

# Optimized Dual-Volumes for Tetrahedral Meshes

Alec Jacobson


University of Toronto and Adobe Research Toronto

## Abstract

Constructing well-behaved Laplacian and mass matrices is essential for tetrahedral mesh processing. Unfortunately, the de facto standard linear finite elements exhibit bias on tetrahedralized regular grids, motivating the development of finite-volume methods. In this paper, we place existing methods into a common construction, showing how their differences amount to the choice of simplex centers. These choices lead to satisfaction or breakdown of important properties: continuity with respect to vertex positions, positive semi-definiteness of the implied Dirichlet energy, positivity of the mass matrix, and unbiased-ness on regular grids. Based on this analysis, we propose a new method for constructing dual-volumes which explicitly satisfy all of these properties via convex optimization.

## 1. Introduction

While the finite-element method with piecewise-linear shape functions may be the *de facto* standard for tetrahedral mesh processing of solids, finite-volume methods are useful for a variety of reasons. In particular, unlike linear finite elements they can result in unbiased operators for tetrahedralized regular grids (see Figure 1). At the heart of the finite-volume method is the definition of local dual-volumes around each vertex of a given tetrahedral mesh. Existing methods attempt to generalize orthogonal Voronoi dual-volumes, which behave well for Delaunay meshes. Unfortunately, when considering general (e.g., non-Delaunay) meshes, each approach gives up one or more desirable properties of its resulting Laplacian and mass matrices: continuity with respect to vertex positions, positive



**Figure 1:** Linear FEM corresponds to dual-volumes formed by connecting simplex barycenters (top, volumes scaled by 50%), but exhibit bias on tetrahedralized grids, e.g., a generalized eigenvector  $\mathbf{L}\mathbf{v}_5 = \lambda_5\mathbf{M}\mathbf{v}_5$  with a pinned base (bottom). Our optimized dual-volumes coincide with circumcentric dual-volumes on a tetrahedralized regular grid (right), avoiding bias.


	$C^0$	$-(\mathbf{L} + \mathbf{L}^T) \succeq 0$	$\mathbf{M} > 0$	GRID
Barycentric	✓	✓	✓	✗
Circumcentric	✓	✗	✗	✓
[AHKS20]	✗	✗	✓	✗
[MMdGD11]	✗	✗	✗	✓
Ours, snapping	✓	✗	✓	✓
Ours, optimized	✓	✓	✓	✓

**Table 1:** Given a mesh with prescribed positions and connectivity we consider existing finite-volume methods. Are their resulting dual volumes continuous functions of vertex positions ( $C^0$ )? Do they yield a positive semi-definite Dirichlet energy ( $-(\mathbf{L} + \mathbf{L}^T) \succeq 0$ ) and generate a positive mass matrix ( $\mathbf{M} > 0$ )? Do they maintain unbiased performance on a tetrahedralized regular grid (GRID)?

semi-definiteness of its implied Dirichlet energy, positivity of the mass matrix, unbiasedness on regular grids (see Table 1).

In this paper, we categorize existing methods according to these criteria. Placing methods into a common construction, their differences amount to the choice of simplex centers, most importantly triangle centers. Each choice drastically affects the discretization of the divergence operator in a finite-volume construction. From their constructions we show how properties of each method arise or break down algebraically.

Based on this analysis, we propose a new method for constructing dual-volumes which explicitly satisfies all of these properties via convex optimization. In contrast to previous methods, our dual-volumes are continuous functions of vertex positions (see Figure 2). As a warm-up, we consider modifying the positive-volume filtering procedure of Alexa et al. [AHKS20] to ensure continuity as tetrahedra morph from acute to non-acute shapes (“Ours snapping” in



**Figure 2:** Previous methods such as [AHKS20] reproduce circum-centric duals for perfect grids, where tetrahedra are on the cusp of becoming non-acute. A gentle nudge and their dual volumes jump to totally different shapes and sizes. Our optimized dual volumes are continuous functions of mesh vertex positions.

Table 1). Unfortunately, this — like [AHKS20] — does not ensure that the Laplacian matrix defines a positive semi-definite Dirichlet energy. Instead, we propose directly optimizing the centers associated with each simplex to explicitly ensure symmetry and negative semi-definiteness of the Laplacian. All properties are satisfied by construction, including continuity as the optimization problem is strictly convex.

Let’s plunge into our problem’s details. We will consider prevalent prior methods after introducing the details of our desired properties and postpone a broader discussion of related work until our final discussion in Section 6.

## 2. Finite Volume Operators

Given a tetrahedral mesh of a solid region  $\Omega \subset \mathbb{R}^3$  with  $n$  vertices, we will follow a general finite-volume construction to build matrices  $\mathbf{M} \in \mathbb{R}^{n \times n}$  and  $\mathbf{L} \in \mathbb{R}^{n \times n}$ .

### 2.1. Mass-matrix

Our goal is to compute a diagonal or *lumped* mass matrix  $\mathbf{M} \in \mathbb{R}^{n \times n}$ , which as a quadratic form computes the inner product of functions defined by per-vertex scalar values  $\mathbf{f} \in \mathbb{R}^n$ .

$$\mathbf{f}^T \mathbf{M} \mathbf{f} \approx \int_{\Omega} f^2 dV \quad (1)$$

We can also consider  $\mathbf{M}$  as a linear operator acting on a *yet to be determined* local dual-volume  $\star_i \subseteq \Omega$  around each vertex  $i$ :

$$(\mathbf{M}\mathbf{f})_i \approx \int_{\star_i} f dV \quad (2)$$

To ensure the diagonality of  $\mathbf{M}$  we now assume that  $f$  is constant over each local volume  $\star_i$ . This local dual volume  $\star_i$  may further be broken into contributions within the primal volume  $\bullet_t$  of each tetrahedron  $t$  (usually only those incident on  $i$ ):

$$(\mathbf{M}\mathbf{f})_i = \left( \sum_t \int_{\bullet_t \cap \star_i} dV \right) \mathbf{f}_i, \quad \mathbf{M}_{ii} := \sum_t \int_{\bullet_t \cap \star_i} dV, \quad (3)$$

with  $\mathbf{M}_{ij} = 0$  for  $i \neq j$ .

### 2.2. Dirichlet Energy / Laplacian Matrix

Our goal is also to compute a sparse Laplacian matrix  $\mathbf{L} \in \mathbb{R}^{n \times n}$  which approximates the continuous Laplacian operator acting on scalar functions or equivalently acting as a quadratic form computing the integral of squared gradients:

$$-\mathbf{f}^T \mathbf{L} \mathbf{f} \approx \int_{\Omega} \nabla f \cdot \nabla f dV = - \int_{\Omega} f \Delta f dV + \int_{\partial\Omega} f \nabla f \cdot \hat{\mathbf{n}} dS \quad (4)$$

where the boundary term vanishes because we assume zero Neumann boundary conditions ( $\nabla f \cdot \hat{\mathbf{n}} = 0$ ), which are this energy’s natural boundary conditions.

