

Kai Hormann, N. Sukumar

*Generalized Barycentric
Coordinates in
Computer Graphics and
Computational Mechanics*

This is a pre-publication version of content to appear in *Generalized Barycentric Coordinates in Computer Graphics and Computational Mechanics* (CRC Press, forthcoming 2017. All rights reserved.)

A primer on Laplacians

Max Wardetzky

Georg-August-University Göttingen, Germany

CONTENTS

5.1	Introduction	78
	5.1.1 Basic properties	78
5.2	Laplacians on Riemannian manifolds	79
	5.2.1 Exterior calculus	79
	5.2.2 Hodge decomposition	81
	5.2.3 The spectrum	82
5.3	Discrete Laplacians	83
	5.3.1 Laplacians on graphs	83
	5.3.2 The spectrum	84
	5.3.3 Laplacians on simplicial manifolds	86
	5.3.4 Strongly and weakly defined Laplacians	87
	5.3.5 Hodge decomposition	88
	5.3.6 The cotan Laplacian and beyond	88
	5.3.7 Discrete versus smooth Laplacians	90

IN THIS CHAPTER WE REVIEW SOME IMPORTANT PROPERTIES OF LAPLACIANS, SMOOTH AND DISCRETE. WE PLACE SPECIAL EMPHASIS ON A UNIFIED FRAMEWORK FOR TREATING SMOOTH LAPLACIANS ON RIEMANNIAN MANIFOLDS ALONGSIDE DISCRETE LAPLACIANS ON GRAPHS AND SIMPLICIAL MANIFOLDS. WE CAST THIS FRAMEWORK INTO THE LANGUAGE OF LINEAR ALGEBRA, WITH THE INTENT TO MAKE THIS TOPIC AS ACCESSIBLE AS POSSIBLE. WE COMBINE PERSPECTIVES FROM SMOOTH GEOMETRY, DISCRETE GEOMETRY, SPECTRAL ANALYSIS, MACHINE LEARNING, NUMERICAL ANALYSIS, AND GEOMETRY PROCESSING WITHIN THIS UNIFIED FRAMEWORK. THE CONNECTION TO GENERALIZED BARYCENTRIC COORDINATES IS ESTABLISHED THROUGH HARMONIC FUNCTIONS THAT INTERPOLATE GIVEN BOUNDARY CONDITIONS.

5.1 INTRODUCTION

The Laplacian is perhaps *the* prototypical differential operator for various physical phenomena. It describes, for example, heat diffusion, wave propagation, steady state fluid flow, and it is key to the Schrödinger equation in quantum mechanics. In Euclidean space, the Laplacian of a smooth function $u : \mathbb{R}^n \rightarrow \mathbb{R}$ is given as the sum of second partial derivatives along the coordinate axes,

$$\Delta u = - \left(\frac{\partial^2 u}{\partial x_1^2} + \frac{\partial^2 u}{\partial x_2^2} + \cdots + \frac{\partial^2 u}{\partial x_n^2} \right),$$

where we adopt the geometric perspective of using a minus sign.

5.1.1 Basic properties

The Laplacian has many intriguing properties. For the remainder of this exposition, consider an open and bounded domain $\Omega \subset \mathbb{R}^n$ and the L^2 inner product

$$(f, g) := \int_{\Omega} fg$$

on the linear space of square-integrable functions on Ω . Let $u, v : \Omega \rightarrow \mathbb{R}$ be two (sufficiently smooth) functions that vanish on the boundary of Ω . Then the Laplacian Δ is a *symmetric* (or, to be precise, a formally self-adjoint) linear operator with respect to this inner product since integration by parts yields

$$(u, \Delta v) = \int_{\Omega} \nabla u \cdot \nabla v = (\Delta u, v). \quad (\text{SYM})$$

Here ∇ denotes the standard gradient operator and $\nabla u \cdot \nabla v$ denotes the standard inner product between vectors in \mathbb{R}^n . The choice of using a minus sign in the definition of the Laplacian makes this operator positive semi-definite since

$$(u, \Delta u) = \int_{\Omega} \nabla u \cdot \nabla u \geq 0. \quad (\text{PSD})$$

If one restricts to functions that vanish on the boundary of Ω , (PSD) implies that the only functions that lie in the kernel of the Laplacian ($\Delta u = 0$) are those functions that vanish on the entire domain. Moreover, properties (SYM) and (PSD) imply that the Laplacian can be diagonalized and its eigenvalues are nonnegative,

$$\Delta u = \lambda u \Rightarrow \lambda \geq 0.$$

Another prominent property of smooth Laplacians is the *maximum principle*. Let $u : \Omega \rightarrow \mathbb{R}$ be *harmonic*, i.e., $\Delta u = 0$. The maximum principle asserts that

$$u \text{ is harmonic} \Rightarrow u \text{ has no strict local maximum in } \Omega, \quad (\text{MAX})$$

where we no longer assume that u vanishes on the boundary of Ω . Likewise, no harmonic function can have a local minimum in Ω .

