

Optimal Voronoi Tessellations with Hessian-based Anisotropy

Max Budninskiy
Caltech

Beibei Liu
Caltech

Fernando de Goes
Pixar

Yiying Tong
MSU

Pierre Alliez
Inria

Mathieu Desbrun
Caltech/Inria

Abstract

This paper presents a variational method to generate cell complexes with local anisotropy conforming to the Hessian of any given convex function and for any given local mesh density. Our formulation builds upon approximation theory to offer an anisotropic extension of Centroidal Voronoi Tessellations which can be seen as a dual form of Optimal Delaunay Triangulation. We thus refer to the resulting anisotropic polytopal meshes as *Optimal Voronoi Tessellations*. Our approach sharply contrasts with previous anisotropic versions of Voronoi diagrams as it employs first-type Bregman diagrams, a generalization of power diagrams where sites are augmented with not only a scalar-valued weight but also a vector-valued shift. As such, our OVT meshes contain only convex cells with straight edges, and admit an embedded dual triangulation that is combinatorially-regular. We show the effectiveness of our technique using off-the-shelf computational geometry libraries.

Keywords: Anisotropic meshing, Bregman diagrams, centroidal Voronoi tessellation, optimal Delaunay triangulation.

Concepts: •Mathematics of computing → Mesh generation;
•Computing methodologies → Shape modeling;

1 Introduction

From seismic waves to plasma filamentation and boundary layers in fluids, many physical problems exhibit solutions changing more significantly in one spatial direction than another. In order to reproduce this type of effects numerically, it is often advantageous to use a mesh with elements stretched along the geometry of the solution so that even low order (e.g., linear) basis functions can nicely resolve the physical phenomenon with only a few elements. While many efforts have successfully extended isotropic meshing methods for the construction of anisotropic triangulations [Chen and Xu 2004; Loseille and Alauzet 2009], the generation of anisotropic complexes formed by convex polyhedra has remained a challenge.

In this paper, we introduce a new meshing technique that generates anisotropic cell complexes of arbitrary 2D or 3D domains. At its core, our approach optimizes the placement and shape of cells by best approximating a convex function through linear finite functions over convex polyhedral elements. This formulation leads to a variational method that extends the machinery of Optimal Delaunay Triangulation to cell complexes, resulting in a new anisotropic version of Centroidal Voronoi Tessellations. Our derivation also reveals new degrees of freedom to control the local anisotropy of polyhedral meshes, while still defining cell complexes that are dual to combinatorially-regular triangulations.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM.

SA '16 Technical Papers, December 05 - 08, 2016, Macao

ISBN: 978-1-4503-4514-9/16/12

DOI: <http://dx.doi.org/10.1145/2980179.2980245>




Figure 1: Anisotropic meshing. In this paper, we show that the construction of an optimal piecewise-linear approximation of a function over a cell complex (left) extends the isotropic notion of Centroidal Voronoi Tessellations (CVT, top) to an anisotropic variant (middle & bottom) we call Optimal Voronoi Tessellation (OVT), to stress its duality to Optimal Delaunay Triangulation (ODT). Cell anisotropy (indicated by tightest ellipses) and density are independently controlled, and the dual triangulation based on cell barycenters is embedded and combinatorially-regular.

1.1 Related Work

We begin by briefly reviewing the main meshing approaches developed in scientific computing and geometry processing.

Isotropic meshing. Meshing techniques typically seek the generation of non-degenerate geometric primitives that offer good condition numbers for common discrete isotropic operators like the Laplacian. Delaunay meshes with local refinements (see, e.g., [Shewchuk 1998]) have been shown most effective at generating such isotropic simplicial meshes. To further improve the quality of the resulting meshes, variational approaches were also introduced. For instance, the use of Centroidal Voronoi Tessellations (CVT) was proposed to generate isotropic cell decompositions of a domain [Du et al. 1999]. A number of papers were later dedicated to accelerate the computations involved in forming CVT meshes [Du and Emelianenko 2006; Liu et al. 2009]. However, the isotropy of Voronoi cells does not lead to isotropic Delaunay simplices in 3D, and slivers can and will occur [Alliez et al. 2005]. To overcome this limitation, the concept of Optimal Delaunay Triangulation (ODT) was introduced as a dual version of CVT meshes [Chen and Xu 2004; Chen 2004]. Its implementation in 3D with details on the sizing field computations and boundary handling was later investigated by Alliez et al. [2005], and a hybrid approach mixing Delaunay refinement and ODT optimization was formulated in [Tournois et al. 2009] to accelerate convergence and guarantee high-quality results. Algorithmic and boundary-handling variants were also proposed in [Chen and Holst 2011; Gao et al. 2012; Chen et al. 2014].

Anisotropic meshing. Considerable effort has been devoted to extend isotropic meshing to the anisotropic case, where the tightest ellipsoid of an n -simplex is bounded above and below by the local ellipsoid representing a given metric. In early works, anisotropy was reached via local vertex optimizations and connectivity updates to explore the space of vertex positions and connectivities, but at a prohibitive computational cost; see, e.g., [Bossen and Heckbert 1996]. For 2D simplicial meshing, greedy bisections were also shown to produce optimally adapted triangulations asymptotically [Mirebeau and Cohen 2011]. For 2D and 3D domains, anisotropic generalizations of the Delaunay refinement paradigm were proposed in [Boissonnat et al. 2008; Boissonnat et al. 2014] with local connectivity between vertices derived from the Delaunay property in the prescribed metric, and refinements triggered to both help reconcile the connectivity of each one-ring neighborhood and control local sizes and shapes. Unfortunately, these approaches often require a significant number of refinements, creating oversampled meshes. The methods presented in [George et al. 2002; Frey and Alauzet 2004] create, instead, meshes for which all edges are of unit length in the desired metric by collapsing short edges and splitting long edges. Particle-based methods [Zhong et al. 2013] were also proposed, but they tend to be far from optimal in 3D where well-spaced vertex distributions no longer guarantee good mesh quality.

Anisotropic CVT. In order to incorporate anisotropy to CVT, various definitions of anisotropic Voronoi diagrams have been proposed [Labelle and Shewchuk 2003; Du and Wang 2005; Cheng et al. 2006; Valette et al. 2008], followed by a series of variants using L_p norms [Lévy and Liu 2010], Jensen divergence [Nielsen and Nock 2011], Mahalanobis distances [Richter and Alexa 2015], and even hexagonal metrics [Sun et al. 2011]. However, these extensions of Voronoi diagrams define bisectors between sites by piecewise smooth manifolds, which bring a slew of practical difficulties. Indeed, Boissonnat et al. [2008] showed that these definitions of anisotropic Voronoi diagrams correspond to the restriction and projection to 3D (resp., 2D) of power diagrams in \mathbb{R}^9 (resp., \mathbb{R}^5). Lévy and Bonneel [2013] proposed a simpler variant where the computation of restricted Voronoi diagrams in a higher-dimensional space based on nearest neighbors queries is used as a means to infer connectivity. In both cases, the efficient sweeping or iterative algorithms to compute Voronoi diagrams cannot be directly applied, and planar approximations of bisectors have to be used. Instead, our work introduces a definition of anisotropic cell complexes that guarantees straight bisectors between sites, which can be exactly and efficiently computed in the original low-dimensional space.

Anisotropic ODT. In parallel to the developments we just reviewed, the work of Chen [2004] extended the concept of isotropic ODT to function-adapted Delaunay triangulations, in which the metric is now defined by the *Hessian of a convex function*. An unpublished note from Boissonnat et al. [2006] later proposed to construct these anisotropic ODT meshes using regular triangulations, where a scalar-valued weight per site is used to encode the anisotropy of the mesh simplices. Recently, local adaption to an arbitrary input anisotropic tensor was presented in [Loseille and Alauzet 2009] using local convex functions, and extended to surfaces in [Fu et al. 2014]. Since local adaption prevents a notion of global optimality, these methods require a combinatorial exploration of connectivity via flips. However, the quality of the resulting simplicial meshes far exceeds that of non-variational methods. Similar to the anisotropic ODT case, our work builds upon approximation theory of convex functions to derive a notion of optimal cell complex through a variational formulation—now using functional approximation defined over polyhedral elements instead of over simplices.

Bregman diagrams. Voronoi diagrams can be further generalized based on Bregman divergence, a notion coming from information theory traditionally used as a measure of distance between probabil-

ity distributions. Given a convex continuously-differentiable function f over \mathbb{R}^n , the Bregman divergence of f between two points \mathbf{x} and \mathbf{y} is defined as:

$$D_f(\mathbf{x}, \mathbf{y}) = f(\mathbf{x}) - f(\mathbf{y}) - \nabla f(\mathbf{y}) \cdot (\mathbf{x} - \mathbf{y}). \quad (1)$$

Notice that this divergence is not a proper distance as it is in general not symmetric, i.e., $D_f(\mathbf{x}, \mathbf{y}) \neq D_f(\mathbf{y}, \mathbf{x})$. Yet, Boissonnat et al. [2010] showed that, for any set of sites $\{\mathbf{x}_i\}_{i=1}^N$, the Bregman divergence defines cells of the form:

$$V_i = \{\mathbf{x} \in \mathbb{R}^n \mid D_f(\mathbf{x}, \mathbf{x}_i) \leq D_f(\mathbf{x}, \mathbf{x}_j) \forall j\}. \quad (2)$$

These cells are always polytopal and convex, and form a tiling of space called a *first-type Bregman diagram* [Boissonnat et al. 2010]. We will demonstrate in this work that these Bregman diagrams are particularly relevant for anisotropic meshing.