Again, we consider the matrix as computing the local integral at each vertex  $i$ :

$$(\mathbf{L}\mathbf{f})_i \approx \int_{\star_i} \Delta f dV \quad (5)$$

This time, we assume that gradient  $\nabla f$  is constant over each local dual volume  $\star_i$ , which will result in a sparse (but not diagonal) matrix  $\mathbf{L}$ . Again the local dual volume  $\star_i$  may be broken into contributions within each primal tetrahedron volume  $\bullet_t$ , upon which divergence theorem may be applied:

$$(\mathbf{L}\mathbf{f})_i = \sum_t \int_{\bullet_t \cap \star_i} \Delta f dV = \sum_t \left( \int_{\partial(\bullet_t \cap \star_i) \cap \Omega} \hat{\mathbf{n}} dS \right) \cdot (\nabla f)_t, \quad (6)$$

where the “ $\cap \Omega$ ” enforces the zero Neumann boundary conditions by *not* including any part of the boundary of the overall solid volume  $\Omega$ . By now  $\mathbf{L}_{ij}$  are fully determined by the choice of local dual volume  $\star_i$ , but we momentarily postpone its explicit formula until after defining our general dual volume construction.


## 3. Dual Volumes

In this paper, we provide a general construction of the local volume contribution  $\bullet_t \cap \star_i$  from each tetrahedron  $t$  to the dual volume of each mesh vertex  $i$  as a hexahedron with vertices at “centers” of all shared simplices (see Figure 3):

- 1 at the vertex  $i$ ,
- 3 at edges of  $t$  incident on  $i$ ,
- 3 at triangle facets of  $t$  incident on  $i$ , and
- 1 at the tetrahedron  $t$ .

For each tetrahedron  $t$  incident on a vertex  $i$ , these centers define the eight corners of a hexahedron  $\bullet_t \cap \star_i$ . To compute the entries of  $\mathbf{M}$  and  $\mathbf{L}$ , we may explicitly create a mesh of this local contribution and compute its volume (for  $\mathbf{M}$ ) or integral of surface normals (for  $\mathbf{L}$ ). For simple choices of centers, we can cook up closed-form expressions for the entries of  $\mathbf{M}$  and  $\mathbf{L}$  without explicitly realizing a mesh of  $\bullet_t \cap \star_i$ .

Nevertheless, the choice of how each center is defined will affect the following criteria we may care about in the mass and Laplacian matrix definitions above:



**Figure 3:** In our general construction, each tetrahedron  $t$  contributes a portion of its volume  $\bullet_t \cap \star_i$  to the local finite dual-volume of each incident vertex  $i$ . This subvolume is determined by centers defined for each simplex. Different methods choose different centers leading to different properties of constructed operators.

### $C^0$ — continuous function of vertex positions

The local volume  $\star_i$  should be a continuous function of tetrahedral mesh vertex positions. This ensures that integrated quantities stored on mesh vertices don't jump in value as the underlying mesh changes continuously, for example, in Lagrangian simulation or mesh deformation methods (see Figure 2).

### $-\mathbf{L} + \mathbf{L}^T \succeq 0$ — positive semi-definite Dirichlet energy

Treating the Laplacian matrix  $\mathbf{L}$  as a quadratic form should result in a semi-definite<sup>†</sup> Dirichlet energy. That is, for any input scalar function  $\mathbf{f} \in \mathbb{R}^n$  we expect a non-negative Dirichlet energy:

$$-\mathbf{f}^T \mathbf{L} \mathbf{f} \geq 0 \quad \forall \mathbf{f} \in \mathbb{R}^n, \quad (7)$$

with equality only for constant functions (see Figure 4). Various methods (e.g., [VL08; Rus07; SOG09; OBS\*12; SACO22; BZC\*23]) rely on explicit spectral decomposition and Dirichlet energy indefiniteness presents a major problem.

Dirichlet energy positive semi-definiteness requires that the sum of  $\mathbf{L}$  and its transpose is negative semi-definite or equivalently that all eigenvalues of the Laplacian plus its transpose are real and non-positive:

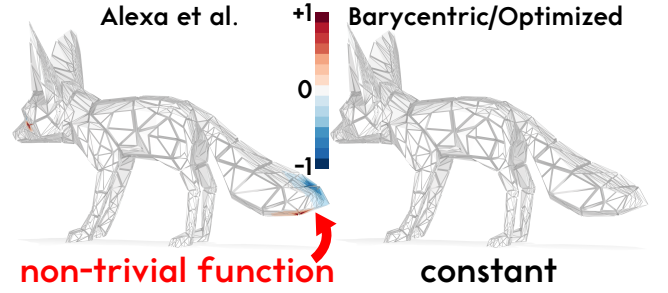
$$\lambda_{\mathbf{L} + \mathbf{L}^T} \leq 0, \quad (8)$$

with a single zero eigenvalue corresponding to the constant function. While symmetry and negative semi-definiteness together are sufficient to ensure this property, they are not necessary.

$$\mathbf{L} = \mathbf{L}^T \quad \text{and} \quad \lambda_{\mathbf{L}} \leq 0 \quad \not\Rightarrow \quad -(\mathbf{L} + \mathbf{L}^T) \succeq 0. \quad (9)$$

Indeed, symmetry alone or non-positive eigenvalues alone does not appear to ensure any practical value of  $\mathbf{L}$ . Meanwhile, an indefinite  $\mathbf{L}$  which does define a positive semi-definite Dirichlet energy is effective for many or most applications.

<sup>†</sup> We use  $\mathbf{X} \succeq 0$  to write that the matrix  $\mathbf{X}$  is positive semi-definite and  $\mathbf{z} \geq 0$  to write that every element of  $\mathbf{z}$  is non-negative.



**Figure 4:** In this numerical example, there exists a non-trivial unit-norm function  $\hat{\mathbf{f}}$  (left) such that the implied Dirichlet energy of [AHKS20] is negative (and the energy is nonconvex):  $\hat{\mathbf{f}}^T (-\mathbf{L}) \hat{\mathbf{f}} \approx -2.96$ . For the barycentric and our optimized centers,  $\mathbf{L}$  is symmetric negative semi-definite and that the smallest energy value is zero from a constant function (right).

The risk of losing Dirichlet energy positive semi-definiteness is an immediate consequence of the choice to discretize at the operator level rather than the quadratic form level. Alternatively, linear FEM or discrete calculus of variations would alternatively discretize the function space and operators present in the energy functional itself, where all terms are squared and thus non-negative.

