Lines: 491

Shape Reconstruction in Information Space

Herbert Edelsbrunner
IST Austria (Institute of Science and Technology Austria)
Am Campus 1, 3400 Klosterneuburg, Austria
edels@ist.ac.at

Katharina Ölsböck
IST Austria (Institute of Science and Technology Austria)
Am Campus 1, 3400 Klosterneuburg, Austria
katharina.oelsboeck@ist.ac.at

Hubert Wagner
IST Austria (Institute of Science and Technology Austria)
Am Campus 1, 3400 Klosterneuburg, Austria
hwagner@ist.ac.at

Abstract

The reconstruction of shape from a point sample is inherently sensitive to the interplay between local point configurations and ambient metric. Applying this viewpoint, we generalize popular Euclidean shape reconstruction methods to Bregman divergences and beyond. We focus on the Alpha and Wrap complexes in the context of the relative entropy and the Fisher metric.

The interest of this work is twofold. First, we use the generalized reconstruction methods, along with persistent homology, to experimentally compare these geometries. Second, the techniques and software we developed are of independent interest. One highlight is that the existing implementations for the Euclidean metric can be reused–although indirectly–in this generalized context. This removes a major roadblock for the development of topological data analysis tools working in non-Euclidean spaces.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases meshing, alpha shapes, wrap complex, persistent homology, non-Euclidean geometries, Bregman geometry, relative entropy, information space

1 Introduction

The motivation for the work reported in this paper is the deeper understanding of the role of the ambient metric in the reconstruction of shape. Specifically, we further generalize geometric and topological data analysis methods from Euclidean geometry to Bregman geometries in which dissimilarity is measured with divergences. By necessity, these methods are sensitive to the dissimilarity defining the ambient geometry, and we exploit this sensitivity to quantify the difference between geometries.

As example geometries, we emphasize those related to information theoretic concepts, such as the *Shannon geometry* and the *Fisher geometry*, in which dissimilarities are defined as the relative entropy (Kullback–Leibler divergence) and the Fisher distance, respectively. These are examples of what we like to call *information spaces*[11].

These geometries are commonly used in data analysis, and we hope this work sheds some light on the differences and commonalities between them. Some particularly pertinent questions are these: Is the Fisher geometry a good approximation of the Shannon geometry? Can we see a significant difference between the Euclidean geometry and the non-Euclidean ones, as predicted by the discrepancy in their practical performances?

We are also interested in the algorithms that underpin the data analysis methods, especially the topological ones. While the Fisher geometry can be handled with Euclidean tools [11], the Shannon geometry used to require customized tools [12]. We show that the

Shannon geometry can also benefit from existing, robust tools, although in this case the application is less direct. We also hope that this development opens new alleys for topological data analysis in information spaces.

Prior work and results. The research reported in this paper is based on two related lines of work, namely the study of Bregman divergences from the computational geometry point of view initiated in [5] and the extension of topological data analysis to Bregman and Fisher geometries started in [12], and [11] respectively.

An important concept in our investigations is the Bregman–Delaunay mosaic, which we formally define as the straight-line dual of the not necessarily straight-line Bregman–Voronoi tessellation obtained by measuring distance with the Bregman divergence from a data point. This mosaic was already defined in [5], and we explain how it can be computed as a weighted Euclidean Delaunay mosaic using standard geometric software. In Euclidean space, the Alpha shapes can be defined as the sublevel sets of the radius function on the Delaunay mosaic, which is a generalized discrete Morse function in the sense of Forman [13] and Freij [14]. As described in [3], the lower sets of the critical simplices of this radius function constitute the Wrap complex, which was introduced as a shape reconstruction tool in [8]. We extend this framework by introducing the rise function on a Bregman–Delaunay mosaic, which provides a convenient measure of the size of a Bregman ball. With these notions, we construct the shape of data in different geometries, and we use them to quantify the difference between the geometries.

We have implemented all the algorithms and use the software to run experiments, comparing Euclidean, Shannon, and Fisher geometries for synthetic data. We find that the Delaunay mosaics and their Alpha and Wrap complexes in these geometries show some but occasionally subtle differences, which we quantify.

Outline. Section 2 provides the necessary background from discrete geometry and combinatorial topology. Section 3 gives the details needed to compute Delaunay mosaics and their Alpha and Wrap complexes in Bregman and Fisher geometries using software for weighted Delaunay mosaics in Euclidean geometry. Section 4 presents computational experiments, and Section 5 discusses the quantification of the difference between Bregman and other geometries. Section 6 concludes this paper.

2 Background

We need background on Bregman divergences, Delaunay mosaics, and discrete Morse functions. Indeed, this paper combines these concepts to get new insights into Bregman–Delaunay mosaics and their scale-dependent subcomplexes.

Bregman divergence. Given a suitable convex function on a convex domain, the best affine approximation at a point defines a dissimilarity measure on the domain; see [6]. We follow [4] in the details of this construction, requiring a technical third condition that guarantees a conjugate function of the same kind. Let $\Omega \subseteq \mathbb{R}^d$ be an open and convex domain. A function $F \colon \Omega \to \mathbb{R}$ is of *Legendre type* if

(i) F is differentiable,
(ii) F is strictly convex,
(iii) ∇F diverges whenever we approach the boundary of Ω.

If the boundary of the domain is empty, which is the case for $\Omega = \mathbb{R}^d$, then Condition (iii) is void. In other words, $\|\nabla F(x)\|$ does not necessarily diverge when $\|x\| \to \infty$. Given points $x, y \in \Omega$, the *Bregman divergence* from x to y associated with F is the difference between F

and the best affine approximation of F at y, both evaluated at x:

$$D_F(x\|y) = F(x) - [F(y) + \langle \nabla F(y), x - y \rangle] \,. \tag{1}$$

Note that $D_F(x\|y) \geq 0$, with equality iff $x = y$. However, the other two axioms of a metric do not hold. the divergence is not necessarily symmetric, and it violates the triangle inequality in all non-trivial cases. In spite of these short-comings, Bregman divergences are useful as measures of dissimilarity. For a given $h \geq 0$, the *primal ball* with center x contains all points y such that the divergence from x to y is at most h, and the *dual ball* contains all points y such that the divergence from y to x is at most h:

$$B_F(x, h) = \{y \in \Omega \mid D_F(x\|y) \leq h\}, \tag{2}$$
$$B_F^*(x, h) = \{y \in \Omega \mid D_F(y\|x) \leq h\}. \tag{3}$$

The geometric intuition for (2) is to cast light onto the graph of F from a point vertically above $x \in \mathbb{R}^d$ in $\mathbb{R}^{d+1}$ and at distance h below the graph of F: the primal ball is the vertical projection of the lit up part of the graph onto $\mathbb{R}^d$. This ball is not necessarily convex. The geometric intuition for (3) is to intersect the graph of F with the tangent hyperplane at x shifted vertically upward by a distance h: the dual ball is the vertical projection of the part of the graph on or below this shifted hyperplane. This ball is necessarily convex.

The conjugate of F can be constructed with elementary geometric means. Specifically, we use the *polarity transform* that maps a point $A = (a, a_{d+1}) \in \mathbb{R}^d \times \mathbb{R}$ to the affine map $A^* \colon \mathbb{R}^d \to \mathbb{R}$ defined by $A^*(x) = \langle a, x \rangle - a_{d+1}$. Similarly, it maps A^* to $A = (A^*)^*$. The graph of F can be described as a set of points or a set of affine maps that touch the graph. The *conjugate function*, $F^* \colon \Omega^* \to \mathbb{R}$, is defined such that polarity maps the points of the graph of F to the tangent affine maps of the graph of F^*, and it maps the tangent affine maps of the graph of F to the points of the graph of F^*. Since A and A^* switch position with gradient, so do F and F^*. More specifically, $\Omega^* = \phi(\Omega)$ and $F^* \colon \Omega^* \to \mathbb{R}$ are given by

$$\phi(x) = \nabla F(x), \tag{4}$$
$$F^*(\phi(x)) = \langle \nabla F(x), x \rangle - F(x), \tag{5}$$
$$\nabla F^*(\phi(x)) = x. \tag{6}$$

The convexity of Ω and Conditions (i), (ii), (iii) imply that Ω^* is convex and F^* satisfies (i), (ii), (iii). In other words, the conjugate of a Legendre type function is again a Legendre type function. Importantly, the Bregman divergences associated with F and with F^* are symmetric: $D_F(x\|y) = D_{F^*}(\phi(y)\|\phi(x))$. Hence, ϕ maps primal balls to dual balls and it maps dual balls to primal balls:

$$B_{F^*}^*(\phi(x), h) = \phi(B_F(x, h)), \tag{7}$$
$$B_{F^*}(\phi(x), h) = \phi(B_F^*(x, h)). \tag{8}$$

Since all dual balls are convex, all primal balls are diffeomorphic images of convex sets. This implies that the common intersection of a collection of primal balls is either empty or contractible, so the Nerve Theorem applies [12].

Examples. An important example of a Legendre type function is $\varpi \colon \mathbb{R}^d \to \mathbb{R}$ defined by mapping x to half the square of its Euclidean norm: $\varpi(x) = \frac{1}{2}\|x\|^2$. It is the only Legendre type function that is its own conjugate: $\varpi = \varpi^*$. The symmetry between the divergences of a Legendre type function and its conjugate thus imply $D_\varpi(x\|y) = D_\varpi(y\|x)$ and $B_\varpi(x, h) = B_\varpi^*(x, h)$. Indeed, it is easy to see that the divergence is half the squared

Euclidean distance, $D_\varpi(x\|y) = \frac{1}{2}\|x-y\|^2$, which is of course symmetric. This particular Legendre type function provides an anchor point for comparison.

The example that justifies the title of this paper is the (negative) *Shannon entropy*, $E\colon \mathbb{R}^d_+ \to \mathbb{R}$, defined by $E(x) = \sum_{i=1}^d [x_i \ln x_i - x_i]$. It is of Legendre type and fundamental to information theory. Its divergence,

$$D_E(x\|y) = \sum_{i=1}^{d} [x_i \ln x_i - x_i \ln y_i - x_i + y_i], \tag{9}$$

is generally referred to as the *relative entropy* or the *Kullback–Leibler divergence* from x to y. The gradient of the Shannon entropy at x is the vector $\nabla E(x)$ with components $\ln x_i$ for $1 \le i \le d$. According to (5), the conjugate of E maps this vector to $\sum_{i=1}^d x_i$. Hence $E^*\colon \mathbb{R}^d \to \mathbb{R}$ is defined by mapping $y \in \mathbb{R}^d$ to $E^*(y) = \sum_{i=1}^d e^{y_i}$.