Compatible primal and dual meshing. Finally, it bears mentioning that meshing is rarely a goal in itself, but a prelude to computations. Many finite-element/finite-volume computational techniques rely not only on the primal discretization of a domain to perform numerical approximations, but also on the presence of an embedded dual mesh. Memari et al. [2012] referred to such a spatial discretization as *primal-dual triangulation*, since they are comprised of a primal triangulation \mathcal{T} and a dual cell complex \mathcal{C} that have *compatible* connectivity, extending the traditional duality between Delaunay triangulations and Voronoi diagrams. Furthermore, they showed that the primal mesh \mathcal{T} of an n -dimensional primal-dual triangulation $(\mathcal{T}, \mathcal{C})$ is necessarily a *combinatorially-regular triangulation* (CRT [Lee 1991]), i.e., a shifted version of a regular triangulation that is an orthogonal dual of \mathcal{C} (see Fig. 3). By relating first-type Bregman diagrams to the extra degrees of freedom offered by primal-dual triangulations, our approach constructs anisotropic cell complexes with an associated embedded primal triangulation, in sharp contrast to most definitions of anisotropic Voronoi diagrams.

1.2 Contributions

While the anisotropic extension of Optimal Delaunay Triangulations has resulted in practical methods to anisotropic simplicial meshing [Chen 2004; Chen et al. 2014], generating *anisotropic cell complexes* of a domain for which elements closely conform to an input metric remains challenging due to the complexity of the current anisotropic extensions to Voronoi diagrams.

In this paper, we present Optimal Voronoi Tessellations (OVT), a new anisotropic extension of Centroidal Voronoi Tessellations based on a dual notion to Optimal Delaunay Triangulations that optimizes the piecewise-linear approximation of convex functions over anisotropic polyhedral elements. Our OVT meshes can thus optimally capture an anisotropy field proportional to the Hessian of any convex function, for arbitrary spatially-varying density functions. Our variational formulation also leads to the construction of first-type Bregman diagrams [Boissonnat et al. 2010], an augmented definition of Voronoi diagrams that assigns to each site a scalar-valued weight and a new vector-valued shift, both derived from an input convex function and a density function. To our knowledge, this approach is the first to exploit Bregman diagrams to directly control anisotropy. Moreover, our scheme satisfies key properties that are absent from previous anisotropic extensions of Voronoi diagrams. In particular, all interior cells are convex with straight bisectors between sites, and the resulting cell complex admits an embedded dual triangulation that is combinatorially-regular, with primal vertices positioned at the centroids of their corresponding dual cells. As a consequence, our method bypasses the need for combinatorial connectivity exploration employed in other anisotropic meshing techniques, thus resulting in an effective solution to polytopal meshing in the restricted case of Hessian-based




Figure 2: Primal interpolation vs. Dual kissing approximation. While ODT interpolates a function from above through simplices (left), OVT is piecewise-linear from below through tangent hyperplanes over convex Bregman cells (right).

anisotropy. Finally, we introduce a practical algorithm that interleaves mesh refinement, site optimization, and connectivity updates, providing controllable density and anisotropy as well as tight approximation bounds for the resulting tessellation.

2 Primer on Approximation Theory

Before delving into our method, we first explain how approximation theory of convex functions has greatly benefited simplicial meshing, and contrast these developments with polyhedral meshing.

2.1 Optimal Delaunay Triangulations

A convex function f over a region $\Omega \subset \mathbb{R}^n$ can be approximated by a piecewise-linear interpolation f_d defined over a triangulation \mathcal{T} tessellating Ω . The approximation error between the functions f and f_d can be quantified in the L^p norm as:

$$Q(f, \mathcal{T}, p) = \|f - f_d\|_{L^p} = \left(\int_{\Omega} |f(\mathbf{x}) - f_d(\mathbf{x})|^p d\mathbf{x} \right)^{1/p} \quad (3)$$

Geometrically, this error Q represents the L^p volume between the graphs of f and f_d in \mathbb{R}^{n+1} , as illustrated in Fig. 2(left). A key result in approximation theory states that the triangulation \mathcal{T} minimizing the approximation error Q for an input convex function f is asymptotically achieved when the anisotropy of the mesh elements matches the local Hessian of f [Nadler 1986; D’Azevedo and Simpson 1989]. This implies that the simplices in \mathcal{T} must be stretched along the eigenvectors of the matrix $\text{Hess}[f]$, with an aspect ratio equal to the square root of the eigenvalues ratio in order to best capture the function f . This result thus indicates that the minimization of Q can drive the generation of isotropic triangulations in the case of $f(x) = |\mathbf{x}|^2$ and, more generally, anisotropic simplicial meshes based on arbitrary convex functions. As we review next, such meshes can be computed by optimizing Q with respect to mesh connectivity and positions of its sites.

Optimizing connectivity. Given a set of sites $X = \{\mathbf{x}_i\}_{i=1}^N$ in \mathbb{R}^n , Chen [2004] showed that the minimizer of Q is the triangulation \mathcal{T} obtained by orthogonally projecting the lower convex hull of the lifted \mathbb{R}^{n+1} points $\{(\mathbf{x}_i, f(\mathbf{x}_i))\}_{i=1}^N$ back onto \mathbb{R}^n . Due to this property, the combinatorial problem of optimizing connectivity is reduced to a simple convex hull algorithm, for which off-the-shelf tools exist. As pointed out in [Boissonnat et al. 2006], the mesh with optimal connectivity can also be computed as a regular triangulation generated by weighted sites $\{(\mathbf{x}_i, w_i = |\mathbf{x}_i|^2 - f(\mathbf{x}_i))\}_{i=1}^N$. Note that, when $f(\mathbf{x}) = |\mathbf{x}|^2$, the weights are zero and the optimal triangulation is Delaunay, hence the term *Optimal Delaunay Triangulation* (ODT). Following [Chen and Xu 2004], we keep the moniker ODT to refer to any *anisotropic* optimal simplicial mesh (i.e., when f is an arbitrary convex function) without adding any qualifiers since ODT meshes are Delaunay in the metric induced by the Hessian of f .

Optimizing sites. For $p = 1$, Chen [2004] also showed that the interpolation error for a convex function f in an n -dimensional triangulation \mathcal{T} with N sites simplifies to:

$$Q(f, \mathcal{T}, 1) = \frac{1}{n+1} \sum_{i=1}^N f(\mathbf{x}_i) |\Omega_i| - \int_{\Omega} f(\mathbf{x}) d\mathbf{x}, \quad (4)$$

where Ω_i denotes the one-ring of n -simplices incident to site i . Therefore, optimal site positions can be computed by equating the gradient of $Q(f, \mathcal{T}, 1)$ to zero, leading to:

$$\nabla_{\mathbf{x}_i} f(\mathbf{x}_i) = - \frac{1}{|\Omega_i|} \sum_{\tau \in \Omega_i} \nabla_{\mathbf{x}_i} |\tau| \left(\sum_{\mathbf{x}_k \in \tau} f(\mathbf{x}_k) \right). \quad (5)$$

Algorithm. Based on these results, an arbitrary simplicial mesh can be optimized into a function-adapted ODT mesh by following a Lloyd-type iterative algorithm that alternates two steps: (1) for fixed site positions, compute the optimal connectivity using a regular triangulation; (2) for fixed connectivity, update the site positions by solving the optimality condition in Eq. (5). At convergence, the primal mesh is a regular triangulation that admits an embedded (and orthogonal) dual cell complex corresponding to the power diagram of the weighted sites $\{(\mathbf{x}_i, w_i)\}_{i=1}^N$.

Sampling conditions. In a follow-up work, Chen et al. [2007] proved that an optimal bound on the L^p interpolation error in \mathbb{R}^n can be obtained if three sampling conditions are verified:

1. variation of the Hessian of f inside each n -simplex is small;
2. n -simplices have near regular shape in the Hessian metric;
3. and n -simplices have near constant volume in the metric $\text{Hess}[f]/[\det \text{Hess}[f]]^{1/(2p+n)}$.

With these requirements, a tight bound on the $L^{p \geq 1}$ interpolation error for a simplicial complex with N sites in \mathbb{R}^n is obtained:

$$\|f - f_d\|_{L^p} \leq CN^{-\frac{2}{n}} \|(\det \text{Hess}[f])^{\frac{1}{n}}\|_{L^{\frac{pn}{2p+n}}}. \quad (6)$$

Extension to arbitrary tensors. Adapting a primal mesh to a convex function is a global optimization problem as we just reviewed. However, one may wish to adapt a mesh to a given *arbitrary tensor field*, indicating the desired anisotropy through its eigenvectors and the mesh density via its determinant (note that this input tensor must be SPD, i.e., symmetric and positive definite). Unless the input tensor field happens to be the Hessian of a function (possibly modulated by a density function ρ), the previous ODT framework does not apply directly. A way to adapt this approximation-theoretical approach to an arbitrary SPD tensor σ was first proposed in [Loiseille and Alauzet 2009] for 2D and 3D flat domains, then extended to surfaces in [Fu et al. 2014]. For each simplex τ , they compute the desired tensor σ_{τ} through a local average of the input tensor field, and derive a local convex function $f_{\tau}(\mathbf{x}) = \frac{1}{2} \mathbf{x}^t \sigma_{\tau} \mathbf{x}$. An element-based residual term $E = \sum_i Q(f_{\tau_i}, \tau_i, 1)$ is introduced to provide a measure of how each local convex function f_{τ} is well approximated by its associated linear interpolant. This objective function is then minimized through both site and connectivity updates (via edge flips in 2D and 2-3, 3-2, 4-4, and 2-2 flips in 3D) to reach a local minimum. Additional checks on edge lengths may trigger refinement or coarsening to enforce that the resulting mesh has approximately unit length in the metric σ . Notice here that this formulation is no longer globally consistent, as it only measures local adaption to local functions. Consequently, mesh connectivity needs to be found through *combinatorial* exploration, instead of using a simple convex hull method as in the convex function case. Moreover, even an element with bounded lengths does not necessarily form a regular shape in the metric σ as it can violate the sampling conditions mentioned previously. Finally, there is no guarantee that the resulting triangulation admits an embedded dual.