Positive semi-definite or not, we can consider the Dirichlet “energy” implied by  $\mathbf{L}$  as a summation of contributions per-tetrahedron:

$$\mathbf{f}^T \mathbf{L} \mathbf{f} = \sum_t \mathbf{f}^T \mathbf{S}_t^T \mathbf{L}_t \mathbf{S}_t \mathbf{f}, \quad (10)$$

where  $\mathbf{S}_t \in \mathbb{R}^{4 \times n}$  selects the entries of  $\mathbf{f}$  corresponding to the vertices of tetrahedron  $t$ , and  $\mathbf{L}_t \in \mathbb{R}^{4 \times 4}$  is the contribution of tetrahedron  $t$  to the Laplacian matrix. By assuming that  $\nabla f$  is constant in each tetrahedron  $t$ , we have reduced the discretization of the integrated Laplacian to that of the per-vertex divergence operator, which in turn can be decomposed into a sum of contributions per-tetrahedron, namely:

$$\mathbf{L}_t = \mathbf{D}_t \underbrace{\mathbf{G}_t}_{\text{fixed by assumption}} \quad (11)$$

where  $\mathbf{D}_t \in \mathbb{R}^{4 \times 3}$  and  $\mathbf{G}_t \in \mathbb{R}^{3 \times 4}$  and the entries are determined according to Eq. (6):

$$\left( \mathbf{S}_t^T \mathbf{D}_t \right)_i = \int_{\partial(t \cap \star_i) \cap \Omega} \hat{n} dS \quad (12)$$

If we require symmetry  $\mathbf{L} = \mathbf{L}^T$  for all input meshes, we are — by reduction — requiring that:

$$\mathbf{D}_t \mathbf{G}_t = \mathbf{G}_t^T \mathbf{D}_t^T = -\mathbf{G}_t \mathbf{A}_t \mathbf{G}_t^T \quad (13)$$

for all tetrahedra  $t$ , where  $\mathbf{A}_t \in \mathbb{R}^{3 \times 3}$  is some symmetric matrix and minus sign appears to foreshadow the interpretation of  $\mathbf{A}_t$  as a piecewise-constant metric or diffusion tensor. Letting  $\mathbf{1}$  be the vector of ones, then because  $\mathbf{G}_t \mathbf{1} = 0$ , we immediately have  $\mathbf{L}_t \mathbf{1} = 0$ . If we further require that  $\mathbf{L}_t \preceq 0$ , then we are requiring that  $\mathbf{A}_t$  is positive definite:

$$\mathbf{A}_t \succ 0. \quad (14)$$

**M > 0 — positive mass matrix**

The mass matrix  $\mathbf{M}$  should be positive for a non-degenerate input mesh. As  $\mathbf{M}_{ii}/\sum_j \mathbf{M}_{jj}$  represents the percentage of the total volume of  $\Omega$ , we expect each local volume  $\star_i$  to be positive. Since  $\mathbf{M}$  is diagonal by construction, this property is true if and only if the mass-matrix  $\mathbf{M}$  is strictly positive definite. Indefinite or worse singular  $\mathbf{M}$  causes all sorts of problems in numerical methods. A sufficient condition for this property is that the local volume contribution in each tetrahedron is positive:  $\int_{\bullet_i \cap \star_i} dV > 0$  for  $t \ni i$ , or more strictly that all centers lie within their respective simplices.

**GRID — unbiased volumes on regular grids**

As a unit-test for “connectivity bias,” the dual volume definition should result in unbiased local volumes on a regular grid where each cube cell is subdivided into six tetrahedra around an arbitrary diagonal of symmetry (see Figure 1).

With conditions in hand, let’s immediately consider some obvious center choices and then those of prevalent prior works (summarized in Table 1).


**3.1. Barycentric dual-volumes**

The barycentric dual-volumes are defined by using the barycenter (a.k.a., centroid or center of mass) of each simplex. For the mass-matrix, this results in an equal division of each tetrahedron volume to its incident vertices:

$$\mathbf{M}_{ii}^{\text{bary}} = \frac{1}{4} \sum_{t \ni i} \int_{\bullet_i} dV \tag{15}$$

Using barycenters results in the same Laplacian matrix as direct application of the linear finite element method: the ubiquitous cotangent Laplacian.

Barycenters are clearly continuous with respect to vertex positions and the cotangent Laplacian is negative semi-definite. In particular,  $\mathbf{D}_t^{\text{bary}} = \mathbf{G}_t^T \mathbf{A}_t^{\text{bary}}$  where  $\mathbf{A}_t^{\text{bary}} = (\int_{\bullet_i} dV) \mathbf{I}$  (see Figure 5). Barycenters always lie inside each respective simplex. This immediately implies that the mass matrix is positive. However, because



**Figure 5:** In our construction, any choice of dual-volume implies a  $3 \times 3$  matrix  $\mathbf{A}_t$  which can be interpreted as a piecewise-constant diffusion tensor. Barycentric dual-volumes, corresponding to linear FEM, imply tensors proportional to the identity: in a way, trusting the connectivity too much. Circumcentric dual-volumes imply tensors which counteract connectivity bias on tetrahedralized grids.

tetrahedra volumes are distributed to each incident vertex based on connectivity not geometry, the barycentric dual-volumes suffer strong bias for a tetrahedralized regular grid (see Figure 1).

**3.2. Circumcentric dual-volumes**


Circumcentric dual-volumes are defined by using the circumcenter of each simplex. For interior vertices of a Delaunay mesh, this choice results in local dual volumes matching the Voronoi diagram of vertices. Circumcenters are a continuous function of vertex positions for non-degenerate meshes.

While not obvious *a priori*, circumcentric dual-volumes result in a symmetric Laplacian matrix. Unlike their 2D analogs for triangle meshes, circumcentric dual-volumes do not result, in general, in the same Laplacian matrix as the linear-FEM or barycentric dual-volume Laplacian:

$$\mathbf{L}^{\text{circum}} \neq \mathbf{L}^{\text{bary}}. \tag{16}$$


That is, in general,  $\mathbf{A}_t^{\text{circum}} \neq \mathbf{A}_t^{\text{bary}}$  for a given tetrahedron  $t$  as  $\mathbf{A}_t^{\text{circum}}$  is, in general, not diagonal. We could alternatively interpret a Laplacian built with circumcentric dual-volumes as a piecewise-linear finite-element method applied to a mesh with non-trivial piecewise-constant metric diffusion tensors  $\mathbf{A}_t^{\text{circum}} \in \mathbb{R}^{3 \times 3}$ ; in other words,  $\mathbf{A}_t^{\text{circum}}$  determines a per-tetrahedron transformation  $\mathbf{J}_t$  applied to each tetrahedron independently before applying linear-FEM construction (see Figure 5). There is no guarantee that the per-tetrahedron transformations could be realized by a vertex deformation of the whole mesh in  $\mathbb{R}^3$ .