A case of special interest is the restriction of the Shannon entropy to the standard simplex, which is a subset of the positive orthant. Writing $x = (x_1, x_2, \ldots, x_d)$ for a point of $\mathbb{R}^d_+$, the *standard* $(d-1)$*-simplex*, denoted Δ^{d-1}, consists of all points x that satisfy $x_1 + x_2 + \ldots + x_d = 1$. We use Δ^{d-1} as the domain of a Legendre type function, which is the reason we introduce Δ^{d-1} as an open set. Finally, write $E_\Delta\colon \Delta^{d-1} \to \mathbb{R}$ for the restriction of the Shannon entropy to the standard simplex. This setting is important because each $x \in \Delta^{d-1}$ can be interpreted as a probability distribution on d disjoint events. Correspondingly, $-E_\Delta(x) = -E(x)$ is the expected efficiency to optimally encode a sample from this distribution. Finally, the relative entropy from x to y is the expected loss in coding efficiency if we use the code optimized for y to encode a sample from x. Projecting the gradient of the unrestricted Shannon entropy into the hyperplane of the simplex passing through origin, we get the gradient of the restriction:

$$\nabla E_\Delta(x) = \begin{bmatrix} \ln x_1 \\ \ln x_2 \\ \vdots \\ \ln x_d \end{bmatrix} - \frac{1}{d}\sum_{i=1}^{d} \ln x_i \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}. \tag{10}$$

Using (4) and (5), we compute the conjugate of E_Δ, which we state in terms of the barycentric coordinates parametrizing $\mathbb{R}^{d-1}$. Specifically, we get $\phi_\Delta(x) = \nabla E_\Delta(x)$ and

$$\begin{aligned} E^*_\Delta(\phi_\Delta(x)) &= \langle \nabla E_\Delta(x), x\rangle - E_\Delta(x) && (11)\\ &= 1 - \tfrac{1}{d}\textstyle\sum_{i=1}^{d} \ln x_i && (12)\\ &= 1 + \ln \textstyle\sum_{i=1}^{d} e^{y_i}, && (13) \end{aligned}$$

in which the $y_i = \ln x_i - \frac{1}{d}\sum_{i=1}^d \ln x_i$ are the coordinates in conjugate space. Indeed, it is not difficult to verify (13) using $\ln \sum_{i=1}^d x_i = 0$ for points in the standard simplex.

Antonelli isometry. A Bregman divergence gives rise to a path metric in which length is measured by integrating the square root of the divergence. As explained in [11], any divergence that decomposes into a term per coordinate implies an isometry between this path metric and the Euclidean metric. By (9), the relative entropy is an example of such a divergence, and the corresponding path metric is known as the *Fisher metric*, which plays an important role in statistics and information geometry [1]. Instead of formalizing the recipe for constructing the Fisher metric from the relative entropy, we present the isometry with Euclidean space, which was first observed by Antonelli et al. [2]. This is the mapping $\jmath \colon \mathbb{R}_+^d \to \mathbb{R}_+^d$ defined by

$$\jmath(x) = (\sqrt{2x_1}, \sqrt{2x_2}, \ldots, \sqrt{2x_d}). \tag{14}$$

By virtue of being an isometry, the distance between points $x, y \in \mathbb{R}_+^d$ under the Fisher metric satisfies $\|x - y\|_{\mathrm{Fsh}} = \|\jmath(x) - \jmath(y)\|$. The path of this length from x to y is the preimage of the line segment from $\jmath(x)$ to $\jmath(y)$, which is generally not straight.

Of special interest is the Fisher metric restricted to the standard simplex. The mentioned isometry maps Δ^{d-1} to $\jmath(\Delta^{d-1})$, which is the positive orthant of the sphere with radius $\sqrt{2}$ and center at the origin in $\mathbb{R}^d$. The shortest path between $x, y \in \Delta^{d-1}$ is thus the preimage of the great-circle arc that connects $\jmath(x)$ and $\jmath(y)$ on the sphere. Since this arc is generally longer than the straight line segment connecting $\jmath(x)$ and $\jmath(y)$ in $\mathbb{R}_+^d$, the distance between x and y under the Fisher metric restricted to Δ^{d-1} is generally larger than in the unrestricted case.

Alpha shapes and Wrap complexes. Two popular shape reconstruction methods based on Delaunay mosaics are the Alpha shapes introduced in [10] and the Wrap complexes first published in [8]. Both extend to generalized discrete Morse functions and therefore to Bregman–Delaunay mosaics and Bregman–Wrap complexes.

Despite working with Bregman divergences, we only require Euclidean weighted Deluanay mosaics. For brevity, standard definitions and properties are available in Appendix A.

Letting D be a simplicial complex and $f \colon D \to \mathbb{R}$ a generalized discrete Morse function, the *Alpha complex* for h is the sublevel set,

$$\mathrm{Alpha}_h(f) = f^{-1}(-\infty, h], \tag{15}$$

and the *Alpha shape* is the underlying space of the Alpha complex. In contrast to the Alpha shape, the assumption that f be a generalized discrete Morse function is essential in the definition of the Wrap complex. Recall that every step of a generalized discrete Morse function is an interval of simplices in the Hasse diagram. We form the *step graph*, $\mathcal{G} = \mathcal{G}_f$, whose nodes are the steps and whose arcs connect step φ to step ψ if there are simplices $P \in \varphi$ and $Q \in \psi$ with an arc from P to Q in the Hasse diagram. By construction, f is strictly increasing along directed paths in the step graph, which implies that the graph is acyclic.

The *lower set* of a node ν in $\mathcal{G}$, denoted $\downarrow \nu$, is the set of nodes φ for which there are directed paths from φ to ν. Similarly, we write $\downarrow N = \bigcup_{\nu \in N} \downarrow \nu$ for the lower set of a collection of nodes, and $\bigcup \downarrow N$ for the corresponding collection of simplices. We are particularly interested in the set of singular intervals, and we recall that each such interval contains a critical simplex of f. We write Sg_f for the set of singular intervals, and $\mathrm{Sg}_f(h) \subseteq \mathrm{Sg}_f$ for the subset whose simplices satisfy $f(Q) \leq h$. The *Wrap complex* for h is the union of steps in the lower sets of the singular intervals with value at most h:

$$\mathrm{Wrap}_h(f) = \bigcup \downarrow \mathrm{Sg}_f(h). \tag{16}$$

There are alternative constructions of the Wrap complex. Starting with the Alpha complex for h, we get the Wrap complex for the same value by collapsing all non-singular intervals that can be collapsed. The order of the collapses is not important as all orders produce the same result, namely $\mathrm{Wrap}_h(f)$. Symmetrically, we may start with the critical simplices of value at most h and add the minimal collection of non-singular intervals needed to get a simplicial complex. The minimal collection is unique and so is the result, $\mathrm{Wrap}_h(f)$. A proof of the equivalence of these three definitions of the Wrap complex is given in Appendix B.

3 Mosaics and Algorithms

In this section, we review Bregman–Delaunay and Fisher–Delaunay mosaics as well as their scale-dependent subcomplexes. All mosaics are constructed using software for weighted Delaunay mosaics in Euclidean geometry, and all subcomplexes are computed by convex optimization. We begin with the mosaics in Bregman geometry.

Bregman–Delaunay mosaics. Let $\Omega \subseteq \mathbb{R}^d$ be open and convex, consider a Legendre type function $F \colon \Omega \to \mathbb{R}$, and let $U \subseteq \Omega$ be locally finite. Following [5, 12], we define the *Bregman–Voronoi domain* of $u \in U$, denoted $\mathrm{dom}_F(u, \Omega)$, as the points $a \in \Omega$ that satisfies $D_F(u\|a) \leq D_F(v\|a)$ for all $v \in U$. The *Bregman–Voronoi tessellation* is the collection of such domains, and the *Bregman–Delaunay mosaic* mosaic records all non-empty common intersections:

$$\mathrm{Vor}_F(U, \Omega) = \{\mathrm{dom}_F(u, \Omega) \mid u \in U\}, \tag{17}$$

$$\mathrm{Del}_F(U, \Omega) = \{Q \subseteq U \mid \bigcap_{u \in Q} \mathrm{dom}_F(u, \Omega) \neq \emptyset\}, \tag{18}$$

and we note that the mosaic is isomorphic to the nerve of the tessellation. To develop geometric intuition, we observe that $\mathrm{Vor}_F(U, \Omega)$ can be obtained by growing primal Bregman balls with centers at the points $u \in U$. When two such balls meet, they freeze where they touch but keep growing everywhere else. Eventually, each ball covers exactly the corresponding domain. Since the primal balls are not necessarily convex, it is not surprising that the faces shared by the domains are not necessarily straight. Nevertheless, the Delaunay mosaic has a natural straight-line embedding as all its cells are vertical projections of lower faces of the convex hull of the points $(u, F(u)) \in \mathbb{R}^{d+1}$. To see this, we note that each cell of the mosaic corresponds to a dual Bregman ball whose boundary passes through the vertices of the cell, and this ball is the vertical projection of the part of the graph of F on or below the graph of an affine function.

Construction. To construct the mosaic, we assume that $U \subseteq \Omega$ is in general position, by which we mean that Conditions I and II are satisfied after transforming $U \subseteq \Omega$ to $X \subseteq \mathbb{R}^d \times \mathbb{R}$ such that $\mathrm{Del}_F(U, \Omega)$ is a subcomplex of the weighted Delaunay mosaic of X. Lifting the points from $\mathbb{R}^d$ to $\mathbb{R}^{d+1}$ and projecting the lower boundary of the convex hull back to $\mathbb{R}^d$, we get the mosaic. We remind the reader that relevant background information can be found in Appendix A, and define $\varpi(a) = \frac{1}{2}\|a\|^2$.