The maximum principle can be derived from another important property of harmonic functions, the *mean value property*. Consider a point $\mathbf{x} \in \Omega$ and a closed ball $B(\mathbf{x}, r)$ of radius r centered at \mathbf{x} that is entirely contained in Ω . Every harmonic function has the property that the value $u(\mathbf{x})$ can be recovered from the average of the values of u in the ball $B(\mathbf{x}, r)$:

$$u(\mathbf{x}) = \frac{1}{\text{vol}(B(\mathbf{x}, r))} \int_{B(\mathbf{x}, r)} u(\mathbf{y}) d\mathbf{y}.$$

A simple argument by contradiction shows that the mean value property implies property (MAX).

The properties mentioned so far play an important role in applications; specifically, in the context of barycentric coordinates, they give rise to *harmonic coordinates* and *mean value coordinates*, see [26, 34, 35]. Below we discuss additional properties of Laplacians. For further reading we refer to the books [5, 25, 43] and the lecture notes [12, 17].

5.2 LAPLACIANS ON RIEMANNIAN MANIFOLDS

The standard Laplacian in \mathbb{R}^n can be expressed as

$$\Delta u = -\text{div } \nabla u,$$

where div is the usual divergence operator acting on vector fields in \mathbb{R}^n . Written in integral form, the negative divergence operator is the (formal) adjoint of the gradient: If X is a vector field and $u : \Omega \rightarrow \mathbb{R}$ is a function that vanishes on the boundary of Ω , then

$$\int_{\Omega} \nabla u \cdot X = \int_{\Omega} u (-\text{div } X).$$

This perspective can be generalized to Riemannian manifolds, which incorporate the notion of curvature. The Laplacian plays an important role in the study of these curved spaces.

5.2.1 Exterior calculus

Although gradient and divergence can readily be defined on Riemannian manifolds, it is more convenient to work with the differential (or *exterior derivative*) d instead of the gradient and with the codifferential d^* instead of divergence.

The differential d is similar to (but not the same as) the gradient. Indeed, given a function $u : \Omega \rightarrow \mathbb{R}$, one has

$$du(X) = \nabla u \cdot X$$

for every vector field X . In particular, the differential does not require the notion of a metric, whereas the gradient does. The codifferential d^* is defined as the formal adjoint (informally, transpose) to d , in the same way as divergence is the adjoint

of the gradient. In contrast to the divergence operator, which acts on vector fields, the codifferential d^* acts on 1-forms. A 1-form is a covector at every point of Ω , i.e., if X is a vector field on Ω and α is a 1-form, then $\alpha(X)$ is a real-valued function on Ω . In order to define the codifferential d^* , consider a 1-form α and a function $u : \Omega \rightarrow \mathbb{R}$ that vanishes on the boundary of Ω . Then

$$\int_{\Omega} du \cdot \alpha = \int_{\Omega} u d^* \alpha,$$

where the dot product is the inner product between covectors induced from the inner product between vectors. Notice that different from the differential d , the codifferential *does* require the notion of a metric. The Laplacian of a function u can be expressed as

$$\Delta u = d^* du,$$

which is equivalent to the representation $\Delta u = -\operatorname{div} \nabla u$ given above.

In order to carry over this framework to manifolds, let M be a smooth orientable manifold with smooth Riemannian metric g . Suppose for simplicity that M is compact and has empty boundary. The Riemannian metric induces a pointwise inner product between tangent vectors on M , which, analogously to the above discussion, induces an inner product between 1-forms. More generally, one works with k -forms for $k \geq 0$. A 0-form, by convention, is a real-valued function on M . A 1-form can be thought of as an oriented 1-volume in the sense that applying a 1-form to a vector field returns a real value at every point. Likewise, a k -form for $k > 1$ can be thought of as an oriented k -volume in the sense of returning a real number at every point when applied to an ordered k -tuple (parallelepiped) of tangent vectors. As a consequence, k -forms can be integrated over (sub)manifolds of dimension k . In the sequel we let Λ^k denote the linear space of k -forms on M .