For a regular grid, circumcentric dual-volumes are unbiased, coinciding with a finite-differences. Unfortunately, if the mesh is non-Delaunay the mass-matrix is not guaranteed to be positive. Circumcenters may slip outside simplices, causing the boundary of the dual volume to self-intersect. This may eventually cause the total signed volume of  $\star_i \cap t$  and even the total signed volume of  $\star_i$  to be negative (see Figure 6).



**Figure 6:** For non-Delaunay tetrahedral meshes such as this heart, circumcentric dual-volumes can result in negative entries in the mass matrix, due to self-intersecting, flipped-inside-out contributions from each element (left). Barycentric (center) and our optimized dual-volumes (right) are always positive.





**Figure 7:** As a tetrahedron continuously deforms from equilateral to non-acute, the circumcenters of the tetrahedron and its triangles pass through the triangle and edge circumcenters, respectively. This inspires our circumcentric snapping generalization of [MDSB02], which ensures continuity of dual-volumes but not symmetry of the Laplacian  $\mathbf{L}$ .

### 3.3. [AHKS20] dual-volumes: barycentric snapping

To combat negative volumes in 2D triangle meshes, Meyer et al. [MDSB02] proposed using circumcentric triangle centers *except* when the triangle is obtuse. In this case, Meyer et al. would snap the center to the midpoint of the edge opposite the obtuse angle. Citing inspiration from the idea of center-filtering at a high-level, Alexa et al. [AHKS20] snap any centers lying outside its simplex to that simplex's *barycenter*. By construction the mass-matrix is positive, but dual-volumes on a  $\epsilon$ -perturbed regular grid are biased as every entry suddenly jumps in shape and size to the biased barycentric dual-volume mass matrix (see Figure 2). Consider a simplex as it morphs from non-obtuse to obtuse: its circumcenter and barycenter will not coincide during the transition. Therefore, the dual volumes of [AHKS20] are not a continuous function of vertex positions.

For *triangle* meshes, the choice of triangle center does not affect the definition of the Laplacian matrix (by Stokes Theorem and that the gradient is constant in each triangle. See, e.g., Sec. 3.3.4 in [BKP\*10]). As noted, edge midpoints are equivalent to barycenters, so the resulting Laplacian matrix agrees with the *symmetric negative semi-definite* linear-FEM / barycentric dual-volume Laplacian.

For tetrahedral meshes, the choice of tetrahedron center also does not affect the Laplacian (again, by applying Stokes Theorem against a constant vector field in each tetrahedron) and of course edge circumcenters are also barycenters, but obtuse triangle circumcenters snapped according to [AHKS20] will affect the Laplacian matrix. As a result, the  $\mathbf{L}^A$  will be non-symmetric in the presence of obtuse triangles. Indeed, for a single tetrahedron with an obtuse triangle the resulting  $\mathbf{D}_t^A$  does not satisfy  $\mathbf{D}_t^A \mathbf{G}_t = \mathbf{G}_t^T (\mathbf{D}_t^A)^T$ . We may solve for its analogous implied “diffusion tensor” *post facto* as

$$\mathbf{A}_t^A = -(\mathbf{G}_t^T)^+ \mathbf{D}_t^A \quad (17)$$

which in this case reveals that  $\mathbf{A}_t^A$  is, in general, not symmetric and  $\mathbf{A}_t^A + (\mathbf{A}_t^A)^T$  is not positive definite. This implies that we do not have, in general, a positive semi-definite Dirichlet energy formed with  $\mathbf{L}^A$  (see Figure 4). Across datasets of larger tetrahedral meshes, we experimentally find that  $\mathbf{L}^A$  is consistently significantly asymmetric. Often,  $\mathbf{L}^A$  will have all non-negative eigenvalues, but  $\mathbf{L}^A + (\mathbf{L}^A)^T$  is indefinite. Meanwhile, other times  $\mathbf{L}^A$  will even have non-trivially imaginary eigenvalues, but  $\mathbf{L}^A + (\mathbf{L}^A)^T$  will be negative semi-definite.

### 3.4. [MMdGD11] optimization over orthogonal dual volumes

Orthogonal duals are a subset of simplex center construction where dual faces are geometrically orthogonal to primal edges. The circumcentric dual-volumes of Section 3.2 are a special case of orthogonal duals, but the space is much larger [Gli05; WMKG07; MMdGD11]. Orthogonal dual constructions for triangle meshes (and less often tetrahedral meshes) have been considered in the past, primarily discussed through the lens of discrete exterior calculus [DhLM05].

Mullen et al. [MMdGD11] propose to optimize over the space of metrics in each local neighborhood so that the dual is orthogonal by construction and other properties such as mass positivity is encouraged by minimization. Mullen et al. may optimize both the metrics and the primal mesh simultaneously. For our consideration, we will consider just the metric optimization. Unfortunately, the optimization problem is non-linear and non-convex, requiring a variant of gradient descent to find a local minimum. This forgoes the guarantee of continuity with respect to vertex positions. While they report *reducing* the amount of negative volumes compared to circumcentric duals, in their examples they do not eliminate negative volume entirely. To our knowledge, there is no available open source implementation of this method to provide explicit counter-examples here.

### 3.5. Continuous dual-volumes via circumcentric snapping

As a warm-up to our main result, we consider modifying the positive-volume filtering of Alexa et al. [AHKS20] to ensure continuity as tetrahedra morph from acute to non-acute shapes (we say a tetrahedron is “acute” if all dihedral angles are less than  $\pi/2$  [ESÜ04]).

Unlike Alexa et al., the original center-filtering of Meyer et al. [MDSB02] ensures continuity with respect to vertex positions because the triangle circumcenter is snapped to the *circumcenter* of its subsimplex (the midpoint of the edge). We can generalize the choice of Meyer et al. [MDSB02] to 3D tetrahedral meshes by snapping circumcenters of simplices recursively to subsimplices. This choice ensures that dual-volumes are continuous with respect to vertex positions. To see this, consider a non-obtuse simplex morphing into an obtuse one. At the moment that its circumcenter exits the simplex it must coincide with a circumcenter of the boundary subsimplex upon which it lies (because it is also equidistant to that subsimplex's corners, see Figure 7).

By design, snapping ensures that the mass matrix is positive, and snapping to circumcenters ensures that the mass matrix of a regular grid remains unbiased like circumcentric dual-volumes. Unfortunately, circumcentric snapping — like [AHKS20] — does not guarantee symmetry or a positive semi-definite Dirichlet energy.

#### 4. Optimized dual-volumes

Drawing inspiration from Meyer et al. [MDSB02] and [AHKS20], we propose to utilize circumcentric dual-volumes when possible. Unlike those methods, we will explicitly ensure a positive semi-definite Dirichlet energy *by constraint*.