2.2 Centroidal Voronoi Tessellations

While primal meshes are often favored in finite-element methods, dual meshes (i.e., polyhedral complexes dual to an embedded triangulation) are in high demand for finite-volume and Petrov-Galerkin methods (see, e.g., [Kuzmin 2010]). Among the various methods for computing dual meshes, Centroidal Voronoi Tessellations (CVT) stand out as minimizers of

$$\mathcal{E}_{\text{CVT}}(X, \mathcal{V}) = \sum_{i=1}^N \int_{V_i} |\mathbf{x} - \mathbf{x}_i|^2 d\mathbf{x}, \quad (7)$$

where \mathcal{V} indicates a cell complex with a cell V_i assigned to each site \mathbf{x}_i [Du et al. 1999]. This objective function is closely related to approximation theory, since it measures the L^1 approximation error between the graph of the canonical function $f(\mathbf{x}) = |\mathbf{x}|^2$ and an underlaid piecewise-linear shape in \mathbb{R}^{n+1} formed by the upper convex envelope of the arrangement of the tangent hyperplanes to the paraboloid at sites \mathbf{x}_i . Therefore, CVT meshes provide a dual concept to *isotropic ODT* meshes that replaces an inscribed interpolation by a circumscribed approximation.

Minimizing \mathcal{E}_{CVT} . As described in [Du et al. 1999], the orthogonal projection of the arrangement of planar facets tangent to the paraboloid at sites creates a cell complex in \mathbb{R}^n equal to the Voronoi diagram of X (Fig. 1(left)). Therefore, the minimization of \mathcal{E}_{CVT} provides a variational definition for the generation of optimal isotropic Voronoi diagrams. Moreover, minimizing \mathcal{E}_{CVT} w.r.t. positions implies that the sites must verify the optimality condition:

$$\nabla_{\mathbf{x}_i} \mathcal{E}_{\text{CVT}} = 2|V_i|(\mathbf{x}_i - \mathbf{b}_i) = 0, \quad (8)$$

where \mathbf{b}_i denotes the barycenter of the associated cell V_i . The optimal positions are thus coincident to cell barycenters, hence the term *Centroidal Voronoi Tessellation* (CVT). Similar to the isotropic ODT, CVT minimizers favor the formation of isotropic cells.

Algorithm. One way to obtain a CVT tessellation from a given set of vertices over a domain Ω is to use a Lloyd-type approach similar to the ODT case, by repeating two steps: (1) computing the Voronoi tessellation for a fixed set of sites; and (2) moving sites to the barycenters of their current Voronoi cells. Several methods have been proposed to accelerate convergence by leveraging the C^2 nature of the objective function [Liu et al. 2009], and incorporate cell size control through a density function $\rho(x)$ (substituting the volume form $d\mathbf{x}$ by $\rho(\mathbf{x})d\mathbf{x}$ in Eq. (7)) [Du et al. 1999; Richter and Alexa 2015]. Note that the approximation error \mathcal{E}_{CVT} and its variants are surprisingly ubiquitous in topics varying from clustering methods in machine learning to information theory.

Anisotropic extensions. Multiple generalizations of Voronoi tessellations to the anisotropic case were proposed in recent years, starting with [Du and Wang 2005]. In contrast to ODT, these extensions do not follow concepts of function approximations, but rely instead on generalizations of the definition of Voronoi diagrams [Labelle and Shewchuk 2003]. Unfortunately, the resulting generalized Voronoi cells have curved bisectors between sites so they can form non-convex cells. Thus, they are more costly to compute, even approximately [Lévy and Bonneel 2013], and not guaranteed to have associated dual triangulations. Instead, our new approach exploits the connection between functional approximation and anisotropy of cells to offer a simple, yet powerful extension.

3 Optimal Voronoi Tessellations

We now describe how the CVT formulation can be modified to offer a new approach to the generation of cell complexes with convex cells and controlled anisotropy, that mimics the approximation-theoretical foundations of ODT applied to the polyhedral setting.

3.1 Definitions

Whereas ODT relied on the approximation properties of piecewise-linear interpolants over simplicial meshes, our goal is to formulate a corresponding picture that offers approximation guarantees on convex polyhedral cells. Towards this end, we first introduce terminology and concepts that will define a finite-dimensional approximation space over which polyhedral meshing can be cast as an optimal piecewise-linear approximation of a convex function.

Bregman diagrams. Let Ω be a domain in \mathbb{R}^n , and $f: \Omega \rightarrow \mathbb{R}$ be an arbitrary convex, twice continuously differentiable function. From a given set of sites $X = \{\mathbf{x}_i\}_{i=1}^N$ in Ω , we can construct a corresponding set of hyperplanes $\{T_i\}_{i=1}^N$ in \mathbb{R}^{n+1} so that $T_i(\mathbf{x})$ is defined as the hyperplane tangent to the graph of $f(\mathbf{x})$ for each site \mathbf{x}_i , i.e.,

$$T_i(\mathbf{x}) = \nabla f(\mathbf{x}_i) \cdot (\mathbf{x} - \mathbf{x}_i) + f(\mathbf{x}_i). \quad (9)$$

The upper envelope of this arrangement of planes, when projected back down to \mathbb{R}^n , forms a tessellation of Ω . More precisely, since tangent hyperplanes are always below the convex function f , the cells of this resulting tessellation are of the form:

$$V_i = \{\mathbf{x} \in \Omega \mid T_i(\mathbf{x}) \geq T_j(\mathbf{x}) \ \forall j\}. \quad (10)$$

Since $\{T_i\}_{i=1}^N$ are hyperplanes, inner boundaries of the resulting cells are straight and inner cells are convex polyhedra. By construction, there are also no orphan cells or hidden sites. Notice that, for the case $f(\mathbf{x}) = |\mathbf{x}|^2$, the cells $\{V_i\}_{i=1}^N$ are simply the Voronoi cells associated to the sites X . Comparing Eq. (10) with Eq. (2), one realizes that such tessellations are first-type Bregman diagrams [Boissonnat et al. 2010]. This fact will soon be shown crucial for their efficient evaluation.

Function approximation. Given sites X and a convex function f , we can leverage the tessellation defined in Eq. (10) and approximate f with a piecewise-linear function over cells:

$$f_d(\mathbf{x}) = \sum_{i=1}^N T_i(\mathbf{x}) \mathbb{1}_{V_i}(\mathbf{x}), \quad (11)$$

where $\mathbb{1}_{V_i}$ is the indicator function of a cell V_i . By construction, the function f_d is a continuous piecewise-linear approximation of f from below, as opposed to the ODT-style piecewise-linear interpolation from above (see Fig. 2). Since this approximant is tangent to the graph of f at each \mathbf{x}_i , it could be argued that it also interpolates the function at each \mathbf{x}_i . We thus refer to an element of this underlaid function space as a *kissing approximant*, to differentiate it from the simplicial interpolation space used in ODT.

Approximation error. We can now evaluate the L^1 approximation error between a convex function f and a kissing approximant f_d defined by the sites X and a cell complex $\mathcal{V} = \{V_i\}_{i=1}^N$:

$$\mathcal{E}_{\text{OVT}}(f, X, \mathcal{V}) = \|f - f_d\|_{L^1} = \sum_{i=1}^N \int_{V_i} (f(\mathbf{x}) - T_i(\mathbf{x})) d\mathbf{x}. \quad (12)$$

Note that, due to Eq. (9), the cells of a Bregman diagram minimize the residual $\mathcal{E}_{\text{OVT}}(f, X, \mathcal{V})$ for a given function f and sites X . Consequently, the optimal connectivity of a function-adapted cell complex is directly deduced from the upper envelope of the sites hyperplanes, similar to the construction of ODT. It is also worth observing that $\mathcal{E}_{\text{OVT}}(|\mathbf{x}|^2, X, \mathcal{V}) \equiv \mathcal{E}_{\text{CVT}}(X, \mathcal{V})$, thus generalizing the CVT energy to arbitrary convex functions.

Cell anisotropy. Simplifying Eq. (12) using Taylor expansion reveals the anisotropic nature of our formulation:

$$\begin{aligned} \mathcal{E}_{\text{OVT}}(f, X, \mathcal{V}) &= \sum_{i=1}^N \int_{V_i} (f(\mathbf{x}) - T_i(\mathbf{x})) d\mathbf{x} \\ &= \frac{1}{2} \sum_{i=1}^N \int_{V_i} |\mathbf{x} - \mathbf{x}_i|_{\text{Hess}[f](\mathbf{x}_i)}^2 d\mathbf{x} + \mathcal{O}([\text{diam } V_i]^{n+3}), \end{aligned} \quad (13)$$

where $\text{diam}(V_i)$ is the maximum width of cell V_i , and $\text{Hess}[f](\mathbf{x}_i)$ is the Hessian of f evaluated at \mathbf{x}_i . Notice that the first term is akin to a CVT energy in the metric induced by the Hessian of f . Moreover, since the first term dominates asymptotically, optimal cells will be stretched in the direction of the eigenvector corresponding to the smaller eigenvalue of the Hessian, with an anisotropy equal to the square root of the ratios of eigenvalues. We thus refer to minimizers of \mathcal{E}_{OVT} as *Optimal Voronoi Tessellations* (OVT) to stress their duality with ODT [Chen et al. 2007]: asymptotically, OVT meshes have Voronoi cells in the metric $\text{Hess}[f]$.