We formalize this method while stating all steps in terms of weighted points in d dimensions:

Step 1: Let $X \subseteq \mathbb{R}^d \times \mathbb{R}$ be the set of weighted points $x(u) = (u, 2\varpi(u) - 2F(u))$, with $u \in U$.

Step 2: Compute the weighted Delaunay mosaic of X in Euclidean geometry, denoted $\mathrm{Del}(X)$.

Step 3: Select $\mathrm{Del}_F(U, \Omega)$ as the collection of simplices in $\mathrm{Del}(X)$ whose corresponding weighted Voronoi cells have a non-empty intersection with Ω^*.

Indeed, the weighted Delaunay mosaic computed in Step 2 may contain simplices that do not belong to the Delaunay–Bregman mosaic of F. To implement Step 3, we note that $\mathrm{Del}_F(U, \Omega)$ is dual to $\mathrm{Vor}_F(U, \Omega)$, which is isomorphic to $\mathrm{Vor}_{F^*}(\phi(U), \Omega^*)$, and this Bregman–Voronoi tessellation is the weighted Voronoi tessellation of X restricted to Ω^*. This tessellation has convex polyhedral cells and is readily available as the dual of $\mathrm{Del}(X)$. Writing $Y(Q) \subseteq X$ for the points $x(u)$ with $u \in Q \subseteq U$ and $\mathrm{dom}(Y)$ for the weighted Voronoi cell that corresponds to $Y \in \mathrm{Del}(X)$, we have

$$\mathrm{Del}_F(U,\Omega) = \{Q \subseteq U \mid \mathrm{dom}(Y(Q)) \cap \Omega^* \neq \emptyset\}. \tag{19}$$

Instead of computing all these intersections, we can collapse $\mathrm{Del}(X)$ to the desired subcomplex and thus save time by looking only at a subset of the mosaic. We explain how the simplices can be organized to facilitate such a collapse. Recalling that $\Omega^* \subseteq \mathbb{R}^d$ is open and convex, we introduce the *signed distance function*, $\theta\colon \mathbb{R}^d \to \mathbb{R}$, which maps every $a \in \mathbb{R}^d$ to plus or minus $r = r(a)$ such that the sphere with center a and radius r touches $\partial\Omega^*$ but does not cross the boundary. Finally, $\theta(a) = r(a)$ if $a \notin \Omega^*$ and $\theta(a) = -r(a)$ if $a \in \Omega^*$. Note that $\Omega^* = \theta^{-1}[-\infty, 0)$ and that $\Omega_t^* = \theta^{-1}[-\infty, t)$ is open and convex for every t. Now construct $\vartheta\colon \mathrm{Del}(X) \to \mathbb{R}$ by mapping $Y \in \mathrm{Del}(X)$ to the maximum $t \in \mathbb{R}$ for which $\mathrm{dom}(Y) \cap \Omega_t^* = \emptyset$. By (19), we get $\mathrm{Del}_F(U,\Omega)$ by removing all simplices Y with $\vartheta(Y) \geq 0$. The crucial observation is that for X in general position, ϑ is a generalized discrete Morse function with a single critical vertex. To see this, we observe that $\mathrm{Vor}(X)$ decomposes Ω_t^* into convex domains for every value t, which by the Nerve Theorem implies that $\vartheta^{-1}(-\infty, t]$ is contractible. Removing the simplices in sequence of decreasing values of ϑ thus translates into a sequence of collapses that preserve the homotopy type of the mosaic.

Rise functions. To introduce scale into the construction of Bregman–Delaunay mosaics, we generalize the radius function from Euclidean geometry to Bregman geometries, changing the name because size is more conveniently measured by height difference in the $(d+1)$-st coordinate direction as opposed to the radius in $\mathbb{R}^d$. Let $\dot{u} = (u, F(u))$ and $\bar{u}\colon \mathbb{R}^d \to \mathbb{R}$ be the point and affine map that correspond to $u \in \Omega$, and let $\upsilon\colon \mathbb{R}^d \to \mathbb{R}$ be the upper envelope of the $\bar{u}$, $u \in U$. We introduce the *rise function*, $\varrho_F\colon \mathrm{Del}_F(U,\Omega) \to \mathbb{R}$, which maps each simplex, Q, to the minimum difference between F^* and υ at points in the conjugate Voronoi cell:

$$\varrho_F(Q) = \inf_{a \in \phi(\mathrm{dom}(Q,\Omega))} [F^*(a) - \upsilon(a)]. \tag{20}$$

It is the infimum amount we have to lower the graph of F^* until it intersects the graph of υ at a point vertically above the Voronoi cell in conjugate space. Without going to the conjugate, we can interpret $\varrho_F(Q)$ in terms of (primal) Voronoi domains and cones of light cast from the $\dot{u}$ onto the graph, which we raise until the cones clipped to within their Voronoi domains have a point in common. This interpretation motivates the name of the function. Comparing (20) with (32), we see that the two agree when $F = \varpi$ and $\Omega = \mathbb{R}^d$. Indeed, we get $F^* = \varpi$ and $\upsilon = \xi$. Furthermore, $\phi(\mathrm{dom}(Q,\Omega)) = \mathrm{dom}(Q,\Omega)$, and taking the infimum is the same as taking the minimum.

For every $h \in \mathbb{R}$, we have a sublevel set, $\mathrm{Del}_{F,h}(U,\Omega) = {\varrho_F}^{-1}(-\infty, h]$, which we refer to as the *Bregman–Alpha complex* of U and F for size h. For $h < 0$, this complex is empty, for $h = 0$, it is a set of vertices namely the points in U, and for sufficiently large positive h, this complex is $\mathrm{Del}_F(U,\Omega)$.

Computation. We compute the rise function following the intuition based on primal Voronoi domains explained below (20). Equivalently, $\varrho_F(Q)$ is the minimum amount we have to raise the graph of F so it has a supporting hyperplane that passes through all points $\dot{u}$, with $u \in Q$, while all other point $\dot{u}$, with $u \in U$, lie on or above the hyperplane.

To turn this intuition into an algorithm, we consider the affine hull of Q and write $\bar{v}\colon \mathrm{aff}\, Q \to \mathbb{R}$ for the affine function that satisfies $\bar{v}(u) = F(u)$ for all $u \in Q$. Let $H\colon \mathrm{aff}\, Q \cap \Omega \to \mathbb{R}$ measure the difference: $H(a) = F(a) - \bar{v}(a)$. Since F is of Legendre type, so is H. We are interested in the infimum of H, which either occurs at a point in $\mathrm{aff}\, Q \cap \Omega$ or at the limit of a divergent sequence. We therefore introduce a numerical routine that returns both, the infimum and the point where it occurs:

```
InfSize (function F, simplex Q):
    (a_Q, h_Q) = (arginf H, inf H);
    return (a_Q, h_Q).
```

Note that the dual Bregman ball centered at $a_Q \in \operatorname{aff} Q \cap \Omega$ and size h_Q contains Q in its boundary, and it may or may not contain points of $U \setminus Q$ in its interior. If it does not, then $\varrho_F(Q) = h_Q$, otherwise, $\varrho_F(Q)$ is the minimum function value of the proper cofaces of Q. To express this more formally, we write $\mathrm{coFacets}(Q)$ for the collection of simplices $R \in \mathrm{Del}(X)$ with $Q \subseteq R$ and $\#R = \#Q + 1$. Since Q gets its value either directly or from a coface, it is opportune to compute the rise function in the order of decreasing dimension:

```
for p = d downto 0 do
    forall p-simplices Q ∈ Del_F(U, Ω) do
        (a_Q, h_Q) = InfSize(F, Q);
        if B*_F(a_Q, h_Q) ∩ [U \ Q] = ∅
            then ϱ_F(Q) = h_Q
            else ϱ_F(Q) = min_{R ∈ coFacets(Q)} ϱ_F(R).
```

Note that this algorithm assigns a value to every simplex in $\mathrm{Del}_F(U, \Omega)$. Indeed, the simplices in $\mathrm{Del}(X)$ that are not in $\mathrm{Del}_F(U, \Omega)$ have been culled in Step 3, as explained above.

Fisher metric. In addition to the Bregman divergences, we consider Delaunay mosaics under the Fisher metric. To construct them, we recall that the mapping $\jmath \colon \mathbb{R}^d_+ \to \mathbb{R}^d_+$ defined by $\jmath(x) = (\sqrt{2x_1}, \sqrt{2x_2}, \ldots, \sqrt{2x_d})$ is an isometry between the Fisher metric and the Euclidean metric. This suggests the following algorithm.

Step 1: Compute the Delaunay mosaic of $\jmath(U)$ in Euclidean space.

Step 2: Remove the simplices from $\mathrm{Del}(\jmath(U))$ whose dual Voronoi cells have an empty intersection with $\mathbb{R}^d_+$.

Step 3: Draw the resulting complex by mapping each point $\jmath(u)$ to the original point $u \in U \subseteq \mathbb{R}^d_+$.

The rise function in Euclidean geometry maps every simplex $\jmath(Q) \in \mathrm{Del}(\jmath(U))$ to the squared radius of the smallest empty circumsphere of $\jmath(Q)$. By isometry, the preimage of this Euclidean sphere is the smallest empty circumsphere of Q under the Fisher metric, and the squared radius is the same. We thus get the rise function on the Fisher–Delaunay mosaic by copying the values of the rise function on the Delaunay mosaic in Euclidean geometry.

The construction of the mosaic for the Fisher metric restricted to the standard simplex, Δ^{d-1}, is only slightly more complicated. As mentioned in Section 2, the isometry maps Δ^{d-1} to $\sqrt{2}\mathbb{S}^{d-1}_+$, which is our notation for the positive orthant of the sphere with radius $\sqrt{2}$ centered at the origin in $\mathbb{R}^d$. The distance between points $u, v \in \Delta^{d-1}$ under the Fisher metric thus equals the Euclidean length of the great-circle arc connecting $\jmath(u), \jmath(v) \in \sqrt{2}\mathbb{S}^{d-1}_+$. The Delaunay mosaic of $\jmath(U)$ under the geodesic distance can be obtained by constructing the convex hull of $\jmath(U) \cup \{0\}$ in $\mathbb{R}^d$ and centrally projecting all faces not incident to 0 onto the sphere. As before, we cull simplices whose dual Voronoi cells have an empty intersection with the positive orthant of the sphere, and we draw the mosaic in Δ^{d-1} by mapping the vertices back to the original points. Furthermore, the rise functions of the mosaics in $\sqrt{2}\mathbb{S}^{d-1}_+$ and in Δ^{d-1} are the same. Note however, that the geodesic radius is the arc-sine of and therefore slightly larger than the straight Euclidean radius in $\mathbb{R}^d$.