Vertex circumcenters are trivially well positioned, and edge circumcenters are non-troublesome midpoints. Meanwhile, the tetrahedra centers do not affect the Laplacian matrix, so they can be snapped individually *post facto*. What remains are the triangle centers.

We propose to optimize the center location of each triangle to be as close as possible to its circumcenter, but to stay within their respective triangles *and* so that the implied diffusion tensor  $\mathbf{A}_t$  of each tetrahedron is positive definite. To ensure the continuity of dual-volumes, two tetrahedra sharing a triangle must agree on the triangle's center. Thus we consider the following optimization over the centers of a tetrahedral mesh with vertex positions in  $\mathbf{v}_i \in \mathbb{R}^3$  for  $i = 1, \dots, n$  and triangle indices gathered into  $\mathbf{F} \in [1, n]^{k \times 3}$ :

$$\min_{\mathbf{C} \in \mathbb{R}^{k \times 3}} \sum_f \|\mathbf{o}_f - \mathbf{c}_f\|^2 \quad (18)$$

$$\text{s.t. } \mathbf{c}_f \in \bullet_f \quad \forall f \quad (19)$$

$$\text{and } (\mathbf{G}_t^T)^+ \mathbf{D}_t(\mathbf{C}) \preceq 0 \quad \forall t \quad (20)$$

where  $\mathbf{o}_f$  is the (known) circumcenter of triangle  $f$  and  $\mathbf{c}_f$  is the optimized center,  $\mathbf{D}_t(\mathbf{C}) : \mathbb{R}^{k \times 3} \rightarrow \mathbb{R}^{4 \times 3}$  is the deterministic and linear function which selects the centers of the four triangles incident on tetrahedron  $t$  from  $\mathbf{C}$  and builds the integrated divergence operator according to the construction in Eq. (12).

Eqs. 18-20 define a strictly convex semi-definite programming (SDP) problem. The problem always has a solution: triangle barycenters are feasible. Therefore, solutions of this feasible strictly convex optimization are a continuous function of its parameters (mesh vertex positions). The constraint in Eq. (20) ensures that the Laplacian matrix is symmetric negative semi-definite, and thus the Dirichlet energy is positive semi-definite. The constraint in Eq. (19) ensures that the mass matrix is positive. When a circumcentric dual-volume satisfies both these properties, our solution will coincide, thus ensuring unbiased dual-volumes on regular grids. Perturbing the mesh vertices slightly will result in a continuous change, degrading gracefully while maintaining definiteness and positivity (see Figure 2).

Large sparse semi-definite programs are tractable to solve, but rather slow. Fortunately, we observe that simply constraining  $\mathbf{A}_t$  to be *symmetric* while ensuring positive dual volumes (via Eq. (19)) appears to always result in a positive semi-definite  $\mathbf{A}_t$ . While we lack a formal or symbolic proof, we constructed a numerical experiment on a unit tetrahedron to hunt for an *negative* definite  $\mathbf{A}_t$  while keeping centers inside their triangles: finding a result would

amount to a counterexample to our claim. However, the SDPT3 solver [TTT03] declared this counterexample-search problem infeasible: agreeing with our hypothesis that symmetry and positive dual volumes are sufficient conditions for positive definiteness. Out of an abundance of caution, for a given mesh we can always verify whether the resulting symmetric  $\mathbf{A}_t$  matrices are positive semi-definite and, if not, resolve the full SDP using a large-sparse-SDP solver like Mosek [AA00]. We have not encountered the need for this yet.

So, we propose to replace Eq. (20) with the following linear equality constraints enforcing symmetry:

$$\mathbf{D}_t(\mathbf{C}) \mathbf{G}_t = \mathbf{G}_t^T \mathbf{D}_t(\mathbf{C})^T \quad \forall t. \quad (21)$$

Eq. (21) appears to be a rank-6 constraint over the entries of  $\mathbf{D}_t$ . However, after symbolical analysis, we find that this constraint is in fact only rank 3 over the entries of  $\mathbf{C}$ . Thus, each tetrahedron contributes 3 linear equality constraints (Eq. (13)), touching only its four incident triangles' centers.

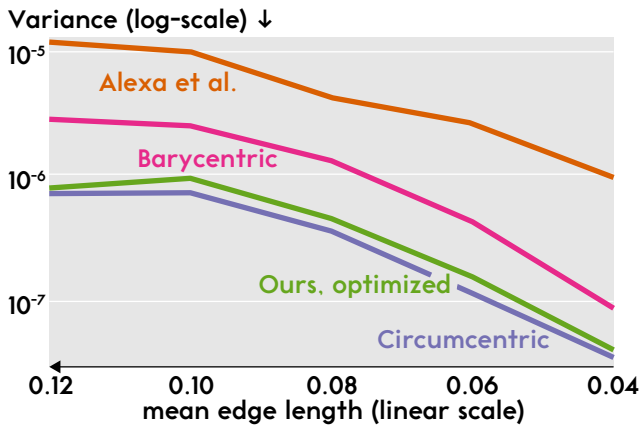
Combined with the quadratic loss in Eq. (18) and the linear inequality constraints in Eq. (19), we have a convex quadratic program (QP). By re-parametrizing  $\mathbf{C}$  as a linear function of barycentric coordinates  $\mathbf{B} \in \mathbb{R}^{k \times 3}$ , we can express the constraints in Eq. (19) as non-negative and partition of unity constraints on  $\mathbf{B}$ . This large, sparse QP can be efficiently solved with a variety of standard solvers [AA00; CKKD20; SBG\*20; Mat24].

#### 4.1. Trivial extension to intrinsic dual-volumes

While the value of intrinsic 2D triangulations has been made exceptionally clear [FSSB07; BS07; SSC19; GSC21; LGC\*23], this value has perhaps not yet extended to 3D tetrahedral meshes. Nevertheless the extension of our general centric construction — and thus also our optimized dual-volumes — to intrinsic dual-volumes is a trivial extension. Our quantities and constraints are defined per-element and our optimization variables can be expressed in terms of (intrinsic) barycentric coordinates. To accommodate an intrinsic tetrahedralization — where mesh edges have lengths but no prescribed vertex positions — one could simply map each tetrahedron to a canonical tetrahedron with matching metric, compute relevant quantities and constraints. Each  $\|\mathbf{o}_f - \mathbf{c}_f\|^2$  in Eq. (18) would split into contributions from each adjacent tetrahedron (expressed in barycentric coordinates).


#### 5. Numerical Experiments

In Figure 8, we replicate the “Dirichlet energy minimization” experiment from [AHKS20], where we generate progressively refined tetrahedral meshes conforming to concentric spheres of radii 1.0, 0.75 and 0.5. We minimize  $-\mathbf{f}^T (\mathbf{L} + \mathbf{L}^T) \mathbf{f}$  subject to boundary conditions of 1.0 and 0.0 on the inner and outer shells respectively, then measure variance of the vertex values on the middle shell. In our experiment, we generate sphere triangle meshes using centroidal Voronoi tessellation, then use Tetgen [S115] to mesh in between them. As shown by Alexa et al., barycentric duals result in higher variance than circumcentric. However, the filtered centers proposed by Alexa et al. result in even higher variance. Our optimized centers result in variance just above circumcentric duals.