3.2 Variational Formulation

So far our approach to generate function-adapted cells in \mathbb{R}^n relied on the notion of Bregman diagrams and the computation of the upper envelope of an arrangement of hyperplanes in \mathbb{R}^{n+1} . Next we employ a simple change of variables introduced in [Boissonnat et al. 2010] that simplifies the construction of a Bregman diagram down to a power diagram. Importantly, this interpretation suggests a Lloyd-based optimization procedure and unveils the existence of a dual (combinatorially-regular) triangulation for each OVT mesh.

Optimizing connectivity. Given a convex function f and sites X , we can rewrite Eq. (12) by introducing auxiliary variables. To each site \mathbf{x}_i , we associate a pair of variables (\mathbf{p}_i, w_i) , where \mathbf{p}_i is a *shifted* site in \mathbb{R}^n and w_i is a scalar weight defined by:

$$\mathbf{p}_i = \frac{1}{2} \nabla f(\mathbf{x}_i), \quad w_i = \frac{1}{4} |\nabla f(\mathbf{x}_i)|^2 + f(\mathbf{x}_i) - \nabla f(\mathbf{x}_i) \cdot \mathbf{x}_i. \quad (14)$$

With these variables, the tangent hyperplanes T_i of f at \mathbf{x}_i defined in Eq. (9) can be rewritten as

$$T_i(\mathbf{x}) = 2\mathbf{p}_i \cdot \mathbf{x} - |\mathbf{p}_i|^2 + w_i. \quad (15)$$

The approximation error \mathcal{E}_{OVT} in Eq. (12) is then reexpressed as:

$$\begin{aligned} \mathcal{E}_{\text{OVT}}(f, X, \mathcal{V}) &= \sum_i \int_{V_i} (f(\mathbf{x}) - T_i(\mathbf{x})) d\mathbf{x} \\ &= \sum_i \int_{V_i} (|\mathbf{x}|^2 - T_i(\mathbf{x})) d\mathbf{x} + \int_{\Omega} (f(\mathbf{x}) - |\mathbf{x}|^2) d\mathbf{x} \\ &= \sum_i \int_{V_i} (|\mathbf{x} - \mathbf{p}_i|^2 - w_i) d\mathbf{x} + \text{constant}. \end{aligned} \quad (16)$$

Up to a constant depending on the function f and the domain Ω , the approximation error \mathcal{E}_{OVT} is reminiscent of a CVT energy, but with two key differences: first, it uses an additional weight per site; second, it is not based on the squared distances to each site \mathbf{x}_i , but to the shifted site \mathbf{p}_i . Eq. (16) thus implies that for a given set of sites X , the optimal f -adapted tessellation \mathcal{V} simply is the *power diagram* of the weighted sites $\{(\mathbf{p}_i, w_i)\}_{i=1}^N$.

Optimizing sites. Our auxiliary variables turn the upper envelope computation in \mathbb{R}^{n+1} described in Sec. 3.1 into a power diagram construction, for which robust and scalable implementations exist in computational geometry libraries [CGAL 2016]. Finding optimal sites X requires minimizing \mathcal{E}_{OVT} with respect to X . Invoking Reynold’s transport theorem [Marsden and Tromba 2003], we get

$$\nabla_{\mathbf{x}_i} \mathcal{E}_{\text{OVT}} = - \int_{V_i} \nabla_{\mathbf{x}_i} T_i(\mathbf{x}) d\mathbf{x} \stackrel{\text{Eq. (9)}}{=} |V_i| \text{Hess}[f](\mathbf{x}_i) (\mathbf{x}_i - \mathbf{b}_i), \quad (17)$$

where $\mathbf{b}_i = \int_{V_i} \mathbf{x} d\mathbf{x} / |V_i|$ is the barycenter of cell V_i . Therefore, the optimal position for each site \mathbf{x}_i coincides with the barycenter of its respective cell. A tessellation minimizing Eq. (12) is thus a *centroidal* Bregman diagram [Boissonnat et al. 2010], confirming that OVT meshes are anisotropic generalizations of CVT.

Algorithm. A Lloyd-based iterative scheme can once again be used to generate optimal cell complexes. More concretely, from an initial set of sites and a convex function f , we repeat the following two steps until convergence: (1) derive a power diagram from the weighted and shifted sites (\mathbf{p}_i, w_i) computed through Eq. (14); and (2) update the sites \mathbf{x}_i by moving them to the barycenter of their own dual cell. Note that each step decreases the objective function \mathcal{E}_{OVT} by construction. The only two major differences with ODT are the use of *shifted* sites \mathbf{p}_i , and the resulting anisotropy of *cells*, rather than *simplices*.

Associated primal triangulations. Since we construct connectivity through a power diagram, there exists an *orthogonal dual* to this diagram, called the regular (or weighted Delaunay) triangulation \mathcal{R} of the weighted points $\{(\mathbf{p}_i, w_i)\}_{i=1}^N$. The connectivity between these weighted sites directly infers a connectivity \mathcal{T} on the original sites \mathbf{x}_i as well. Since \mathcal{R} and \mathcal{T} share the same connectivity and \mathcal{R} is regular, \mathcal{T} is by definition a combinatorially-regular triangulation [Lee 1991]. Moreover, the shifted sites \mathbf{p}_i derive from the original sites \mathbf{x}_i by a deformation with a locally positive definite Jacobian asymptotically since $\partial \mathbf{p}_i / \partial \mathbf{x}_i = \frac{1}{2} \text{Hess}[f](\mathbf{x}_i)$. As the regular triangulation based on the weighted sites is guaranteed to be embedded and the Hessian of f is positive definite, the mesh \mathcal{T} is an embedded triangulation. We thus have constructed a primal-dual triangulation [Memari et al. 2012]. Note that this property is missing from existing anisotropic Voronoi diagrams [Labelle and Shewchuk 2003; Du and Wang 2005].

3.3 Practical conditions of optimality

While Eqs. (16) and (17) provide *necessary* conditions on the connectivity and the positions of sites with respect to their associated cells, they are not sufficient to ensure proper adaption of the convex cell to the function f : sites may be trapped in suboptimal positions corresponding to saddle points and local extrema of \mathcal{E}_{OVT} . Consequently, we need additional criteria that guide and accelerate the optimization procedure to make sure the current tessellation is not far from optimal. For convenience, we denote the *average* value of the Hessian of f within a cell V as \bar{H}_V , that is,

$$\bar{H}_V = \frac{1}{|V|} \int_V \text{Hess}[f](\mathbf{x}) d\mathbf{x}.$$

Sampling conditions. We establish a few crucial conditions on the cells V_i to certify they are well adapted to the input function f . We will show that these sampling conditions offer a tight bound on the approximation error between the piecewise-linear kissing approximation f_d defined over a cell complex and the continuous function f , thus providing the counterpart of the error bound analysis of ODT meshes [Chen et al. 2007] now applied to cell complexes. A




Figure 3: Triangulations dual to OVT. An Optimal Voronoi Tessellation (left) is dual, by construction, to a regular triangulation \mathcal{R} (center, zoomed-out) formed by the (shifted) weighted points $\{(\mathbf{p}_i, w_i)\}_{i=1}^N$. The connectivity of \mathcal{R} can be transferred to the original (centroidal) sites $\{\mathbf{x}_i\}_{i=1}^N$, creating an embedded combinatorially-regular triangulation \mathcal{T} (right).

complete derivation for $L^{p \geq 1}$ is offered in the Supplemental Material; we only state here the resulting sampling conditions for $p = 1$ for brevity:

1. all the cells are star-shaped with respect to their sites.
2. the variations of $\text{Hess}[f]$ are uniformly bounded within each cell, i.e., $\exists \alpha_0, \alpha_1 > 0$ such that $\forall V_i, \forall \mathbf{x} \in V_i$,

$$\alpha_0 \bar{H}_{V_i} \preceq \text{Hess}[f](\mathbf{x}) \preceq \alpha_1 \bar{H}_{V_i} \quad (18)$$

where $A \preceq B$ iff $(A - B)$ is negative semi-definite.

3. cells are close to being spherical in the Hessian metric, i.e., $\exists \beta_0 > 0$ such that $\forall V_i$,

$$\max_{\mathbf{x}, \mathbf{y} \in V_i} \left[\sqrt{(\mathbf{x} - \mathbf{y})^t \bar{H}_{V_i} (\mathbf{x} - \mathbf{y})} \right] \left[\sqrt{\det \bar{H}_{V_i} |V_i|} \right]^{-\frac{1}{n}} \leq \beta_0. \quad (19)$$

4. all cell volumes in the $(\det \bar{H}_{V_i})^{\frac{1}{n+2}} \bar{H}_{V_i}$ metric are close to equal, i.e., $\exists \beta_1 > 0$ such that $\forall V_i \in \mathcal{V}_X \subset \mathcal{V}_f$

$$\frac{\max_i \left[(\det \bar{H}_{V_i})^{\frac{1}{n+2}} |V_i| \right]}{\min_i \left[(\det \bar{H}_{V_i})^{\frac{1}{n+2}} |V_i| \right]} \leq \beta_1. \quad (20)$$

Condition 1 is automatically enforced for interior cells since, by construction, our power cells are convex; however, we add this condition to ensure that our bound applies properly to boundary cells as well (see Sec. 4 for details on boundary treatment). Condition 2 ensures that cells are not covering large variations of f , while Condition 3 enforces shape regularity in the metric $\text{Hess}[f]$ by imposing that cell diameters are close to cell volumes to the power $1/n$, i.e., that each cell is nearly round in the \bar{H}_V metric.