4 Computational Experiments

We illustrate the Bregman–Alpha and Bregman–Wrap complexes while comparing them to the conjugate, the Fisher, and the Euclidean constructions.

Example in positive quadrant. Let X be a set of 1000 points uniformly distributed according to the Fisher metric in $(0,2]^2 \subseteq \mathbb{R}^2_+$. To sample X, we use the isometry, $\jmath \colon \mathbb{R}^2_+ \to \mathbb{R}^2_+$, between the Fisher and the Euclidean metric mentioned in Section 2. Specifically, we sample 1000 points uniformly at random according to the Euclidean metric in $(0,2]^2$, and we map each point with coordinates x_1, x_2 to $\jmath^{-1}(x_1,x_2) = \frac{1}{2}(x_1^2, x_2^2)$, which is again a point in $(0,2]^2$. To compute the Delaunay mosaic in Fisher geometry, we construct the (Euclidean) Delaunay mosaic of $\jmath(X)$ and draw this mosaic with the vertices at the points in X. Recall however that the domain is $\Omega = \mathbb{R}^d_+$ and not $\mathbb{R}^d$. A simplex whose corresponding Voronoi cell has an empty intersection with the positive orthant thus does not belong to the mosaic, which is restricted to Ω. We identify these simplices and remove them from the Delaunay mosaic as described in Section 3.

Figure 1 displays the Bregman–Alpha complex in Shannon geometry for threshold 0.004. Infinitesimally, the relative entropy agrees with the squared Fisher metric, so the uniform distribution of the points translates into a fairly uniform arrangement of random holes in the complex. The closer we get to the left or the lower side of the square, the denser the points get and the more anisotropically aligned with the sides the edges and triangles get.

For comparison, Figure 2 shows the Bregman–Alpha complex in conjugate Shannon geometry, in Fisher geometry, in Euclidean geometry, and in weighted Euclidean geometry. The primal and the dual balls behave similarly, which explains the similarity of the complexes in Figure 1 and in Figure 2(a). It should however be mentioned that the underlying triangulation in 2(a) occasionally folds, which is caused by moving the vertices from the conjugate points (for which we have a straight-line embedding) to the original points. Not surprisingly, there is also a striking similarity to the reconstruction in Fisher geometry 2(b). The Bregman–Alpha complex in Euclidean geometry 2(c) is just the usual Alpha complex of the points. It clearly shows that the density decreases along the diagonal. The complex in 2(d) mixes aspects of Shannon and Euclidean geometry. In particular, it reuses the mosaic in Figure 1 and assigns weights to the points such that this triangulation is the weighted Delaunay mosaic of the weighted points in Euclidean geometry. The corresponding rise function reflects the difference between the Shannon entropy and the squared Euclidean norm. Indeed, the rise function increases along the diagonal, which explains why the reconstructed complex is almost the entire mosaic, with cells along the left and bottom sides of the square domain missing.

We see very similar reconstructions in Figures 3 and 4, which show the Bregman–Wrap complexes for the same set of points and the same threshold. By construction, each Wrap complex is a homotopy equivalent subcomplex of the corresponding Alpha complex. The biggest difference occurs in weighted Euclidean geometry, in which we reuse the mosaic in Shannon geometry but filter with the rise function obtained from the squared Euclidean norm. The corresponding Bregman–Wrap complex consists of a single vertex near the upper right corner of the square domain; see Figure 4(d). This reconstruction reflects the simple relation between the Shannon entropy and the halved squared Euclidean norm: $\varpi(x) - E(x)$ is monotonically increasing from left to right and from bottom to top. This translates into a discrete gradient that introduces a flow with a single critical cell, namely the vertex near the upper right corner.

Example in standard triangle. Motivated by our interest in information-theoretic applications, we repeat the above experiment within the standard triangle, Δ^2, which consists of all points $(x_1, x_2, x_3) \in \mathbb{R}^3_+$ that satisfy $x_1 + x_2 + x_3 = 1$. Every point in Δ^2 can be interpreted as a probability distribution on three disjoint events, which is indeed the most relevant scenario for the application of the relative entropy. To sample a set Y of 1000 points uniformly at random according to the Fisher metric in Δ^2, we use again $\jmath$, now restricted to Δ^2, whose image is the positive orthant of the sphere with radius $\sqrt{2}$ centered at the origin of $\mathbb{R}^3$. Sampling 1000 points uniformly at random according to the geodesic distance on the sphere, we take the convex hull of $\jmath(Y) \cup \{0\}$ and get the mosaic by mapping the vertices to the points in $Y = \jmath^{-1}(\jmath(Y))$. Before drawing the faces in Δ^2, we remove 0 and all incident faces, as well as all faces whose corresponding Voronoi cells have an empty intersection with $\mathbb{R}^2_+$.

Recall that the squared Fisher metric matches the relative entropy in the infinitesimal regime, which explains the random appearance of the reconstruction in Figure 5 for which we set the threshold to 0.0025. As in the above example, the reconstruction in Shannon geometry is similar to those in conjugate Shannon geometry in Figure 6(a) and in Fisher geometry in Figure 6(b). To interpret the reconstruction in 6(d), we observe that the difference between the Shannon entropy and the squared Euclidean norm has a minimum at the center and no other critical points in the interior of the triangular domain. Accordingly, the reconstruction removes simplices near the corners and the three sides first. More drastically, the Bregman–Wrap complex for the same data removes all simplices except for a single critical edge near the center; see Figure 8(d).

5 Quantification of Difference

We take a data-centric approach to quantifying the differences between the geometries. Given a common domain, Ω, and a finite set of points, $U \subseteq \Omega$, we compare the corresponding mosaics and rise functions.

Mosaics. The Delaunay mosaics of U depend on the local shape of the balls defined by the metric or the divergence. Letting D and E be two Delaunay mosaics with vertex sets U, we compare them by counting the common cells:

$$J(D, E) = 1 - \frac{\#(D \cap E)}{\#D + \#E - \#(D \cap E)}, \tag{21}$$

which is sometimes referred to as the *Jaccard distance* between the two sets. It is normalized so that $J = 0$ iff $D = E$ and $J = 1$ iff D and E share no cells at all. In our application, the two mosaics share all vertices, so J is necessarily strictly smaller than 1. To apply this measure to the constructions in Section 4, we write D_0, D_1, D_2, D_3, D_4 for the mosaics in Figures 9 and 10, and we write E_0, E_1, E_2, E_3, E_4 for the mosaics in Figures 13 and 14. All mosaics are different, except for $D_0 = D_4$ and $E_0 = E_4$. The Jaccard distances are given in Table 1. We see that the mosaics in conjugate Shannon geometry and in Fisher geometry are most similar to each other and less similar to the mosaic in Shannon geometry. The mosaic in Euclidean geometry is most dissimilar to the others. See Figures 9, 10 and 13, 14 for visual confirmation.

Rise functions. Different rise functions on the same mosaic can be compared by counting the inversions, which are the pairs of cells whose orderings are different under the two functions. Recall that $D_0 = D_4$ and $E_0 = E_4$, let $d_0 \colon D_0 \to \mathbb{R}$ and $e_0 \colon E_0 \to \mathbb{R}$ be the rise functions in Shannon geometry, and let $d_4 \colon D_4 \to \mathbb{R}$ and $e_4 \colon E_4 \to \mathbb{R}$ be the rise functions in

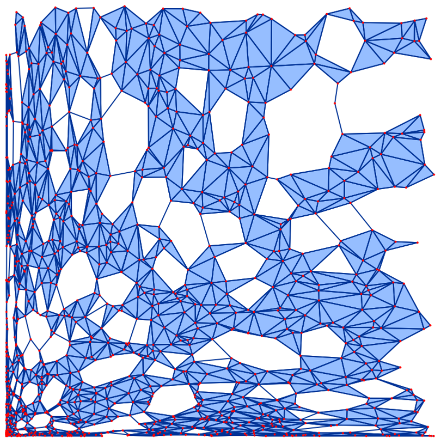

Figure 1: The Bregman–Alpha complex in Shannon geometry of a set X of 1000 points uniformly distributed according to the Fisher metric in $(0, 2]^2$ and a threshold $h = 0.004$.

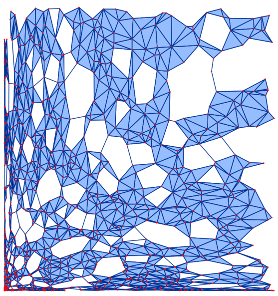

(a) Conjugate Shannon.

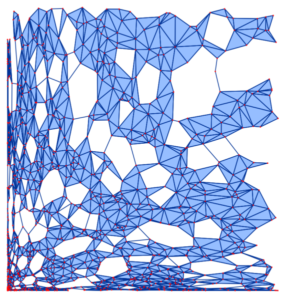

(b) Fisher.

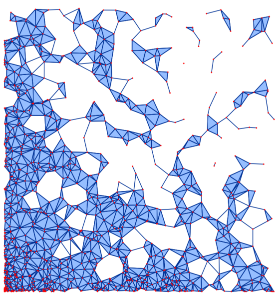

(c) Euclidean.

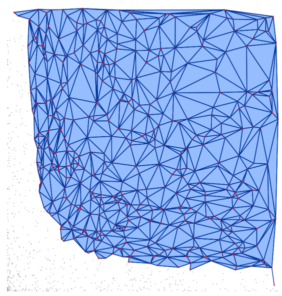

(d) Weighted Euclidean.

Figure 2: The reconstructions in four different geometries for the same points and the same threshold as in Figure 1.

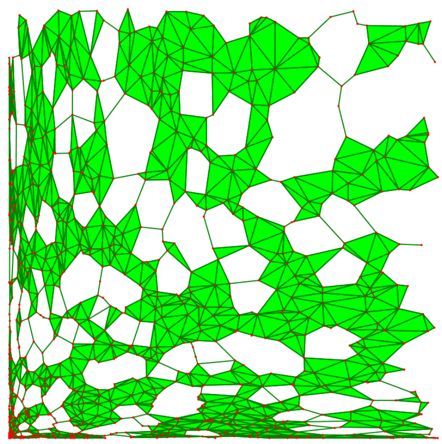

Figure 3: The Bregman–Wrap complex in Shannon geometry of the same points and the same threshold as in Figure 1.