**Figure 8:** Compared to circumcentric, barycentric dual-volumes are known to produce more variance when used to solve a radially symmetric Laplace problem. The filtered centers of Alexa et al. [AHKS20] result in even higher variance. Our optimized centers result in variance just above circumcentric duals. In all cases, variance decreases with mesh refinement.

We made almost no effort optimizing the runtime performance of our numerical implementation. We implemented the optimization of face centers in Matlab with lots of `for` loops. Nevertheless, the bottleneck is solving the large-sparse quadratic program (using Mosek [AA00]). The number of variables in the optimization (three times the number of triangles  $3k$ ) closely tracks  $30\times$  the number of mesh vertices in the previous experiment. The inset plot shows that sub-quadratic performance in the number of vertices, far better than worst-case complexity for convex quadratic programs indicating that sparsity is paying off. Timings are recorded on a 2020 MacBook Pro laptop (2.3GHz Quad-core Intel Core i7, 32GB RAM).



Once constructed numerically, our Laplacian  $\mathbf{L}$  and mass matrix  $\mathbf{M}$  have the same sparsity and thus same runtime implications for downstream tasks as others methods (but unlike alternatives with larger stencils, e.g., [BBA21]). In Figure 4, we compare the numerically computed eigenvectors of the Laplacian built from our dual volumes to that of Alexa et al. [AHKS20], whose lack of semi-definiteness manifests as numerical debris in the eigenvectors. Laplacian eigenmodes are a critical downstream task used in a variety of applications areas in geometry processing [VL08; Rus07; SOG09; OBS\*12; SACO22; BZC\*23].

Our optimized dual-volumes are continuous with respect to vertex positions. In Figure 9, we consider a downstream design analysis task where the tetrahedral mesh connectivity of a wrench is transported along with procedural changes to its design *à la* [SXZ\*17]. The squared norm of a sinusoidal function  $\mathbf{f}$  is measured across the changing designs:  $\mathbf{f}^T \mathbf{M} \mathbf{f}$ . The line plots show relative differences of each method's value compared to that of barycentric




**Figure 9:** Consider computing  $\mathbf{f}^T \mathbf{M} \mathbf{f}$  as the vertex positions of a tetrahedral mesh change with its design (bottom) Barycentric, circumcentric, and our optimized dual-volumes are continuous function of mesh vertex positions, and so is this value computed with their respective mass matrices. The dual volumes of Alexa et al. [AHKS20] are discontinuous, producing a noisy value as the design changes. The line plot (top) shows the relative difference of each method's value compared to that of barycentric dual volumes as the design changes.

dual volumes as the design changes. The dual volumes of Alexa et al. [AHKS20] produce a noisy discontinuous function, while ours — like barycentric and circumcentric — are continuous with respect to change in vertex positions.

## 6. Discussion & Conclusion

As casually observed by Alexa et al. [AHKS20], tetrahedral meshes are often not Delaunay in the wild. *Constrained* Delaunay meshes (e.g., [Si15; DPVA23]) do not enforce the local Delaunay criteria at constrained (e.g., boundary) faces. Meanwhile, optimization-based meshes (e.g., [HZG\*18; HSW\*20]) will trade Delaunay-ness for improved aspect ratios or other important mesh quality measures. Indeed, considering the 10,000 outputs of TETWILD run on the Thing10k dataset [ZJ16], we see that most tetrahedra circumcenters lie outside, and roughly one quarter of face circumcenters lie outside (inset). Consequently, circumcentric dual-volumes result in negative mass matrix entries for 74% of the meshes across the dataset. The barycentric snapping method of Alexa et al. [AHKS20] results in asymmetric Laplacian matrix  $\mathbf{L}$  for 100% of the models, and of these Cholesky decomposition of  $\mathbf{L} + \mathbf{L}^T$  fails on 57%.

### % of circumcenters outside triangles and tets



The general construction considered in this paper explicitly constructs the dual volume. Besides making constraints and analysis explicitly clear, ensuring a geometric realization could prove essential for applications which rely on sampling or interpolation within dual-volumes. Elcott et al. [ETK\*07] use primal-dual meshes for fluids and generalized barycentric coordinates for in-

terpolation. Batty et al. [BXH10] extend this to non-conforming Delaunay meshes, where circumcentric dual volumes are positive and convex shaped. Our optimized dual-volumes are positive but not necessarily convex; on the other hand, we do not require a Delaunay mesh, so meshes can be conforming. While convenient for velocity interpolation, convexity is not strictly necessary and generalized barycentric coordinates have since improved for non-convex polyhedra [JMD\*07; HS08; ZDL\*14; CDH23].

In this paper, we considered construction and optimization of dual-volumes assuming a contract with the user that the primal mesh remains fixed. Prior works have also considered (simultaneous) optimization of the primal mesh or its intrinsic connectivity to improve Laplacian and mass-matrix operators for 2D triangulations [dMMD14] and tetrahedral meshes [SMWF22] or both [MMdGD11; Ale19; ABE\*20; Ale20; MKMD23]. Alexa [Ale19] shows how unique the situation is for triangulations in 2D and that many clear concepts or algorithms do not extend to 3D and beyond. While we believe our optimization generalizes in theory to higher dimensional meshes, we leave this exploration to future work.

The space of orthogonal duals is too small to include the barycentric duals (and thus linear FEM, the *de facto* standard). Our general construction with simplex centers and space of optimization explicitly includes these, while others do not [WMKG07; MMdGD11].

When considering the orthogonal dual construction of the Laplacian for 2D triangle meshes, the choice of triangle center at the circumcenter seems to have great importance: Wardetzky et al. [WMKG07] write “the cotan weights ... arise from assigning dual vertices to circumcenters of primal triangles” [WMKG07]). However, in the more general view of constructing dual-volumes by connecting simplex centers, the choice of triangle center is demonstrably irrelevant. For the cotan weights to arise, only the choice of edge centers at midpoints matters (see also [BKP\*10]).

In 3D, the situation is very different. Edge centers remain canonically midpoints, but triangle centers are neither irrelevant nor canonically determined. Fortunately, they span an interesting enough space to find dual-volumes that fulfil a variety of properties. Like the triangle center in 2D, the tetrahedron center in 3D is irrelevant to the Laplacian matrix, but does affect the mass matrix.

The taxonomy of Wardetzky et al. [WMKG07] is not quite expressive enough to capture the locality or lack thereof of our proposed approach. The sparsity pattern of our Laplacian matrix matches the combinatorial mesh, but the entries are effectively determined via a global optimization (similar to [MMdGD11]).