Connection to Gershov's conjecture. Finally, Condition 4 guarantees that the cell volume V_i is proportional to a (dimension-dependent) power of the determinant of the local Hessian of f :

$$|V_i| \sim (\det \bar{H}_{V_i})^{-\frac{1}{n+2}}. \quad (21)$$

While this requirement may not look intuitive at first, one can derive this property as follows. Given that our objective function \mathcal{E}_{OVT} is, to first order, an anisotropic CVT energy (Eq. (13)), we can invoke Gershov's conjecture (see, e.g., [Du et al. 1999]) to deduce that the integrals $\int_{V_i} \|\mathbf{x} - \mathbf{x}_i\|_{\text{Hess}[f](\mathbf{x}_i)}^2 dx$ over OVT cells are asymptotically all equal. If we denote by $h(\mathbf{x})$ the local sizing field of a given OVT (controlling the local average distance from a cell barycenter to its boundaries), the aforementioned integral is thus asymptotically converging to $h^2_{\text{Hess}[f](\mathbf{x})} |V| \sim h^2(\mathbf{x}) (\det \text{Hess}[f](\mathbf{x}))^{1/n} |V|$. Moreover, the local Euclidean volume of a cell in \mathbb{R}^n is proportional to h^n , so the cell volume must indeed satisfy Eq. (21). This implies that the local cell density for a f -adapted polyhedral mesh is completely determined—but we will introduce in Sec. 3.4 an approach that alters this behavior to enforce any user-input cell density.

Approximation error bound. If these sampling conditions hold, we can derive an approximation error bound by first taking a Taylor expansion of f with second order remainder within each cell, then computing the L^1 norm of the error, to conclude that:

Theorem: *If a cell complex with N cells satisfies the four conditions (1)-(4) for a convex C^2 function f , the corresponding piecewise-linear kissing approximation f_a of f satisfies:*

$$\|f - f_a\|_{L^1} \leq CN^{-\frac{2}{n}} \|(\det \text{Hess}[f])^{\frac{1}{n}}\|_{L^{\frac{n}{n+2}}} \quad (22)$$

where C depends on $\alpha_0, \alpha_1, \beta_0, \beta_1$ and the dimension n .

For a complete proof (which holds for $L^p, 1 \leq p \leq \infty$), refer to the Supplementary Material. Our error bound thus coincides with the ODT optimal bound given in [Chen et al. 2007] (see Eq. (6)), reinforcing our claim of duality between our formulation and ODT.

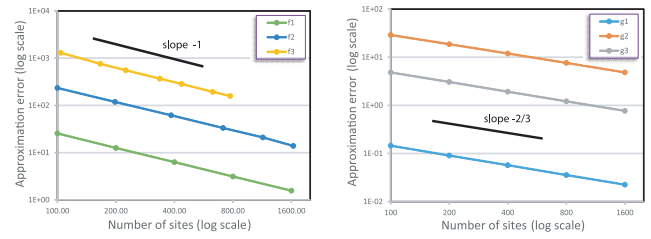


Figure 4: Approximation error rates. As predicted by Eq. (22), the approximation error of an OVT in \mathbb{R}^n behaves as $N^{-2/n}$. The plots were done for $f_1(\mathbf{x}) = 25x^2 + y^2$, $f_2(\mathbf{x}) = x^4 + y^4 + (x+3)^2 + (y+3)^2$, $f_3(\mathbf{x}) = x^2 + 10^{-5}y^2 + y^4$; $g_1(\mathbf{x}) = 4x^2 + y^2 + z^2$, $g_2(\mathbf{x}) = x^4 + y^4 + z^4 + (x-3)^2 + (y-3)^2 + (z-3)^2$, $g_3(\mathbf{x}) = x^2 + y^2 + z^4$.




Figure 5: Density control. A given convex function (here, $f(\mathbf{x}) = -\log x - \log y$) determines both anisotropy and size of the OVT cells (left). However, by imposing a sizing field h as a density modulation, we can adapt cell volumes to an arbitrary local feature size (center) or even force them to have equal volumes (right), without affecting the local anisotropy of the cell decomposition.

3.4 Density control

Mesh density control is an important part of any meshing algorithm: being able to control where the mesh needs to be finer or coarser is crucial in many applications, such as fluid animation (to handle boundary layers), elasticity (to handle various stiffness coefficients in heterogeneous objects), or even geometry processing (to reduce site count by focusing only on targeted regions). Meshing traditionally relies on an input *sizing field* $h(\mathbf{x})$ defined over the domain to which the relative size of mesh elements must conform. We can modify our approach to offer density control as well.

Density-modulated L^1 norm. While adapting a mesh to a convex function f completely determines both the anisotropy of the resulting polyhedral elements (matching the eigenstructure of the Hessian of f) and the mesh density (as stated by Eq. (21)), we can further modify our approach to add a positive scalar field $\rho: \Omega \mapsto \mathbb{R}^+$ to modulate the L^1 norm. This amounts to adopting an L^1_ρ norm in the definition of \mathcal{E}_{OVT} in Eq. (12) by substituting a modulated local volume form $\rho(x)dx$ for the original dx . That is, we now compute a ρ -modulated f -adapted cell complex as a minimizer of:

$$\mathcal{E}_{\rho\text{-OVT}}(f, X, \mathcal{V}) = \|f - f_a\|_{L^1_\rho} = \int_{\Omega} \rho(\mathbf{x}) (f(\mathbf{x}) - f_a(\mathbf{x})) d\mathbf{x}. \quad (23)$$

Note that this extension is in line with the density-modulated ODT introduced in [Chen and Xu 2004] and further exploited in [Chen et al. 2014]. Additionally, modifying the volume form does *not* affect the anisotropic properties of formulation: it only acts as a multiplier to the Hessian in Eq. (12), and thus does not alter eigenvectors or ratios of eigenvalues.

Density-modulated OVT. The use of a density-modulated OVT modifies only a few of the statements made previously:

- a. Barycenters are now defined as $\mathbf{b}_i = 1/|V_i|_\rho \int_{V_i} \mathbf{x} \rho(\mathbf{x}) dx$ where $|V_i|_\rho = \int_{V_i} \rho(\mathbf{x}) dx$ is the modulated volume of a cell.

b. Defining the average modulation $\bar{\rho}_{V_i}$ within a cell, the fourth sampling condition must be rewritten as:

$$\frac{\max_i \left[(\bar{\rho}_{V_i}^n \det \bar{H}_{V_i})^{\frac{1}{n+2}} |V_i| \right]}{\min_i \left[(\bar{\rho}_{V_i}^n \det \bar{H}_{V_i})^{\frac{1}{n+2}} |V_i| \right]} \leq \beta_1, \quad (24)$$

since Gershgorin's conjecture now imposes:

$$|V_i| \sim (\bar{\rho}_{V_i}^n \det \bar{H}_{V_i})^{-\frac{1}{n+2}}. \quad (25)$$

c. The error bound given in Eq. (22) is now updated to involve $\|\rho \cdot (\det \text{Hess}[f])^{1/n}\|_{L^{\frac{n}{n+2}}}$.

Local sizing field enforcement. Given that the local Euclidean volume of a f -adapted cell is of the form of Eq. (25), we can directly guarantee that our cell complexes closely follow a given (strictly-positive) spatially-varying sizing field $h(\mathbf{x})$ by setting our density modulation function ρ as:

$$\rho(\mathbf{x}) = (\det \text{Hess}[f](\mathbf{x}))^{-1/n} |h(\mathbf{x})|^{-(n+2)}. \quad (26)$$

This modulation ρ will compensate for the Hessian influence on the size of mesh elements in order to obtain the user-prescribed mesh density function h , as illustrated in Fig. 5.

3.5 Discussion

Our variational approach based on convex function approximation and density control mirrors quite closely the structure, properties, and algorithmic simplicity of the classical CVT [Du et al. 1999] and ODT [Chen and Xu 2004] techniques. However, it now produces function-adapted anisotropic cell complexes with their associated embedded combinatorially-regular triangulation, without resorting to non-straight Voronoi diagram definitions of [Labelle and Shewchuk 2003; Du and Wang 2005]. While we are limited to anisotropic fields of the type $\rho \text{Hess}[f]$, more general solutions require expensive connectivity exploration [Loseille and Alauzet 2011] or the construction and projection of restricted Voronoi diagrams in higher-dimensions [Lévy and Bonneel 2013], which degrade computational efficiency. Since our OVT formulation involves power diagrams of *auxiliary* (weighted) sites $\{\mathbf{p}_i\}_{i=1}^N$ in \mathbb{R}^n (Eq. (14)), it cannot be seen as a variant or special case of other existing meshing approaches based on Voronoi or power diagrams [Du et al. 1999], even if extra coordinates are introduced [Lévy and Bonneel 2013; Dassi et al. 2015]: edges of the dual triangulation are not orthogonal to the cell facets, in contrast to the case of regular triangulations and power diagrams. Instead, our approach leverages first-type Bregman diagrams [Boissonnat et al. 2010] to derive scalar-valued weights $\{w_i\}_{i=1}^N$ and vector-valued shifts $\{\mathbf{p}_i - \mathbf{x}_i\}_{i=1}^N$ that define a function-adapted anisotropic primal-dual triangulation [Memari et al. 2012].

4 Implementation

Armed with these theoretical foundations, we now describe the implementation of our anisotropic meshing method. From an input convex function f and a sizing field h , we use a strategy similar to Tournois et al. [2009]: we interleave local enforcement of sampling conditions with connectivity and site updates to improve convergence rates and guarantee optimality of the results. We use the CGAL library [CGAL 2016] for all computational-geometric operations. Pseudocode is given in Alg. 1.

Initialization and data structures. From the input surface mesh describing the boundary of Ω , we begin with the construction of a *background mesh* through Delaunay refinement of a constrained 3D Delaunay mesh of the domain. Refinements are driven by the scalar fields f and h so that the resulting interior mesh nodes provide quadrature samples (i.e., the background mesh is made denser

where the gradient of f is large). Besides the values of both f and the modulating function ρ (found from h and f using Eq. (23)), values of the gradient and the Hessian of f (computed via finite differences) are also stored on the nodes. This background mesh is then used to localize and interpolate any query for $\rho(\mathbf{x})$, $f(\mathbf{x})$, $\nabla f(\mathbf{x})$ or $\text{Hess}[f](\mathbf{x})$. Finally, we pick the initial sites \mathbf{x}_i to be a subset of the background mesh nodes to start from site positions that roughly captures the right density h .