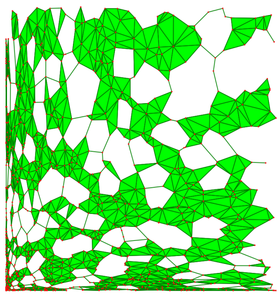

(a) Conjugate Shannon.

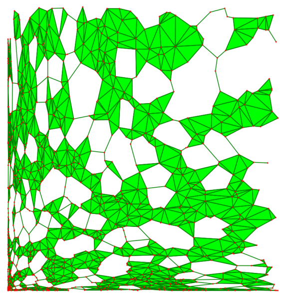

(b) Fisher.

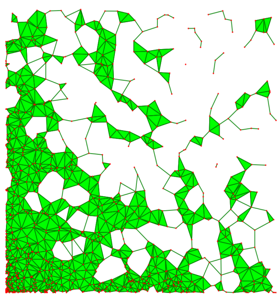

(c) Euclidean.

(d) Weighted Euclidean.

Figure 4: The reconstructions in four different geometries for the same points and the same threshold as in Figure 3.

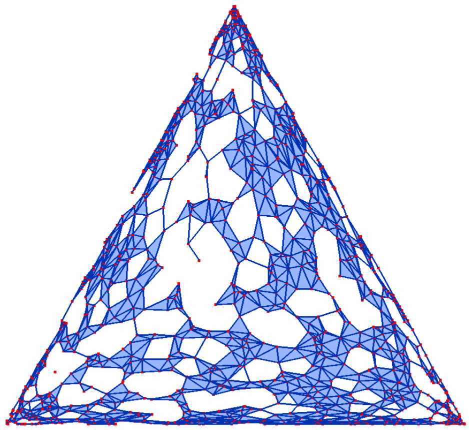

Figure 5: The Bregman–Alpha complex in Shannon geometry of a set Y of 1000 random points in Δ^2 with threshold $h = 0.0025$.

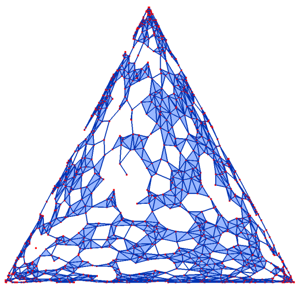

(a) Conjugate Shannon.

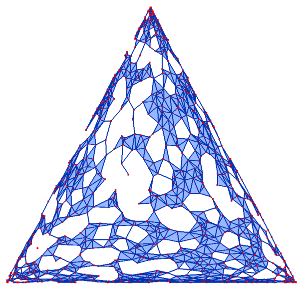

(b) Fisher.

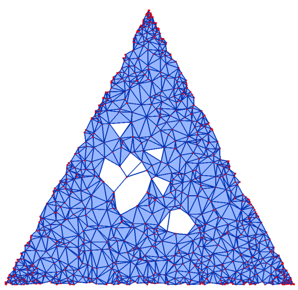

(c) Euclidean.

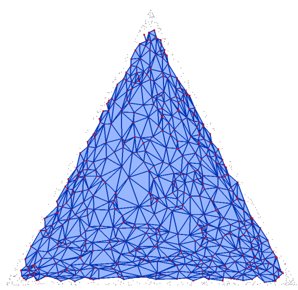

(d) Weighted Euclidean.

Figure 6: The reconstructions in four different geometries for the same points and threshold as in Figure 5.

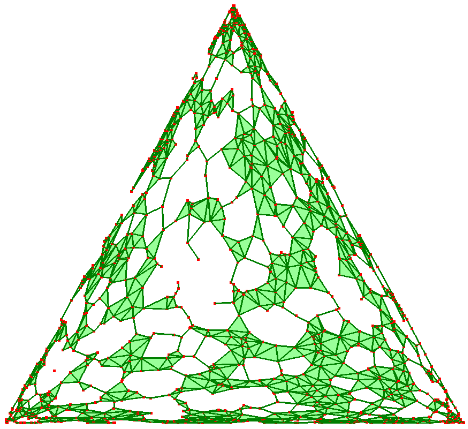

Figure 7: The Bregman–Wrap complex in Shannon geometry of the same points and threshold as in Figure 5.

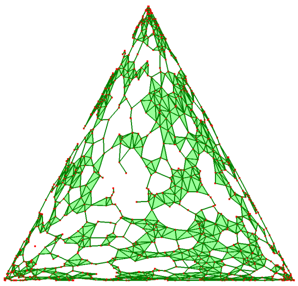

(a) Conjugate Shannon.

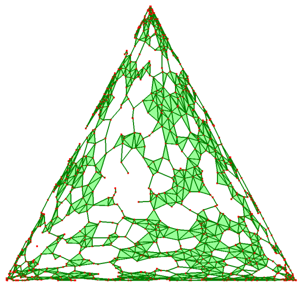

(b) Fisher.

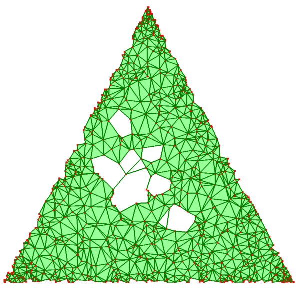

(c) Euclidean.

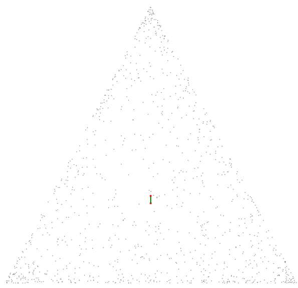

(d) Weighted Euclidean.

Figure 8: The reconstructions in four different geometries for the same points and threshold as in Figure 7.

J	D_0	D_1	D_2	D_3	D_4
D_0	0.00	0.06	0.04	0.48	0.00
D_1		0.00	0.02	0.47	0.06
D_2			0.00	0.47	0.04
D_3				0.00	0.48
D_4					0.00
	E_0	E_1	E_2	E_3	E_4
E_0	0.00	0.10	0.06	0.52	0.00
E_1		0.00	0.04	0.51	0.10
E_2			0.00	0.51	0.06
E_3				0.00	0.52
E_4					0.00

Table 1: The Jaccard distances between the Delaunay mosaics in Shannon, conjugate Shannon, Fisher, Euclidean, and weighted Euclidean geometries for points in the positive quadrant on the *top* and in the standard triangle on the *bottom*.

weighted Euclidean geometry. The normalized number of inversions are

$$I(d_0, d_4) = 0.476, \tag{22}$$

$$I(e_0, e_4) = 0.467. \tag{23}$$

In words, slightly fewer than half the pairs are inversions, both for d_0, d_4 and for e_0, e_4. This is plausible because d_4 orders the cells along the diagonal while d_0 preserves the random character of the point sample; see Figures 11 and 12(d). Similarly, e_4 orders the cells radially, from the center of the standard triangle to its periphery, while e_0 preserves again the random character of the sample; see Figures 15 and 16(d).

We can compare the rise functions also visually, by color-coding the 2-dimensional cells, and this works even if the mosaics are different. Specifically, we shade the triangles by mapping small to large rise function values to dark to light color. In Figures 11, 12(a), and 12(b), this leads to randomly mixed dark and light triangles, while in Figures 12(c) and 12(d) there are clear but opposing gradients parallel to the diagonal. Similarly, in 16(c) we see the rise function decrease from the center to the boundary of the standard triangle, and in 16(d) we see it increasing from the center to the boundary. In addition, we compare general rise functions by computing their persistence diagrams; see [9]. Writing $\mathrm{Dgm}(d)$ for the persistence diagram of function d, we quantify the difference with the bottleneck between the diagrams:

$$B(d, e) = W_\infty(\mathrm{Dgm}(d), \mathrm{Dgm}(e)). \tag{24}$$

As explained in [9], the bottleneck distance is 1-Lipschitz, that is: $B(d, e) \leq \|d - e\|_\infty$, but $d \neq e$ does not necessary imply $B(d, e) \neq 0$. The bottleneck distances between the $d_i \colon D_i \to \mathbb{R}$ and the $e_i \colon E_i \to \mathbb{R}$ are given in Table 2. In part this comparison agrees with the Jaccard distances between the mosaics given in Table 1. The most obvious disagreements are for d_0, d_4 and for e_0, e_4, in which quite different functions are defined on identical mosaics.

6 Discussion

We formulate two popular Euclidean shape reconstruction methods within the framework of discrete Morse functions and show how this generalizes the methods to data in Bregman and Fisher geometries without the need to develop customized software. Turning the table, we use these generalized shape reconstruction methods to compare different geometries

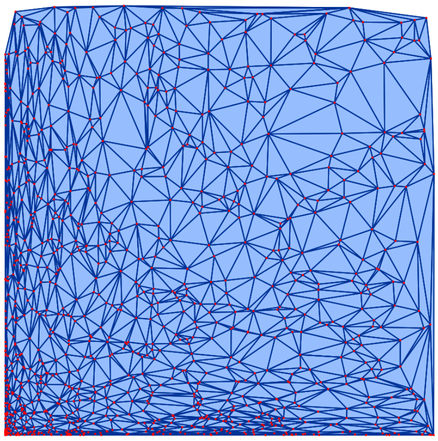

Figure 9: The Bregman–Delaunay mosaic in Shannon geometry for the same set of points as used in Figures 1 to 4.

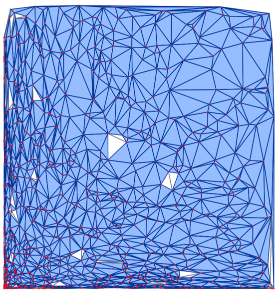

(a) Conjugate Shannon.

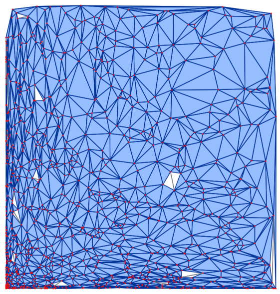

(b) Fisher.