Alexa et al. [AHKS20] muse about linearly blending between barycentric and circumcentric duals. Using a single parameter for the entire mesh would result in a space that could ensure positivity, definiteness and continuity, but not unbiased-ness on (local) grid structures. It is not obvious how to spatially vary this blend so that properties are ensured locally; our optimization works on a large space by considering effectively two parameters per triangle (barycentric coordinates of its center).

Wardetzky et al. [WMKG07] propose a general construction of symmetric 2D triangle mesh Laplacians by spanning a parameterization over the metrics of each edge. This is analogous to our ob-

servation that once  $\nabla f$  is assumed to be piecewise constant in an element, what remains is the choice of divergence operator  $\mathbf{D}$  or diffusion tensor  $\mathbf{A}$ . Perhaps the construction of Wardetzky et al. could be generalized to tetrahedral meshes, but it would seem to correspond to directly parameterizing  $\mathbf{A}$  rather than working with simplex centers. By optimizing centers directly we ensure that dual-volumes are continuous across triangles and embedded. In this paper, we focus on the effect of these properties on the Laplacian and mass matrices. It would be interesting to explore our optimized dual-volumes for other operators or uses (e.g., fluids [BXH10]).

Employing convex-optimization to *build* the Laplacian and mass matrices will often be more expensive than a single subsequent sparse linear system or truncated eigenvalue decomposition. In some applications, ensuring the features in Table 1 may be worth the wait. In any case, our result also serves as an existence proof that we hope fuels future research into more efficient methods. Finally, any invocation of a feature chart like Table 1 should be met with a critical speculation about which other feature columns should be considered in the future. We hope the analysis and novel method presented here will continue to fuel research into dual-volumes. To this end, we have open-sourced implementation as part of the gptoolbox [Jac\*24] library.

## Acknowledgements

Our research is funded in part by NSERC Discovery (RGPIN-2022-04680), the Ontario Early Research Award program, the Canada Research Chairs Program, a Sloan Research Fellowship, the DSI Catalyst Grant program and gifts by Adobe Inc. We are grateful for the feedback of the Banh Mi weekly meeting attendees at University of Toronto, insightful conversations with Oded Stein & Nick Sharp, and early draft comments from Towaki Takikawa, Abhishek Madan, & Eitan Grinspun. The “Fennec Fox” model in Figure 4 is by Physics\_Dude by and used according to the CC BY 4.0 license.

## References

- [AA00] ANDERSEN, ERLING D. and ANDERSEN, KNUD D. “The MOSEK Interior Point Optimizer for Linear Programming: An Implementation of the Homogeneous Algorithm”. *High Performance Optimization* (2000), 197–232 6, 7.
- [ABE\*20] ABDELKADER, AHMED, BAJAJ, CHANDRAJIT L., EBEIDA, MOHAMED S., et al. “VoroCrust: Voronoi Meshing Without Clipping”. *ACM Trans. Graph.* 39.3 (2020), 23:1–23:16 8.
- [AHKS20] ALEXA, MARC, HERHOLZ, PHILIPP, KOHLBRENNER, MAXIMILIAN, and SORKINE-HORNUNG, OLGA. “Properties of Laplace Operators for Tetrahedral Meshes”. *Comput. Graph. Forum* 39.5 (2020), 55–68 1–3, 5–8.
- [Ale19] ALEXA, MARC. “Harmonic triangulations”. *ACM Trans. Graph.* 38.4 (July 2019). ISSN: 0730-0301 8.
- [Ale20] ALEXA, MARC. “Conforming weighted delaunay triangulations”. *ACM Trans. Graph.* 39.6 (Nov. 2020) 8.
- [BBA21] BUNGE, A., BOTSCH, M., and ALEXA, M. “The Diamond Laplace for Polygonal and Polyhedral Meshes”. *Computer Graphics Forum* 40.5 (2021), 217–230 7.
- [BKP\*10] BOTSCH, MARIO, KOBELT, LEIF, PAULY, MARK, et al. *Polygon Mesh Processing*. A K Peters, 2010 5, 8.



