triangulations. *Journal of Computational Mathematics*, 22:299–308, 2004.
- [22] Siu-Wing Cheng, Tamal K Dey, Edgar A Ramos, and Rephael Wenger. Anisotropic surface meshing. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 202–211. Society for Industrial and Applied Mathematics, 2006.
- [23] Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. Graph.*, 32(5):152:1–152:11, 2013.
- [24] E. F. D’Azevedo and R. B. Simpson. On optimal interpolation triangle incidences. *SIAM J. Sci. Statist. Comput.*, 10(6):1063–1075, 1989.
- [25] Qiang Du and Desheng Wang. Anisotropic centroidal Voronoi tessellations and their applications. *SIAM Journal on Scientific Computing*, 26(3):737–761, 2005.
- [26] Ramsay Dyer, Gert Vegter, and Mathijs Wintraecken. Riemannian simplices and triangulations. 2014. Preprint: arXiv:1406.3740.
- [27] Ramsay Dyer, Hao Zhang, and Torstan Möller. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum (Special Issue of Symp. Geometry Processing)*, 27(5):1393–1402, 2008.
- [28] Ramsay Dyer, Hao Zhang, Torsten Möller, and Andrew Clements. An investigation of the spectral robustness of mesh Laplacians. Technical Report TR 2007-17, Simon Fraser University, 2007. SFU-CMPT.

-
- [29] Herbert Edelsbrunner and Nimish R. Shah. Triangulating topological spaces. *Int. J. Comput. Geometry Appl.*, 7(4):365–378, 1997.
- [30] Xiao-Ming Fu, Yang Liu, John Snyder, and Baining Guo. Anisotropic simplicial meshing using local convex functions. *ACM Trans. Graph.*, 33(6):182:1–182:11, 2014.
- [31] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *ACM SIGGRAPH*, pages 209–216, 1997.
- [32] John F Hughes. Differential geometry of implicit surfaces in 3-space-a primer. Research report, Brown University Providence, Rhode Island, 2003.
- [33] Hermann Karcher. Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*, 30:509–541, 1977.
- [34] Wilfrid S. Kendall. Probability, convexity, and harmonic maps with small image I: Uniqueness and fine existence. *Proceedings of the London Mathematical society*, s3-61 (Issue 2):371–406, 1990.
- [35] Ender Konukoglu, Maxime Sermesant, Olivier Clatz, Jean-Marc Peyrat, Hervé Delingette, and Nicholas Ayache. A recursive anisotropic fast marching approach to reaction diffusion equation: Application to tumor growth modeling. In *Proceedings of the 20th International Conference*, pages 687–699, 2007.
- [36] Francois Labelle and Jonathan Richard Shewchuk. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SoCG*, pages 191–200, 2003.
- [37] Greg Leibon and David Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *SoCG*, pages 341–349, 2000.
- [38] Gregory Leibon. *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization*. PhD thesis, UCSD, 1999. arXiv:math/0011016v1.
- [39] Yang Liu, Hao Pan, John Snyder, Wenping Wang, and Baining Guo. Computing self-supporting surfaces by regular triangulation. *ACM Trans. Graph.*, 32(4):92:1–92:10, 2013.
- [40] Yang Liu, Wenping Wang, Bruno Lévy, Feng Sun, Dong-Ming Yan, Lin Lu, and Chenglei Yang. On centroidal voronoi tessellation—energy smoothness and fast computation. *ACM Trans. Graph.*, 28(4):101:1–101:17, 2009.
- [41] Stuart Lloyd. Least squares quantization in pcm. *IEEE Trans. Inf. Theor.*, 28(2):129–137, 2006.
- [42] Jean-Marie Mirebeau. Optimal meshes for finite elements of arbitrary order. *Constructive approximation*, 32(2):339–383, 2010.
- [43] Jean-Marie Mirebeau. Anisotropic fast-marching on cartesian grids using lattice basis reduction. *SIAM Journal on Numerical Analysis*, 52(4), 2014.
- [44] Patrick Mullen, Pooran Memari, Fernando de Goes, and Mathieu Desbrun. Hot: Hodge-optimized triangulations. *ACM Trans. Graph.*, 30(4):103:1–103:12, 2011.
- [45] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Comp. Geom.*, 39(1-3):419–441, 2008.