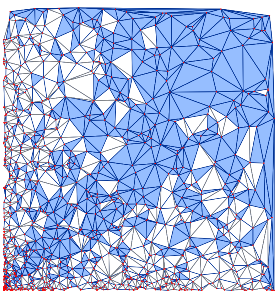

(c) Euclidean.

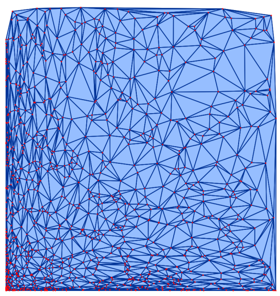

(d) Weighted Euclidean.

Figure 10: Four Delaunay mosaics whose triangles and edges are colored depending on whether or not they belong to the Shannon–Delaunay mosaic in Figure 9.

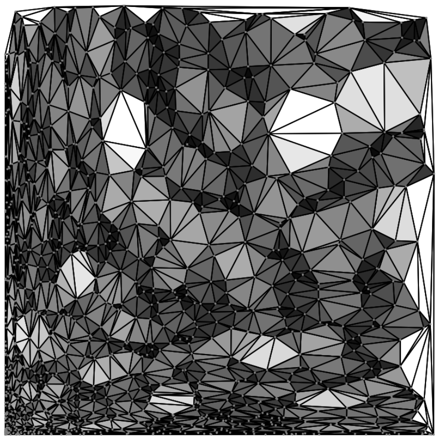

Figure 11: A color-coded Bregman–Delaunay mosaic in Shannon geometry. The set of points is the same as in Figure 9.

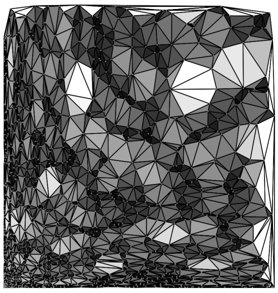

(a) Conjugate Shannon.

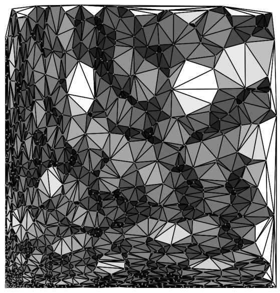

(b) Fisher.

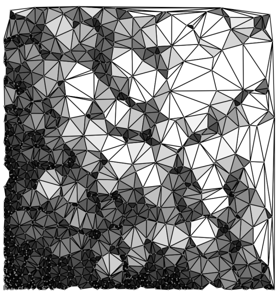

(c) Euclidean.

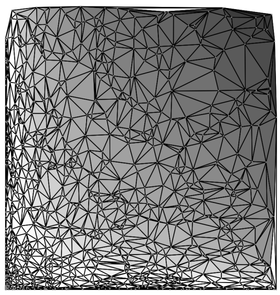

(d) Weighted Euclidean.

Figure 12: The color-coded Delaunay mosaics for the same set X as in Figure 11.

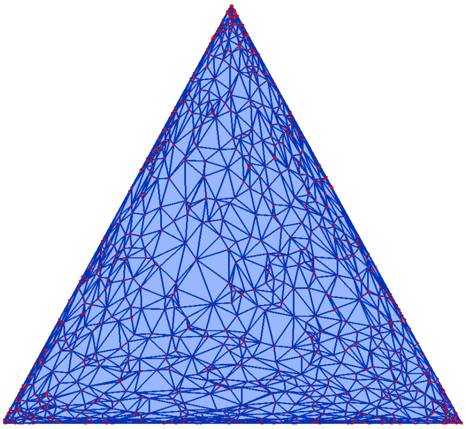

Figure 13: The Bregman–Delaunay mosaic in Shannon geometry for the same set of points as used in Figures 5 to 8.

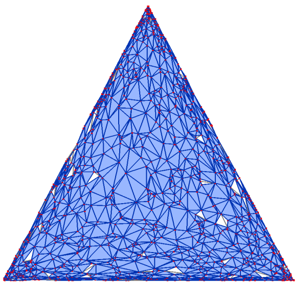

(a) Conjugate Shannon.

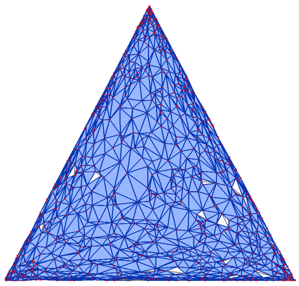

(b) Fisher.

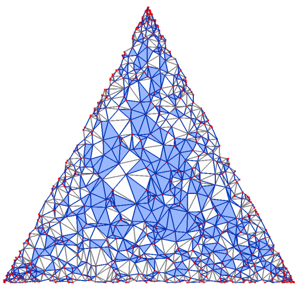

(c) Euclidean.

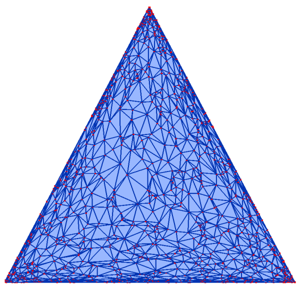

(d) Weighted Euclidean.

Figure 14: Four Delaunay mosaics whose triangles and edges are colored depending on whether or not they belong to the Shannon–Delaunay mosaic in Figure 13.

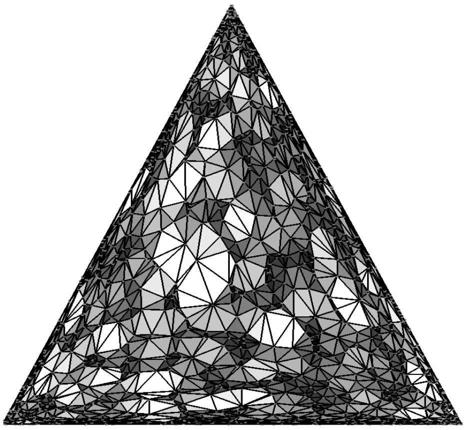

Figure 15: The color-coded Bregman–Delaunay mosaic in Shannon geometry of the same set of points as in Figure 13.

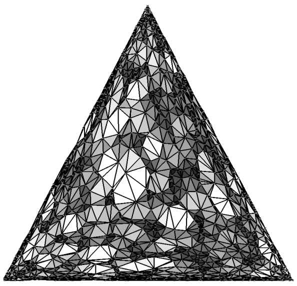

(a) Conjugate Shannon.

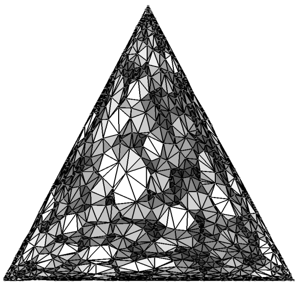

(b) Fisher.

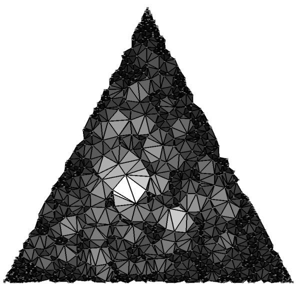

(c) Euclidean.

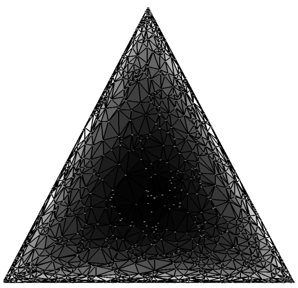

(d) Weighted Euclidean.

Figure 16: The color-coded Delaunay mosaics for the same set as in Figure 15.

B	d_0	d_1	d_2	d_3	d_4
d_0	0.0000	0.0028	0.0004	0.0126	0.0048
d_1		0.0000	0.0028	0.0126	0.0048
d_2			0.0000	0.0126	0.0048
d_3				0.0000	0.0126
d_4					0.0000
	e_0	e_1	e_2	e_3	e_4
e_0	0.0000	0.0006	0.0003	0.0031	0.0035
e_1		0.0000	0.0003	0.0030	0.0034
e_2			0.0000	0.0030	0.0034
e_3				0.0000	0.0023
e_4					0.0000

Table 2: The bottleneck distances between the persistence diagrams of the rise functions on the Delaunay mosaics in Shannon, conjugate Shannon, Fisher, Euclidean, and weighted Euclidean geometries for points in the positive orthant on the *top* and points in the standard triangle on the *bottom*.

experimentally. Our experimental approach to study geometries is a first step in this direction. It is prudent to ask how it can be improved and whether there are more effect experimental approaches to understand metric spaces.

- Can the sensitivity of Delaunay mosaics to the dissimilarity be quantified probabilistically, as the expected Jaccard distance for random point processes?
- Are homotopies between filtrations better measures of the dissimilarity between filtrations than the normalized number of inversions?
- Persistence has been used before to compare metric spaces [7], and it would be interesting to know whether there are deeper connections to our work.

On a practical note, our comparison suggests that the Shannon and Fisher geometries are quite similar, at least in low dimensions. Is this true in higher dimensions? How does this generalize to other Bregman divergences and the corresponding generalized Fisher metrics? To what extent can the Fisher space replace the Shannon space in various applications?

Finally, we mention a concrete question concerning the Delaunay mosaics in Fisher geometry: is the drawing we get by mapping the vertices to the corresponding points and connecting these point with straight edges, flat triangles, etc. necessarily a geometric realization of the mosaic?

References

1. S. Amari and H. Nagaoka. *Methods of Information Geometry.* Amer. Math. Soc., Providence, Rhode Island, 2000.
2. P.L. Antonelli et al. The geometry of random drift I-VI. *Adv. Appl. Prob.* **9-12** (1977-80).
3. U. Bauer and H. Edelsbrunner. The Morse theory of Čech and Delaunay complexes. *Trans. Amer. Math. Soc.*, **369** (2017), 3741–3762.
4. H.H. Bauschke and J.M. Borwein. Legendre functions and the method of random Bregman projections. *J. Convex Analysis* **4** (1997), 27-67.
5. J.-D. Boissonnat, F. Nielsen and R. Nock. Bregman Voronoi diagrams. *Discrete Comput. Geom.* **44** (2010), 281–307.
6. L.M. Bregman. The relaxation method of finding the common point of convex sets and its applications to the solution of problems in convex programming. *USSR Comput. Math. Math. Phys.* **7** (1967), 200–217.
7. F. Chazal, D. Cohen-Steiner, L.J. Guibas, F. Mémoli and S.Y. Oudot. Gromov–Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum* **28** (2009), 1393–1403.
8. H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. In *Discrete and Computational Geometry. The Goodman–Pollack Festschrift*, 379–404, eds. B. Aronov, S. Basu, J. Pach and M. Sharir, Springer-Verlag, 2003.