Lloyd optimization and refinements. Optimizing the sites via Lloyd iterations is directly implemented by alternating between site and connectivity updates. Fig. 6 shows the improvements on the anisotropy and size of the cells as the number of Lloyd iterations increase. As discussed in Sec. 3.3, such a Lloyd optimization steps can be interspersed with local refinement. In practice, we thus alternate between a round of 10 site and connectivity updates (i.e., we use MaxIterations=10 in Alg. 1) and a round of enforcement of sampling conditions, until convergence. Condition 2 (Eq. (18)) is enforced indirectly by the following simple computational test per cell V_i : $\max_{\mathbf{x}, \mathbf{y} \in V_i} \|\text{Hess}[f](\mathbf{x}) - \text{Hess}[f](\mathbf{y})\|_F \leq \gamma$, where all pairs (\mathbf{x}, \mathbf{y}) of background mesh nodes in V_i are used in the computation of the max, and where γ is set to 5% of the (precomputed) total variation of the Hessian over the entire domain Ω . For Condition 3 (Eq. (19)), we use $\beta_0 = 1.6$, which corresponds to a rectangular shape of ratio 2 : 1 to guarantee that the mesh elements in the metric induced by the Hessian of f are close to isotropic. Similarly, we set $\beta_1 = 2$ in Condition 4 (Eq. (20)) to make sure that volumes of the cells are all within a factor of 2 of each other. If any of the conditions is not met, the offending cell is pushed into a list of cells to be split. Splitting is then performed by inserting a new site next to each offender. Observe that right after the split, newly created cells will inevitably have shapes violating Condition 3, however, the next optimization phase will automatically and quickly reposition them appropriately. Coarsening via deletion can also be added, but using only our progressive addition of sites has been shown effective in all our tests.

Newton-based optimization. Our objective function can be more efficiently minimized through Newton steps [Liu et al. 2009] for which expressions of both gradient and Hessian of \mathcal{E}_{OVT} w.r.t. \mathbf{x}_i are given in the Supplemental Material. In our implementation we revert to LBFGS iterations after a few rounds of Lloyd's iterations for efficiency, and use an energy-based line search. Following [Liu et al. 2009], we use the seven previous gradients to approximate the Hessian without ever resetting the approximation to be most efficient in practice. Fig. 7 depicts a typical behavior of the error reduction as a function of time, showing that once we have reached a decent energy reduction after 12 Lloyd iterations, quasi-Newton steps reduce the approximation error much more efficiently.

Boundary handling. Meshing requires careful handling of boundaries. In fact, most variational meshing approaches exhibit the worst elements near the border as the presence of domain bound-

Algorithm 1 Optimal Voronoi Tessellation

Input: Domain Ω , convex function f , sizing field h .

Initialize sites \mathbf{x}_i

repeat

repeat // *Optimize*

 Compute weighted points (\mathbf{p}_i, w_i) via Eq. (14)

 Generate cells \mathcal{V} by computing their power diagram

 Update sites \mathbf{x}_i via Eq. (17)

until Convergence or MaxIterations reached

 // *Refine if needed*

 Split cells violating any of the sampling conditions (Sec. 3.3)

until Sampling and Convergence conditions met

Output: Anisotropic cells \mathcal{V} and respective sites X .




Figure 6: Optimization via Lloyd iterations. For a constant anisotropy of 8:1 and a radial density modulation, an initial placement of sites generates a random distribution of cells with approximately correct density (left). After one, five, then 100 Lloyd iterations, the cells get aligned to the expected anisotropy and with the proper density, as demonstrated by displaying the tightest ellipse for each cell (right).




Figure 7: Lloyd vs. Quasi-Newton. Plot of approximation error for the function $x^2 + 25y^2$ with 2000 sites vs. computation time, where the L-BFGS optimization starts after 12 Lloyd iterations.

aries adds competing constraints to accommodate. This is particularly true for ODT, where several boundary treatments have been proposed to improve site placement on the boundary [Alliez et al. 2005; Chen and Holst 2011; Gao et al. 2012]. However, boundary handling for cell complexes turns out to be significantly simpler than for the simplicial case: based on Eq. (10), boundary cells simply need to be clipped by the boundary of the domain Ω , and all other results remain intact. This is in sharp contrast to ODT, where site updates/projections at the boundary are necessary to prevent shrinking. For efficient clipping we use the 3D halfspace intersection algorithm from the “Voronoi Covariance Measure” component of CGAL [2016], but alternatives such as [Rycroft 2009; Yan et al. 2013] can also be used. Note, however, that boundary cells may still end up being not connected or having their barycenters outside of the domain if cell sizes are not adapted to the domain. While this does not pose a problem from an approximation theory viewpoint, such situations are not always desirable in practice. In that case, one can project the sites lying outside the domain back onto $\partial\Omega$ through a Euclidean- or Hessian-based projection at each step. Enforcing Condition 1 in Sec. 3.3 also provides a means to remove these cases through refinement, and we end up with optimal cells at the boundary that are simply-connected and star-shaped with respect to their barycenters. The use of a Lipschitz sizing field μ_{fs} computed through local feature size [Alliez et al. 2005] is yet another approach to automatically guarantee geometrically adapted cells at the boundary as we will discuss in Sec. 5, see Fig. 5. The combinatorially-regular triangulation \mathcal{T} also requires cleanup at the boundary, where primal $(n-1)$ -simplices that correspond to dual edges situated entirely outside the domain must be removed.

5 Results

We now present a series of tests to validate both our theoretical results and our implementation choices.

Anisotropy and sizing. We first tested how well our approach controls the resulting anisotropy and size of cells. An optimization procedure for the case of constant anisotropy and space varying radial density is shown in Fig. 6. The Lloyd algorithm reaches the desired density and anisotropy in as few as 5 iterations, and the quality improves further with each additional iteration (see energy

plot in Fig. 7). Additionally, our method can easily adapt cells to quickly varying anisotropy: as shown in Fig. 9: the sharp transition between regions with drastically different anisotropy is well captured by the shape of the cells in the resulting meshes. To evaluate the robustness of our optimization procedure, we also show results for increasing anisotropy but identical sizing in Fig. 8. As depicted by the tightest ellipses bounding the cells, the resulting anisotropic mesh follows the sampling density ratio along horizontal and vertical directions, even when the ratio is above 30. Our experiments were successful with anisotropic aspect ratios up to 200, but meshes become visually difficult to evaluate. Finally, all our examples have, by construction, an associated dual triangulation \mathcal{T} that is embedded with vertices on cell barycenters.

Quality measures. To illustrate the evolution of mesh quality during optimization, we tracked three key measures of the cells V_i :

- Hessian variation $\max_{\mathbf{x}, \mathbf{y} \in V_i} \|\text{Hess}[f](\mathbf{x}) - \text{Hess}[f](\mathbf{y})\|_F$,
- shape ratio $\max_{\mathbf{x}, \mathbf{y} \in V_i} \left[\sqrt{(\mathbf{x} - \mathbf{y})^t \bar{H}_{V_i} (\mathbf{x} - \mathbf{y})} \right] \left[\sqrt{|\det \bar{H}_{V_i}|} |V_i| \right]^{-\frac{1}{n}}$,
- and modified area $(\bar{\rho}_{V_i}^n \det \bar{H}_{V_i})^{\frac{1}{n+2}} |V_i|$.

By binning these measures by occurrence during Lloyd iterations, we show in Fig. 13 the behavior of the resulting distributions corresponding to the sampling criteria (1)-(3) described in Sec. 3.3. As expected, the distributions steadily shift to the left and become more peaked, demonstrating that the optimization procedure adapts the cell complex to the predicted anisotropy and sizing. However, the anisotropy may not be captured perfectly when an insufficient number of sites is used: as mentioned earlier, optimization can locally get stuck in saddle points and local extrema of the objective function \mathcal{E}_{OVT} . Convergence speed, quality of functional approximation, and cell anisotropy and sizing are thus improved by interleaving rounds of optimization and refinement as described in Alg. 1. Moreover, quasi-Newton optimization leverages the smoothness of the energy function and typically achieves in 25 steps what 200+ Lloyd iterations would take when used for large site counts, at a third of the total computation cost—see Fig. 7. Unsurprisingly, this speedup can be less pronounced if the metric contains sudden variations of anisotropy (Fig. 13) as the smoothness of the energy is affected. We also tested the convergence rates of the approximation error for a same domain with an increasing number of vertices N . The bound in Eq. (22) indicates an error inversely proportional to $N^{2/n}$, and our numerical tests confirms this rate as exhibited in log-log plots of the error w.r.t N in Fig. 4.