-
- [46] G. Peyré and L. Cohen. Geodesic remeshing using front propagation. *International Journal on Computer Vision*, 69(1):145–156, 2006.
- [47] Gabriel Peyré, Mickael Péchaud, Renaud Keriven, and Laurent D. Cohen. Geodesic methods in computer vision and graphics. *Found. Trends. Comput. Graph. Vis.*, 5:197–397, 2010.
- [48] Raif M. Rustamov. Barycentric coordinates on surfaces. *Eurographics Symposium on Geometry Processing*, 29(5), 2010.
- [49] Jonathan R. Shewchuk. What is a good linear finite element? - interpolation, conditioning, anisotropy, and quality measures. <http://www.cs.berkeley.edu/~jrs/papers/elemj.pdf> (last viewed 2008), 2002.
- [50] Kenji Shimada, Atushi Yamada, and Takayuki Itoh. Anisotropic triangulation of parametric surfaces via close packing of ellipsoids. *International Journal of Computational Geometry & Applications*, 200.
- [51] Jane Tournois, Camille Wormser, Pierre Alliez, and Mathieu Desbrun. Interleaving delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. *ACM Trans. Graph.*, 28(3):75:1–75:9, 2009.
- [52] Xiaoning Wang, Xiang Ying, Yong-Jin Liu, Shi-Qing Xin, Wenping Wang, Xianfeng Gu, Wolfgang Mueller-Wittig, and Ying He. Intrinsic computation of centroidal Voronoi tessellation (CVT) on meshes. *Computer-Aided Design*, 58:51–61, 2015.
- [53] Hassler Whitney. *Geometric Integration Theory*. Princeton University Press, 1957.
- [54] Mathijs H.M.J. Wintraecken. *Ambient and Intrinsic Triangulations and Topological Methods in Cosmology*. PhD thesis, Rijksuniversiteit Groningen, 9 2015.
- [55] Sang Wook Yoo, Joon-Kyung Seong, Min-Hyuk Sung, Sung Yong Shin, and Elaine Cohen. A triangulation-invariant method for anisotropic geodesic map computation on surface meshes. *Visualization and Computer Graphics, IEEE Transactions on*, 18(10):1664–1677, 2012.
- [56] Zichun Zhong, Xiaohu Guo, Wenping Wang, Bruno Lévy, Feng Sun, Yang Liu, and Weihua Mao. Particle-based anisotropic surface meshing. *ACM Trans. Graph.*, 32(4):99:1–99:14, 2013.
- [57] Zichun Zhong, Liang Shuai, Miao Jin, and Xiaohu Guo. Anisotropic surface meshing with conformal embedding. *Graph. Models*, 76(5):468–483, 2014.

Contents

1	Introduction	3
1.1	Contribution	3
2	Related work	4
3	Background	5
3.1	Nets	5
3.2	Riemannian metric	5
3.3	Geodesics	6
3.3.1	Riemannian Voronoi diagrams	7
3.4	Protection and power-protection	7
4	Properties of power protected nets in Euclidean space	7
4.1	Voronoi angle bounds	8
4.2	Separation of the Voronoi vertices	9
5	Stability of the power protected net with respect to metric perturbations	11
5.1	Stability of the power protected net property with respect to metric transformation	11
5.2	Stability of the net property under metric perturbation	11
5.3	Stability of the protection property under metric perturbation	12
6	The discretized Riemannian Voronoi diagram	18
6.1	Computing geodesics	18
6.2	Canvas generation	18
6.3	Generation of the diagram	18
6.4	Extracting the nerve of the DRVD	19
7	Fast DRVD computations	19
7.1	The local Delaunay algorithm	20
7.2	Using star sets as canvasses	20
8	Overview of the theoretical sections	21
9	Equating the nerves of the DRVD and the RVD	22
9.1	Equating the nerves of the DRVD and the RVD: Euclidean metric field	22
9.2	Equating the nerves of the DRVD and the RVD: uniform metric field	26
9.3	Equating the nerves of the DRVD and the RVD: arbitrary metric field.	27
9.4	Approximate geodesic computations	30
10	Riemannian Delaunay triangulation	30
10.1	Riemannian simplices	30
10.2	2-manifolds	31
11	Straight edge realization of the RVD	35
12	Extension to surfaces	38
13	Constructing the point set \mathcal{P}	38

14 Results	39
15 Metric fields	42
15.1 Curvature	42
15.2 Hyperbolic shock	43
15.3 Swirl	44
16 Conclusion	45



**RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399