9 H. Edelsbrunner and J.L. Harer. *Computational Topology. An Introduction.* Amer. Math. Soc., Providence, Rhode Island, 2010.
10 H. Edelsbrunner, D.G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **IT-29** (1983), 551–559.
11 H. Edelsbrunner, Z. Virk and H. Wagner. Topological data analysis in information space. *In* "Proc. 35th Ann. Sympos. Comput. Geom., 2019", 31:1–31:14.
12 H. Edelsbrunner and H. Wagner. Topological data analysis with Bregman divergences. *In* "Proc. 33rd Ann. Sympos. Comput. Geom., 2017", 39:1–39:16.
13 R. Forman. Morse theory for cell complexes. *Adv. Math.* **134** (1998), 90–145.
14 R. Freij. Equivariant discrete Morse theory. *Discrete Math.* **309** (2009), 3821–3829.
15 G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième Mémoire: Recherches sur les paralléloèdres primitifs. *J. Reine Angew. Math.* **134** (1908), 198–287.

Acknowledgements

We thank Anton Nikitenko for first observing that the Wrap complex can be characterized as stated in Claim (ii) of the Wrap Complex Lemma, and Ondrej Draganov for correcting a critical mistake in one of our formulas in Section 2.

A Standard background on Delaunay mosaics and related topics.

We recall standard definitions related to Delaunay mosaics, the corresponding liftings and projections, as well as discrete Morse theory.

Delaunay mosaics. In this paper, the ability to assign real weights to points is essential, so we go straight to the weighted generalizations of the Voronoi tessellation and the Delaunay mosaic. A *weighted point* is a pair $x = (\mathrm{pt}(x), \mathrm{wt}(x)) \in \mathbb{R}^d \times \mathbb{R}$, in which $\mathrm{pt}(x)$ is its *location* and $\mathrm{wt}(x)$ is its *weight*. The *power distance* of $a \in \mathbb{R}^d$ from x is $\pi_x(a) = \|\mathrm{pt}(x) - a\|^2 - \mathrm{wt}(x)$. It is common to think of the weighted point as a ball with center $\mathrm{pt}(x)$ and squared radius $\mathrm{wt}(x)$. With this interpretation, $\pi_x(a)$ is negative inside, zero on the boundary, and positive outside the ball. Given a locally finite set of weighted points, $X \subseteq \mathbb{R}^d \times \mathbb{R}$, the *(weighted) Voronoi domain* of $x \in X$ consists of all points a for which x minimizes the power distance, and the *(weighted) Voronoi tessellation* of X is the collection of such domains:

$$\mathrm{dom}(x) = \{a \in \mathbb{R}^d \mid \pi_x(a) \leq \pi_y(a), \forall y \in X\}, \tag{25}$$

$$\mathrm{Vor}(X) = \{\mathrm{dom}(x) \mid x \in X\}. \tag{26}$$

A *(weighted) Voronoi cell* is the common intersection of Voronoi domains, and we write $\mathrm{dom}(Q) = \bigcap_{x \in Q} \mathrm{dom}(x)$. Note that the affine hull of $\mathrm{dom}(Q)$ contains a unique point, denoted a_Q, that minimizes the power distance to the weighted points in Q. Indeed, a_Q is at the intersection of the affine hull of $\mathrm{dom}(Q)$ and the affine hull of the locations $\mathrm{pt}(x)$, $x \in Q$. Let $\#Q$ be the cardinality of Q. We are primarily interested in the generic case, when every non-empty Voronoi cell, $\mathrm{dom}(Q)$, satisfies the following two *general position conditions*:

I. the dimension of $\mathrm{dom}(Q)$ is $d + 1 - \#Q$,
II. a_Q does not belong to the boundary of $\mathrm{dom}(Q)$.

By Condition I, $\mathrm{dom}(Q) = \emptyset$ whenever $\#Q > d + 1$. Condition I also implies that every non-empty Voronoi cell is the intersection of a unique collection of Voronoi domains. The

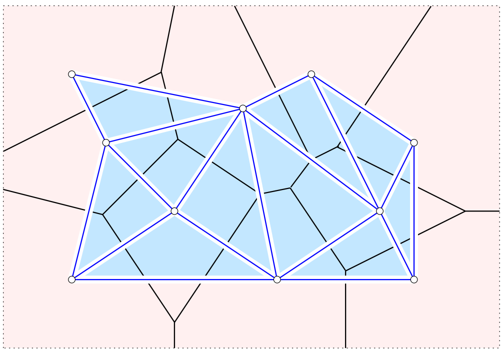

Figure 17: The Voronoi tessellation restricted to the open rectangular region and its dual restricted Delaunay mosaic.

(weighted) Delaunay mosaic is the collection of polytopes spanned by subsets of X that define non-empty Voronoi cells. It is convenient to identify such a subset, Q, with the polytope it

spans, which is the convex hull of the locations of the weighted points in Q. In the assumed generic case, all polytopes are simplices and the Delaunay mosaic is a simplicial complex geometrically realized in $\mathbb{R}^d$, which we denote as $\mathrm{Del}(X)$. Most of the time, we restrict our attention to an open convex region, $\Omega \subseteq \mathbb{R}^d$, we assume $X \subseteq \Omega \times \mathbb{R}$, and we write $\mathrm{dom}(Q, \Omega) = \mathrm{dom}(Q) \cap \Omega$. Correspondingly, the *restricted Voronoi tessellation* and the *restricted Delaunay mosaic* are

$$\mathrm{Vor}(X, \Omega) = \{\mathrm{dom}(x, \Omega) \mid x \in X\}, \tag{27}$$

$$\mathrm{Del}(X, \Omega) = \{Q \subseteq X \mid \mathrm{dom}(Q, \Omega) \neq \emptyset\}; \tag{28}$$

see Figure 17.

Lifting and projecting. The Voronoi tessellation and the Delaunay mosaic can both be constructed as the projection of the boundary complexes of convex polyhedra in $\mathbb{R}^{d+1}$. To explain this, recall that $\varpi(a) = \frac{1}{2}\|a\|^2$ and map every weighted point, x, to the point $\dot{x} \in \mathbb{R}^{d+1}$ and to the affine map $\bar{x} \colon \mathbb{R}^d \to \mathbb{R}$ defined by

$$\dot{x} = (\mathrm{pt}(x), \varpi(\mathrm{pt}(x)) - \tfrac{1}{2}\mathrm{wt}(x)), \tag{29}$$

$$\bar{x}(a) = \varpi(\mathrm{pt}(x)) + \langle \mathrm{pt}(x), a - \mathrm{pt}(x)\rangle + \tfrac{1}{2}\mathrm{wt}(x) \tag{30}$$

$$= \langle \mathrm{pt}(x), a\rangle - \tfrac{1}{2}\|\mathrm{pt}(x)\|^2 + \tfrac{1}{2}\mathrm{wt}(x). \tag{31}$$

The map is chosen so that the solution to $\varpi(a) - \bar{x}(a) = 0$ is the sphere with center $\mathrm{pt}(x)$ and squared radius $\mathrm{wt}(x)$. The point is chosen so that a point on the graph of ϖ lies on or below the graph of $\bar{x}$ iff this point is visible from $\dot{x}$, by which we mean that the entire line segment connecting $\dot{x}$ with this point lies below the graph of ϖ.

Let $\xi \colon \mathbb{R}^d \to \mathbb{R}$ be the pointwise maximum of the affine maps, $\xi(a) = \max_{x \in X} \bar{x}(a)$, and note that it is piecewise linear and convex. As observed already by Georges Voronoi [15], the vertical projection of its linear pieces gives the Voronoi tessellation of X in $\mathbb{R}^d$. To get a similar construction of the Delaunay mosaic, we take the convex hull of the points $\dot{x} \in \mathbb{R}^{d+1}$. We call a hyperplane that touches the polytope without intersecting its interior a *support plane*, and the intersection of the polytope with a support plane a *face* of the polytope. For points in general position, all faces are simplices. A *lower face* is the intersection of the polytope with a non-vertical support plane such that the polytope lies above the hyperplane. In analogy to the relation observed by Voronoi, the vertical projection of the lower faces of the convex hull gives the Delaunay mosaic of X in $\mathbb{R}^d$.

The interpretations of the Voronoi tessellation and the Delaunay mosaic as projections of boundary complexes of convex polyhedra provide geometrically intuitive interpretations of a function that plays a crucial role in this paper. Recall that each simplex, $Q \in \mathrm{Del}(X)$, corresponds to a Voronoi cell, $\mathrm{dom}(Q)$. The *radius function*, or more precisely the half squared radius, $\varrho \colon \mathrm{Del}(X) \to \mathbb{R}$, maps Q to the minimum difference between ϖ and ξ at points in the Voronoi cell:

$$\varrho(Q) = \min_{a \in \mathrm{dom}(Q)} [\varpi(a) - \xi(a)]. \tag{32}$$

In words, $\varrho(Q)$ is the amount we have to lower the graph of ϖ until it intersects the graph of ξ at a point vertically above $\mathrm{dom}(Q)$. The function value is also the minimax difference between ϖ and any affine map that satisfies $\bar{y}(\mathrm{pt}(x)) \leq \varpi(\mathrm{pt}(x)) - \frac{1}{2}\mathrm{wt}(x)$ for all $x \in X$ and with equality for all $x \in Q$. Specifically, we minimize the maximum $\bar{y}(a) - \varpi(a)$, in which the maximization is over all $a \in \mathbb{R}^d$, and the minimization is over all affine maps, $\bar{y} \colon \mathbb{R}^d \to \mathbb{R}$, that satisfy the conditions stated above.