- [BS07] BOBENKO, ALEXANDER I. and SPRINGBORN, BORIS. “A Discrete Laplace-Beltrami Operator for Simplicial Surfaces”. *Discret. Comput. Geom.* 38.4 (2007), 740–756 6.
- [BXH10] BATTY, CHRISTOPHER, XENOS, STEFAN, and HOUSTON, BEN. “Tetrahedral Embedded Boundary Methods for Accurate and Flexible Adaptive Fluids”. *Comput. Graph. Forum* 29.2 (2010) 8.
- [BZC\*23] BENCHEKROUN, OTMAN, ZHANG, JIAYI ERIS, CHAUDHURI, SIDDARTHA, et al. “Fast Complementary Dynamics via Skinning Eigenmodes”. *ACM Trans. Graph.* 42.4 (2023), 106:1–106:21 3, 7.
- [CDH23] CHANG, QINGJUN, DENG, CHONGYANG, and HORMANN, KAI. “Maximum Likelihood Coordinates”. *Comput. Graph. Forum* 42.5 (2023), i–viii 8.
- [CKKD20] CHESHMI, KAZEM, KAUFMAN, DANNY M., KAMIL, SHOAB, and DEHNAVI, MARYAM MEHRI. “NASOQ: numerically accurate sparsity-oriented QP solver”. *ACM Trans. Graph.* 39.4 (2020), 96 6.
- [DHLM05] DESBRUN, MATHIEU, HIRANI, ANIL N., LEOK, MELVIN, and MARSDEN, JERROLD E. *Discrete Exterior Calculus*. 2005. arXiv: [math/0508341](https://arxiv.org/abs/math/0508341) [[math.DG](https://arxiv.org/abs/math/0508341)] 5.
- [dMMD14] DE GOES, FERNANDO, MEMARI, POORAN, MULLEN, PATRICK, and DESBRUN, MATHIEU. “Weighted Triangulations for Geometry Processing”. *ACM Trans. Graph.* 33.3 (2014) 8.
- [DPVA23] DIAZZI, LORENZO, PANOZZO, DANIELE, VAXMAN, AMIR, and ATTENE, MARCO. “Constrained Delaunay Tetrahedrization: A Robust and Practical Approach”. *ACM Trans. Graph.* 42.6 (2023), 181:1–181:15 7.
- [ESÜ04] EPPSTEIN, DAVID, SULLIVAN, JOHN M., and ÜNGÖR, ALPER. “Tiling space and slabs with acute tetrahedra”. *Comput. Geom.* 27.3 (2004), 237–255 5.
- [ETK\*07] ELCOTT, SHARIF, TONG, YIYING, KANSO, EVA, et al. “Stable, circulation-preserving, simplicial fluids”. *ACM Trans. Graph.* 26.1 (2007), 4 7.
- [FSSB07] FISHER, MATTHEW, SPRINGBORN, BORIS, SCHRÖDER, PETER, and BOBENKO, ALEXANDER I. “An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing”. *Computing* 81.2-3 (2007), 199–213 6.
- [Gli05] GLICKENSTEIN, DAVID. “Geometric triangulations and discrete Laplacians on manifolds”. (2005). arXiv: [math / 0508188](https://arxiv.org/abs/math/0508188) [[math.MG](https://arxiv.org/abs/math/0508188)] 5.
- [GSC21] GILLESPIE, MARK, SHARP, NICHOLAS, and CRANE, KEENAN. “Integer coordinates for intrinsic geometry processing”. *ACM Trans. Graph.* 40.6 (2021), 252:1–252:13 6.
- [HS08] HORMANN, KAI and SUKUMAR, N. “Maximum Entropy Coordinates for Arbitrary Polytopes”. *Comput. Graph. Forum* 27.5 (2008), 1513–1520 8.
- [HSW\*20] HU, YIXIN, SCHNEIDER, TESEO, WANG, BOLUN, et al. “Fast tetrahedral meshing in the wild”. *ACM Trans. Graph.* 39.4 (2020), 117 7.
- [HZG\*18] HU, YIXIN, ZHOU, QINGNAN, GAO, XIFENG, et al. “Tetrahedral meshing in the wild”. *ACM Trans. Graph.* 37.4 (2018), 60 7.
- [Jac\*24] JACOBSON, ALEC et al. *gptoolbox: Geometry Processing Toolbox*. <http://github.com/alecjacobson/gptoolbox>. 2024 8.
- [JMD\*07] JOSHI, PUSHKAR, MEYER, MARK, DEROSE, TONY, et al. “Harmonic coordinates for character articulation”. *ACM Trans. Graph.* 26.3 (2007), 71 8.
- [LGC\*23] LIU, HSUEH-TI DEREK, GILLESPIE, MARK, CHISLETT, BENJAMIN, et al. “Surface Simplification using Intrinsic Error Metrics”. *ACM Trans. Graph.* 42.4 (2023), 118:1–118:17 6.
- [Mat24] MATHWORKS. “MATLAB and Statistics Toolbox Release 2024a”. (2024) 6.
- [MDSB02] MEYER, MARK, DESBRUN, MATHIEU, SCHRÖDER, PETER, and BARR, ALAN H. “Discrete Differential-Geometry Operators for Triangulated 2-Manifolds”. *Vis. Math.* 4.3 (2002), 35–57 5, 6.
- [MKMD23] MITCHELL, SCOTT A., KNUPP, PATRICK M., MACKAY, SARAH, and DEAKIN, MICHAEL F. “Not so HOT Triangulations”. *Comput. Aided Des.* 158 (2023), 103497 8.
- [MMdGD11] MULLEN, PATRICK, MEMARI, POORAN, de GOES, FERNANDO, and DESBRUN, MATHIEU. “HOT: Hodge-optimized triangulations”. *ACM Trans. Graph.* 30.4 (2011) 1, 5, 8.
- [OBS\*12] OVSJANIKOV, MAKS, BEN-CHEN, MIRELA, SOLOMON, JUSTIN, et al. “Functional maps: a flexible representation of maps between shapes”. *ACM Trans. Graph.* 31.4 (2012), 30:1–30:11 3, 7.
- [Rus07] RUSTAMOV, RAIF M. “Laplace-Beltrami eigenfunctions for deformation invariant shape representation”. *Proc. SGP*. Ed. by BELYAEV, ALEXANDER G. and GARLAND, MICHAEL. Vol. 257. 2007, 225–233 3, 7.
- [SACO22] SHARP, NICHOLAS, ATTAIKI, SOUHAIB, CRANE, KEENAN, and OVSJANIKOV, MAKS. “DiffusionNet: Discretization Agnostic Learning on Surfaces”. *ACM Trans. Graph.* 41.3 (2022), 27:1–27:16 3, 7.
- [SBG\*20] STELLATO, BARTOLOMEO, BANJAC, GORAN, GOULART, PAUL, et al. “OSQP: an operator splitting solver for quadratic programs”. *Math. Program. Comput.* 12.4 (2020), 637–672 6.
- [Si15] SI, HANG. “TetGen, a Delaunay-Based Quality Tetrahedral Mesh Generator”. *ACM Trans. Math. Softw.* 41.2 (2015), 11:1–11:36 6, 7.
- [SMWF22] STRÖTER, DANIEL, MUELLER-ROEMER, JOHANNES SEBASTIAN, WEBER, DANIEL, and FELLNER, DIETER W. “Fast harmonic tetrahedral mesh optimization”. *Vis. Comput.* 38.9 (2022), 3419–3433 8.
- [SOG09] SUN, JIAN, OVSJANIKOV, MAKS, and GUIBAS, LEONIDAS J. “A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion”. *Comput. Graph. Forum* 28.5 (2009), 1383–1392 3, 7.
- [SSC19] SHARP, NICHOLAS, SOLIMAN, YOUSUF, and CRANE, KEENAN. “Navigating intrinsic triangulations”. *ACM Trans. Graph.* 38.4 (2019), 55:1–55:16 6.
- [SXZ\*17] SCHULZ, ADRIANA, XU, JIE, ZHU, BO, et al. “Interactive design space exploration and optimization for CAD models”. *ACM Trans. Graph.* 36.4 (2017), 157:1–157:14 7.
- [TTT03] TÜTÜNCÜ, REHA H., TOH, KIM-CHUAN, and TODD, MICHAEL J. “Solving semidefinite-quadratic-linear programs using SDPT3”. *Math. Program.* 95.2 (2003), 189–217 6.
- [VL08] VALLET, BRUNO and LÉVY, BRUNO. “Spectral Geometry Processing with Manifold Harmonics”. *Comput. Graph. Forum* 27.2 (2008), 251–260 3, 7.
- [WMKG07] WARDETZKY, MAX, MATHUR, SAURABH, KÄLBERER, FELIX, and GRINSPUN, EITAN. “Discrete laplace operators: no free lunch”. *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*. SGP '07. 2007, 33–37. ISBN: 9783905673463 5, 8.
- [ZDL\*14] ZHANG, JUYONG, DENG, BAILIN, LIU, ZISHUN, et al. “Local barycentric coordinates”. *ACM Trans. Graph.* 33.6 (2014), 188:1–188:12 8.
- [ZJ16] ZHOU, QINGNAN and JACOBSON, ALEC. “Thing10K: A Dataset of 10, 000 3D-Printing Models”. *CoRR* abs/1605.04797 (2016). arXiv: [1605.04797](https://arxiv.org/abs/1605.04797) 7.