3D tests. By design, our algorithm extends naturally to 3D meshing and offers the same control over density and anisotropy. Figs. 10 and 11 show how our method can go from the typical isotropic CVT mesh for $f(\mathbf{x}) = |\mathbf{x}|^2$ (left) to anisotropic cell complexes (with aspect ratio 2, 5, and 10 respectively) on a simple spherical domain: the cutaway views show that the desired aspect ratio is

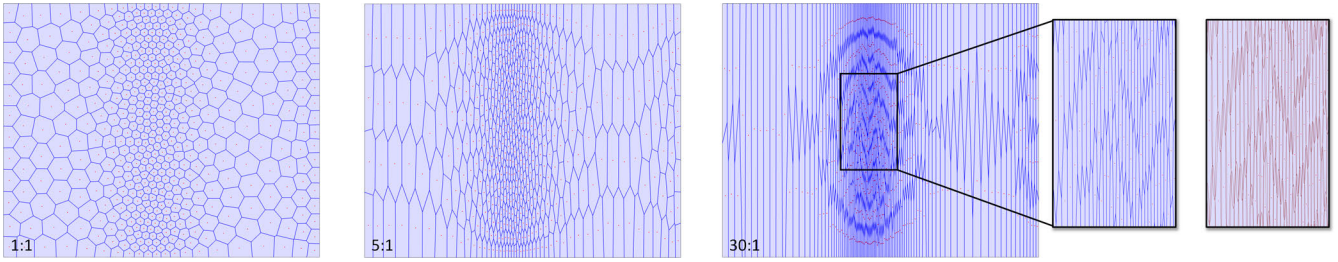


Figure 8: Control over 2D anisotropy. For a given sizing field, we can independently control the anisotropy of an OVT mesh. From left to right: isotropic (CVT), 5:1 and 30:1 (closeups, with tightest ellipses depicted in red).




Figure 9: Sharp anisotropy change. Optimal Voronoi Tessellation for $f(\mathbf{x}) = x^2 + 10^{-5}y^2 + y^4$ (left) over a disk of radius 0.25 centered at the origin produces a sudden change of anisotropy (center), which is well captured even for a relatively small number of sites (right, closeup with cells in blue and tightest ellipses in red).

achieved. Fig. 12 shows spatially-varying anisotropy instead (using a 3D extension to the anisotropy tensor used in Fig. 13). Once again, control over anisotropy is effective: the cell shapes follow the designated anisotropy metric without any user intervention.

Design of domain-adapted anisotropy and density. For density adaptation, we found that the sizing field construction of Alliez et al. [2005] (which computes a maximal K -Lipschitz function $\mu_{\text{Lfs}}(\mathbf{x})$ that does not exceed the *local feature size* (lfs) on the boundary of the domain; see their Eq. (9)) is a very convenient way to adapt the local volume of cells to the geometry of the domain, see Fig. 5. The right choice of sizing field is, however, very application-dependent. The design of an anisotropic function is also entirely dependent on the application, and a function can always be constructed so that its Hessian best fits, in the Frobenius norm, a given metric. If the anisotropy is prescribed only at a few locations in the domain, quadratic programming can be used to find the smoothest function f with proper Hessian: the condition of convexity is a simple *linear* constraint per $(n-1)$ -simplex. The user can thus design very anisotropic 2D and 3D meshes which, once used in a simulation, will offer controlled, realistic tearing and cracking patterns—for instance along fiber directions for wooden structures (Fig. 10).

Timing. As a variational method, our algorithm strives to achieve high quality anisotropic meshes at the cost of longer computation time. However, our algorithm is strictly equivalent to existing CVT




Figure 10: Control over 3D anisotropy. Optimal Voronoi Tessellations of a sphere: isotropic (CVT, inset), 2:1:1, 4:1:1, and 10:1:1.

methods in terms of numerical complexity. In 2D, a typical Lloyd iteration (e.g., Fig. 13) with 2K sites takes around 0.1s, while 100K sites require about 4s. Our Lloyd algorithm running on a 3D example (e.g., Fig. 12) with 30K sites takes around 40s per iteration.

6 Limitations

It should be pointed out that our anisotropic meshing comes with a few challenges and limitations in practice, as we discuss next.

Anisotropy vs. domain shape. If the specified convex function f from which our formulation is derived is arbitrary, anisotropy and local geometry may conflict, since there may not be enough room locally near a boundary to fit a prescribed anisotropy, preventing the approach from generating well-adapted meshes even with numerous refinements. Fu et al. [2014] addressed this issue for simplicial anisotropic meshing by simply preventing further refinement if an edge length in the local metric becomes too small. This is a practical fix, but the generated output will no longer conform to the Hessian of the input function. Instead, our use of *density control* allows us to reduce the sizing near boundary (through the Lipschitz sizing field from [Alliez et al. 2005], see Sec. 5) without changing the anisotropy, hence eliminating this issue completely. However, we pay the price of ending up with more cells so that anisotropy of cells and boundary geometry do not conflict.

General tensor-adapted meshing. In computational physics, simulation of phenomena that are anisotropic due to the physical or mechanical properties of the domain calls for the generation of a mesh adapted to a given symmetric and positive definite (SPD) tensor field σ . In general, this anisotropic tensor is not conformally equivalent to the Hessian of a function (thus, not of the form $\rho \text{Hess}[f]$), so one cannot directly use the machinery described above for this task. In order to construct an algorithm that can generate convex cells adapted to arbitrary SPD tensors by leveraging the properties of our anisotropic technique, we can adapt to our dual setting an approach based on local convex functions [Loseille and Alauzet 2009; Fu et al. 2014] as follows. A local convex function f_i is now associated to each site \mathbf{x}_i , and each local cell is made to best fit this function by minimizing the modified energy, with T_i^i being the hyperplane tangent to f_i at \mathbf{x}_i :

$$\mathcal{E}_{\sigma\text{-OVT}} = \sum_{i=1}^N \int_{V_i} \rho(\mathbf{x}) \left(f_i(\mathbf{x}) - T_i^i(\mathbf{x}) \right) d\mathbf{x}$$

Compared to its primal counterpart, we need to deal with maintaining a cell complex while the per-cell energy is being minimized. A simple approach consists in working with a primal mesh based on the sites \mathbf{x}_i , for which each n -simplex corresponds to a dual vertex of the cell complex. Connectivity changes are thus done on the primal mesh, and dual vertices are consequently updated based on the optimal positions deduced from the local functions. However, searching for the optimal connectivity now needs to be achieved through local flips [Loseille and Alauzet 2009; Fu et al. 2014].

This fact brings a series of theoretical and practical issues. First, like in the primal case, there is no guarantee that such a procedure can reach a tiling of the domain close to the optimal cell complex: local flips can limit the exploration of connectivity. But more importantly, the constant testing of potential flips to achieve energy decrease is significantly slowing down the meshing process. Finding a principled and robust approach allowing arbitrary anisotropic metrics thus remains an open problem.

7 Conclusion and Perspectives

In this paper, we introduced the notion of anisotropic cell complexes as an approximation theoretical dual to Optimal Delaunay Triangulations. We showed how our formulation is a simple and powerful extension to Centroidal Voronoi Tessellations, where now *shifted* and *weighted* sites are used to derive first-type Bregman cell complexes with controllable Hessian-based anisotropy and density. We contrasted our new, simple construction of polyhedral meshes in 2D and 3D with recent anisotropic generalization of Voronoi diagrams by pointing out that the resulting complex has straight-edge convex cells, and admits a dual triangulation that is combinatorially-regular. We also established tight approximation error bounds for our optimal cells in terms of the function, dimension of the domain, and number of vertices. We believe that our work opens a number of important avenues to explore.

Versatility of OVT. First, we note that our formulation can also be used as foundations for other optimizations. For instance, one can improve the quality of function approximation by shifting the local tangent hyperplanes to reduce the L^2 (or even the Sobolev) norm through post-treatment as in [App and Reif 2010]. Control over anisotropy and volume of the convex cells are also bound to find many applications given how CVT meshes have proven to be useful in clustering, sampling, and information theory.

Generality of shifted power diagrams. Second, our introduction of “shifted” power diagrams provides a practical backdrop to the theoretical concepts proposed in [Memari et al. 2012]. Since power diagrams have found recent adoption in graphics for diverse topics from meshing [Mullen et al. 2011] and point sampling [de Goes et al. 2012], to masonry [de Goes et al. 2013; Liu et al. 2013] and fluid simulation [de Goes et al. 2015], this extension is bound to lead to other probably unexpected applications.

Towards variational meshing for arbitrary metrics. Finally, an important result of computational geometry is that every simple cell complex in dimension 3 and above is a power diagram [Aurenhammer 1987]. In arbitrary dimension, a primal-dual triangulation $(\mathcal{T}, \mathcal{C})$ (i.e., a primal triangulation \mathcal{T} and a dual cell complex \mathcal{C} with compatible connectivity) exists if and only if \mathcal{T} is combinatorially-regular. This implies that one can *restrict* a search for any notion




Figure 11: Anisotropic Kitten. The kitten model is decomposed into cells, with increasing anisotropy (1:1, 2:1, 4:1, and 8:1).




Figure 12: Anisotropic Bimba. Bimba for $f(\mathbf{x}) = x^4 + y^4 + z^4 + (x-3)^2 + (y-3)^2 + (z-3)^2$, with a neck-centered radial sizing field. Various exterior and cutaway views to show anisotropy.

of optimal cell complex to shifted power diagrams only: in 3D, this is in fact not a restriction but a convenient way to parameterize the space of generic cell complexes, while in 2D it does reduce the space of cell complexes slightly—but does guarantee the existence of an associated primal triangulation. We plan to exploit this observation to construct a global approach to primal and dual anisotropic meshing for arbitrary tensors, which would entirely bypass the typical (but costly) flip-based combinatorial search of optimal connectivity for arbitrary input metrics.

Acknowledgements

The authors wish to thank Pooran Memari for early discussions, and the reviewers for their feedback. We also thank Shang-Hua Teng for interesting follow-up discussions. This work was partially supported through NSF grants CCF-1011944, IIS-0953096, CMMI-1250261 and III-1302285, and through the European Research Council (Starting Grant IRON: Robust Geometry Processing, agreement 257474). 3D meshes were provided courtesy of AIM@SHAPE. MB and MD gratefully acknowledge the Inria




Figure 13: Quality measures during optimization. We track the distribution of three quality criteria during optimization. In descending order: initial, iteration 1, 10, and 50 (resp. approximation error 163.9, 72.2, 50.4, and 46.4). Bottom image shows result after only 20 steps of interleaved refinement and optimization.