Discrete Morse theory. Assuming general position, the radius function on the Delaunay mosaic enjoys structural properties, which we now formalize. Let K be a simplicial complex and $P, Q \in K$ two simplices. For a *monotonic function*, $f \colon K \to \mathbb{R}$, $P \subseteq Q$ implies $f(P) \leq f(Q)$. The *Hasse diagram* of K is the directed graph whose nodes are the simplices and whose arcs are the codimension 1 face relations: every arc ends at a p-simplex and starts at a $(p-1)$-dimensional face of this simplex. By construction, the values of a monotonic function are non-decreasing along directed paths in the Hasse diagram. A *level set* of f is a maximal collection of simplices with shared function value, $f^{-1}(r) \subseteq K$, and we call a maximal connected subset of a level set a *step*. For simplices $P \subseteq R$ in K, call $\psi = \{Q \in K \mid P \subseteq Q \subseteq R\}$ an *interval*, $P = \mathrm{lb}(\psi)$ its *lower bound*, $R = \mathrm{ub}(\psi)$ its *upper bound*, and note that $\#\psi = 2^{\#R - \#P}$. According to an inessential modification of the original formulation by Forman [13], f is a *discrete Morse function* if every step is an interval of size 1 or 2. A slightly weaker condition was introduced by Freij [14], calling f a *generalized discrete Morse function* if every step is an interval. The corresponding partition of K into intervals is called the *generalized discrete gradient* of f.

The singleton intervals are special, which is expressed by calling the simplices they contain and the corresponding values *critical*. To motivate this terminology, consider two contiguous values, $r < s$, and the corresponding sublevel sets, $K_r = f^{-1}(-\infty, r]$ and $K_s = f^{-1}(-\infty, s]$. By assumption, no simplex maps to a value strictly between r and s, which implies that the difference between the two complexes is the level set at s. This level set is a disjoint union of steps, and because f is generalized discrete Morse, a disjoint union of mutually separated intervals. When we add the simplices of such an interval to K_r, then the homotopy type changes if the interval consists of a single, critical simplex, and it remains unchanged if the interval consists of two or more simplices. The operation of removing a non-singular interval is called a *collapse*. If all intervals in $f^{-1}(s)$ are non-singular, then we write $K_s \searrow K_r$ to express that K_r can be obtained from K_s by collapsing all intervals in the difference. More generally, if $(r, t]$ contains no critical value of f, then $K_t \searrow K_r$; see Forman [13].

B Equivalence of Definitions

This appendix proves that the three definitions of the Wrap complex offered in Section 2 are indeed equivalent. Given a generalized discrete Morse function $f \colon D \to \mathbb{R}$, we recall that $\mathrm{Wrap}_h(f) \subseteq \mathrm{Alpha}_h(f)$ are the Wrap and the Alpha complexes of f for h, and $\mathrm{Sg}_f(h)$ is the collection of singular steps whose critical simplices have function value at most h.

▶ **1** (Wrap Complex Lemma). *Let $f \colon D \to \mathbb{R}$ be a generalized discrete Morse function on a simplicial complex. Then*

(i) $\mathrm{Wrap}_h(f)$ is the smallest complex $K \subseteq D$ that satisfies $\mathrm{Alpha}_h(f) \searrow K$, in which we restrict the collapses to intervals of f.

(ii) $\mathrm{Wrap}_h(f)$ is the smallest subcomplex of D that contains $\bigcup \mathrm{Sg}_f(h)$ and is a union of intervals of f.

Proof. Consider two steps, φ and ψ, in the step graph $\mathcal{G}$ of f. If there is an arc from φ to ψ, then φ contains a proper face of a simplex in ψ. This implies that if M is a collection of steps such that $K = \bigcup M$ is a complex, then $\psi \in M$ implies $\varphi \in M$. If both belong to M, then φ cannot be collapsed. On the other hand, if $\varphi \in M$ and no successor of φ in $\mathcal{G}$ belongs to M, then we can collapse φ; that is: $K \setminus \varphi$ is a complex. To prepare the proofs of (i) and (ii), we let M be a collection of steps such that

1. $K = \bigcup M$ contains a critical simplex iff $f(Q) \leq h$;

2. K is a complex;
3. there is no step $\varphi \in M$ such that $K \searrow K \setminus \varphi$.

First we claim that the three properties specify M uniquely. To prove this claim, let $\varphi_0 \in M$ be non-singular and let $\varphi_0, \varphi_1, \ldots, \varphi_k$ be maximal such that $\varphi_i \in M$ is a successor of φ_{i-1} in $\mathcal{G}$ for $1 \leq i \leq k$. We note that $k \geq 1$ because φ_0 cannot be collapsed, and φ_k is singular because the sequence is maximal. To get a contradiction, we assume that $N \neq M$ is another collection of steps that satisfies Properties 1, 2, 3. Suppose first that N contains a step $\mu_0 \notin M$, and consider a maximal sequence $\mu_0, \mu_1, \ldots, \mu_\ell$ such that $\mu_j \in N$ is a successor of μ_{j-1} in $\mathcal{G}$ for $1 \leq j \leq \ell$. Since $\mu_0 \notin M$, the step is necessarily non-singular, which implies $\ell \geq 1$ and μ_ℓ singular. But then there is a first step along this sequence, μ_j, that belongs to M. Since there is an arc from μ_{j-1} to μ_j and $\mu_{j-1} \notin M$, this contradicts that M is a complex. Suppose second that N contains no such step μ_0, but M contains a step $\varphi_0 \notin N$. By the symmetric argument, this implies that N is not a complex, again a contradiction. We conclude that the collection M that satisfies Properties 1, 2, 3 is unique.

Second we claim that the unique complex that satisfies Properties 1, 2, 3 is $\mathrm{Wrap}_h(f)$. By definition, the Wrap complex contains all critical simplices that satisfy $f(Q) \leq h$. The value of Q is the maximum of any step in the lower set of its singular interval, which implies that $\mathrm{Wrap}_h(f)$ contains no critical simplex with value larger than h and therefore satisfies Property 1. Property 2 is satisfied because all faces of a simplex in a step that are not in the step belong to predecessors of the step. Indeed, the directed path from a face to a simplex in the Hasse diagram maps to a possibly shorter directed path from the step of the face to the step of the simplex in $\mathcal{G}$. To see that Property 3 is satisfied as well, we note that every non-singular step $\varphi_0 \subseteq \mathrm{Wrap}_h(f)$ has a directed path to a singular interval and can therefore not be collapsed. We conclude that $\mathrm{Wrap}_h(f)$ is the only union of steps that satisfies Properties 1, 2, 3.

To prove (i), we note that $\mathrm{Alpha}_h(f)$ satisfies Properties 1 and 2, so it cannot satisfy Property 3 unless it is equal to $\mathrm{Wrap}_h(f)$. We can therefore collapse non-singular intervals. The process must halt, and the only way it can halt is when it reaches the unique union of steps that satisfies Properties 1, 2, 3, which is $\mathrm{Wrap}_h(f)$.

To prove (ii), we observe that $\mathrm{Wrap}_h(f)$ contains $\bigcup \mathrm{Sg}_f(h)$ and is a union of steps. To see that it is the smallest such complex, suppose there is another complex, $L = \bigcup N$, that has this property and there exists a step $\varphi \subseteq \mathrm{Wrap}_h(f) \setminus L$. As argued above, this contradicts that L is a complex, which implies (ii). ◻

C Algorithm for Discrete Gradient

It is easy to see that the rise function defined in Section 3 is monotonic. As proved in [12], it also satisfies the more stringent requirements of a generalized discrete Morse function provided $U \subseteq \Omega$ is in general position. The generalized discrete gradient of this function is a partition of the Delaunay mosaic into intervals, and this partition is instrumental in the construction of subcomplexes discussed in Section 4.

The construction of this partition is complicated by the impossibility of computing the rise function exactly, at least for general Legendre type functions. Given a numerical approximation, $g\colon K \to \mathbb{R}$, our goal is therefore to first recover the generalized discrete Morse function that g approximates. Given a tolerance, $\varepsilon \geq 0$, we give an algorithm that computes such a function $f\colon K \to \mathbb{R}$ with $\|f - g\|_\infty \leq \varepsilon$ and such that the corresponding partition is minimal in a restricted sense. To prepare the algorithm, we define the *gap* of a subset $\varphi \subseteq K$ as the maximum difference of function values:

$$\operatorname{gap}\varphi = \max_{P,Q\in\varphi,\ P\subseteq Q} [g(Q) - g(P)]. \tag{33}$$

If g is monotonic, then all gaps are non-negative. Otherwise, let $-\varepsilon_0$ be the smallest (largest negative) gap between pairs $P \subseteq Q$, set $g(P) = \min\{g(P), g(Q)\}$ whenever $P \subseteq Q$, and note that this makes g monotonic while changing the value of any simplex by at most ε_0. We will therefore assume that g is monotonic. Letting V be a partition of K into intervals, we call an interval $\psi \notin V$ *compatible* with V if

(i) ψ is the union of intervals in V;

(ii) every pair of simplices $P \subseteq Q$ with $P \in \psi$ and $Q \notin \psi$ implies $g(\mathrm{ub}(\psi)) \leq g(Q)$,

in which $\mathrm{ub}(\psi)$ is the upper bound of the interval. The algorithm constructs the discrete gradient of f by adding compatible intervals to an initially trivial partition of K, namely the one in which every simplex belongs to its own set in the partition. The function itself is computed by spreading the function value of the upper bound to the other simplices in the interval. Let $\psi_1, \psi_2, \ldots, \psi_m$ be the collection of all intervals of K, sorted by gap, let $\varepsilon \geq 0$ be a fixed threshold, and initialize i to 1 and V to the trivial partition of K.

```
while i ≤ m and gap ψ_i ≤ ε do
    if ψ_i compatible with V then
        remove all φ ∈ V with φ ⊆ ψ_i from V;
        add ψ_i to V;
        forall P ∈ ψ_i do set f(P) = g(ub(ψ_i));
    i = i + 1.
```

Condition (i) guarantees that the computed V is a partition of K into intervals. Condition (ii) makes sure that no relation is reversed, which implies that the computed function, $f\colon K \to \mathbb{R}$, is monotonic and that V is a refinement of its partition into steps. Finally, $0 \leq f(P) - g(P) \leq \varepsilon$ for all simplices $P \in K$, as claimed. Without assuming that g be monotonic, the upper bound on the distance between the two functions is $\varepsilon + \varepsilon_0$.

A slight improvement of the algorithm takes into account that an interval can change from incompatible to compatible. By keeping track of this property throughout the algorithm, we can add an interval to the partition even after it was rejected earlier.