International Chair program and all the members of the TITANE team for their support as well.

References

ALLIEZ, P., COHEN-STEINER, D., YVINEC, M., AND DESBRUN, M. 2005. Variational tetrahedral meshing. In *ACM SIGGRAPH*, vol. 24(3), 617–625.

APP, A., AND REIF, U. 2010. Piecewise linear orthogonal approximation. *SIAM J. Num. Anal.* 48, 3, 840–856.

AURENHAMMER, F. 1987. A criterion for the affine equivalence of cell complexes in \mathbb{R}^d and convex polyhedra in \mathbb{R}^{d+1} . *Disc. & Comput. Geom.* 2, 1, 49–64.

BOISSONNAT, J.-D., COHEN-STEINER, D., AND YVINEC, M. 2006. Comparison of algorithms for anisotropic meshing and adaptive refinement. *Tech. Rep. ACS-TR-362603*, INRIA.

BOISSONNAT, J.-D., WORMSER, C., AND YVINEC, M. 2008. Anisotropic diagrams: Labelle-Shewchuk approach revisited. *Theor. Comput. Sci.* 408, 2–3, 163–173.

BOISSONNAT, J.-D., NIELSEN, F., AND NOCK, R. 2010. Bregman Voronoi diagrams. *Disc. & Comput. Geom.* 44, 2, 281–307.

BOISSONNAT, J.-D., SHI, K.-L., TOURNOIS, J., AND YVINEC, M. 2014. Anisotropic Delaunay meshes of surfaces. *ACM Trans. Graph.*

BOSSEN, FRANK, J., AND HECKBERT, P. S. 1996. A pliant method for anisotropic mesh generation. In *Int. Meshing Roundtable*, 63–76.

CGAL. 2016. *CGAL 4.8 User and Reference Manual*. CGAL Editorial Board, <http://www.cgal.org>.

CHEN, L., AND HOLST, M. J. 2011. Efficient mesh optimization schemes based on Optimal Delaunay Triangulations. *Comput. Methods Appl. Mech. Engrg.* 200, 967–984.

CHEN, L., AND XU, J. 2004. Optimal Delaunay triangulations. *J. of Comp. Mathematics* 22, 2, 299–308.

CHEN, L., SUN, P., AND XU, J. 2007. Optimal anisotropic simplicial meshes for minimizing interpolation errors in L^p -norm. *Math. of Comput.* 76, 257, 179–204.

CHEN, Z., WANG, W., LVY, B., LIU, L., AND SUN, F. 2014. Revisiting Optimal Delaunay Triangulation for 3D graded mesh generation. *SIAM J. Sci. Comput.* 36, 3, 930–954.

CHEN, L. 2004. Mesh smoothing schemes based on Optimal Delaunay Triangulations. In *Int. Meshing Roundtable*, 109–120.

CHENG, S.-W., DEY, T. K., RAMOS, E. A., AND WENGER, R. 2006. Anisotropic surface meshing. In *Symp. Disc. Alg.*, 202–211.

DASSI, F., SI, H., PEROTTO, S., AND STRECKENBACH, T. 2015. Anisotropic finite element mesh adaptation via higher dimensional embedding. *Procedia Engineering* 124, 265–277.

D’AZEVEDO, E. F., AND SIMPSON, R. B. 1989. On optimal interpolation triangle incidences. *SIAM J. Sci. Stat. Comput.* 10, 6, 1063–1075.

DE GOES, F., BREEDEN, K., OSTROMOUKHOV, V., AND DESBRUN, M. 2012. Blue noise through optimal transport. *ACM Trans. Graph.* 31, 6, Art. 171.

DE GOES, F., ALLIEZ, P., OWHADI, H., AND DESBRUN, M. 2013. On the equilibrium of simplicial masonry structures. *ACM Trans. Graph.* 32, 4, Art. 93.

DE GOES, F., WALLEZ, C., HUANG, J., PAVLOV, D., AND DESBRUN, M. 2015. Power particles: An incompressible fluid solver based on power diagrams. *ACM Trans. Graph.* 34, 4, Art. 50.

DU, Q., AND EMELIANENKO, M. 2006. Acceleration schemes for computing centroidal Voronoi tessellations. *Numer. Linear Algebr.* 13, 2–3, 173–192.

- DU, Q., AND WANG, D. 2005. Anisotropic centroidal Voronoi tessellations and their applications. *SIAM J. Sci. Comp.* 26, 3, 737–761.
- DU, Q., FABER, V., AND GUNZBURGER, M. 1999. Centroidal Voronoi tessellations. *SIAM Review* 41, 4, 637–676.
- FREY, P., AND ALAUZET, F. 2004. Anisotropic metrics for mesh adaptation. In *Comput. Fluid Solid Mech.*, K. Bathe, Ed., 24–28.
- FU, X.-M., LIU, Y., SNYDER, J., AND GUO, B. 2014. Anisotropic simplicial meshing using local convex functions. *ACM Trans. Graph.* 33, 6, Art. 182.
- GAO, Z., YU, Z., AND HOLST, M. 2012. Quality tetrahedral mesh smoothing via boundary-optimized Delaunay triangulation. *Comput. Aided Geom. Design* 29, 9, 707–721.
- GEORGE, P., FREY, P., AND ALAUZET, F. 2002. Automatic generation of 3D adapted meshes. *World Congress on Comput. Mech.*
- KUZMIN, D. 2010. *A Guide to Numerical Methods for Transport Equations*. University of Erlangen-Nuremberg.
- LABELLE, F., AND SHEWCHUK, J. R. 2003. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *Symp. Comp. Geom.*, 191–200.
- LEE, C. 1991. Regular triangulations of convex polytopes. *Applied Geom. & Disc. Math.* 4, 443–456.
- LÉVY, B., AND BONNEEL, N. 2013. Variational anisotropic surface meshing with Voronoi parallel linear enumeration. In *Int. Meshing Roundtable*. 349–366.
- LÉVY, B., AND LIU, Y. 2010. L_p Centroidal Voronoi Tessellation and its applications. *ACM Trans. Graph.* 29, 4, Art. 119.
- LIU, Y., WANG, W., LÉVY, B., SUN, F., YAN, D.-M., LU, L., AND YANG, C. 2009. On Centroidal Voronoi Tessellation – energy smoothness and fast computation. *ACM Trans. Graph.* 28, 4, Art. 101.
- LIU, Y., PAN, H., SNYDER, J., WANG, W., AND GUO, B. 2013. Computing self-supporting surfaces by regular triangulation. *ACM Trans. Graph.* 32, 4, Art. 92.
- LOSEILLE, A., AND ALAUZET, F. 2009. Optimal 3D highly anisotropic mesh adaptation based on the continuous mesh framework. In *Int. Meshing Roundtable*. 575–594.
- LOSEILLE, A., AND ALAUZET, F. 2011. Continuous mesh framework, Part I: Well-posed continuous interpolation error. *SIAM J. Num. Anal.* 49, 1, 38–60.
- MARSDEN, J. E., AND TROMBA, A. 2003. *Vector Calculus*, 5th ed. W. H. Freeman.
- MEMARI, P., MULLEN, P., AND DESBRUN, M. 2012. Parametrization of generalized primal-dual triangulations. In *Int. Meshing Roundtable*. 237–253.
- MIREBEAU, J.-M., AND COHEN, A. 2011. Greedy bisection generates optimally adapted triangulations. *Math. Comp.* 81, 811–837.
- MULLEN, P., MEMARI, P., DE GOES, F., AND DESBRUN, M. 2011. HOT: Hodge-optimized triangulations. *ACM Trans. Graph.* 30, 4, Art. 103.
- NADLER, E. 1986. Piecewise linear best L_2 approximation on triangulations. In *Approx. Theory V*, C. K. Chui, L. L. Schumaker, and J. D. Ward, Eds. Academic Press, 499–502.
- NIELSEN, F., AND NOCK, R. 2011. Skew Jensen-Bregman Voronoi diagrams. In *Transactions on Computational Science XIV, Special Issue on Voronoi Diagrams and Delaunay Triangulation*, M. L. Gavrilova, C. J. K. Tan, and M. A. Mostafavi, Eds. Springer, 102–128.
- RICHTER, R., AND ALEXA, M. 2015. Mahalanobis centroidal Voronoi tessellations. *Comp. & Graph.* 46, 48–54.
- RYCROFT, C. H. 2009. VORO++: A three-dimensional Voronoi cell library in C++. *Chaos: An Interdisciplinary Journal of Non-linear Science* 19, 4.
- SHEWCHUK, J. 1998. Tetrahedral mesh generation by Delaunay refinement. In *Symp. on Comp. Geometry*, 86–95.
- SUN, F., CHOI, Y.-K., WANG, W., YAN, D.-M., LIU, Y., AND LÉVY, B. 2011. Obtuse triangle suppression in anisotropic meshes. *Comput. Aided Geom. Design* 28, 9, 537–548.
- TOURNOIS, J., WORMSER, C., ALLIEZ, P., AND DESBRUN, M. 2009. Interleaving Delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. *ACM Trans. Graph.* 28, 3, Art. 75.
- VALETTE, S., CHASSERY, J.-M., AND PROST, R. 2008. Generic remeshing of 3D triangular meshes with metric-dependent discrete Voronoi diagrams. *IEEE Trans. Vis. Comp. Graph.* 14, 2, 369–381.
- YAN, D.-M., WANG, W., LÉVY, B., AND LIU, Y. 2013. Efficient computation of clipped Voronoi diagram. *Computer-Aided Design* 45, 4, 843 – 852.
- ZHONG, Z., GUO, X., WANG, W., LÉVY, B., SUN, F., LIU, Y., AND MAO, W. 2013. Particle-based anisotropic surface meshing. *ACM Trans. Graph.* 32, 4, Art. 99.