Analogous to the L^2 inner product between function in \mathbb{R}^n , let

$$(\alpha, \beta)_k := \int_M g(\alpha, \beta) \operatorname{vol}_g$$

denote the L^2 inner product between k -forms α and β on M , where, by slight abuse of notation, we let $g(\alpha, \beta)$ denote the (pointwise) inner product induced by the Riemannian metric.

The differential $d : \Lambda^k \rightarrow \Lambda^{k+1}$ maps k -forms to $(k+1)$ -forms for $0 \leq k \leq \dim M$, where one sets $d\alpha = 0$ for any k -form with $k = \dim M$. One can define the differential acting on k -forms by postulating *Stokes' theorem*,

$$\int_U d\alpha = \int_{\partial U} \alpha,$$

for every k -form α and every (sufficiently smooth) submanifold $U \subset M$ of dimension $(k+1)$ with boundary ∂U . If one asserts this equality as the defining property of the differential d , then it immediately follows that $d \circ d = 0$ since the boundary of a boundary of a manifold is empty ($\partial(\partial U) = \emptyset$).

The codifferential d^* , taking $(k+1)$ -forms back to k -forms, is the (formal) adjoint

of d with respect to the L^2 inner products on k - and $(k+1)$ -forms. It is defined by requiring that

$$(d\alpha, \beta)_{k+1} = (\alpha, d^*\beta)_k$$

for all k -forms α and all $(k+1)$ -forms β . Finally, the *Laplace–Beltrami operator* $\Delta : \Lambda^k \rightarrow \Lambda^k$ acting on k -forms is defined as

$$\Delta\alpha := dd^*\alpha + d^*d\alpha.$$

Notice that this expression reduces to $\Delta u = d^*du$ for 0-forms (functions) on M . It follows almost immediately from the definition of the Laplacian that a k -form α is *harmonic* ($\Delta\alpha = 0$) if and only if α is closed ($d\alpha = 0$) and co-closed ($d^*\alpha = 0$).

From a structural perspective it is important to note that properties (SYM), (PSD), and (MAX) mentioned earlier remain true (among various other properties) in the setting of Riemannian manifolds. For further details on exterior calculus and the Laplace–Beltrami operator, we refer to [43].

5.2.2 Hodge decomposition

Every sufficiently smooth k -form α on M admits a unique decomposition

$$\alpha = d\mu + d^*\nu + h,$$

known as the *Hodge decomposition* (or Hodge–Helmholtz decomposition), where μ is a $(k-1)$ -form, ν is a $(k+1)$ -form and h is a harmonic k -form. This decomposition is unique and orthogonal with respect to the L^2 inner product on k -forms,

$$0 = (d\mu, d^*\nu)_k = (h, d\mu)_k = (h, d^*\nu)_k,$$

which immediately follows from the fact that $d \circ d = 0$ and the fact that harmonic forms satisfy $dh = d^*h = 0$. The Hodge decomposition can be thought of as a (formal) application of the well-known fact from linear algebra that the orthogonal complement of the kernel of a linear operator is equal to the range of its adjoint (transpose) operator.

By duality between vector fields and 1-forms, the Hodge decomposition for 1-forms carries over to a corresponding decomposition for vector fields into curl-free and divergence-free components, which has applications for fluid mechanics [3] and Maxwell’s equations for electromagnetism [29].

Geometrically, the Hodge decomposition establishes relations between the Laplacian and global properties of manifolds. Indeed, the linear space of harmonic k -forms is finite-dimensional for compact manifolds and isomorphic to $H^k(M; \mathbb{R})$, the k -th cohomology of M . As an application of this fact, consider a compact orientable surface without boundary. Then the dimension of the space of harmonic 1-forms is equal to twice the genus of the surface, that is, this dimension is zero for the 2-sphere, two for the two-dimensional torus, four for a genus two surface (pretzel) and so on. Hence the Laplacian provides global information about the topology of the underlying space.

For a thorough treatment of Hodge decompositions, including the case of manifolds with boundary, we refer to [45].

5.2.3 The spectrum

One cannot speak about the Laplacian without discussing its spectrum. On a compact orientable manifold without boundary, it follows from the inequality

$$(\Delta u, u)_0 = (du, du)_1 \geq 0$$

that the spectrum is nonnegative and that the only functions in the kernel of the Laplacian are constant functions. Thus zero is a trivial eigenvalue of the Laplacian with a one-dimensional space of eigenfunctions. The next (non-trivial) eigenvalue $\lambda_1 > 0$ is much more interesting. By the min-max principle, this eigenvalue satisfies

$$\lambda_1 = \min_{(u,1)_0=0} \frac{(du, du)_1}{(u, u)_0},$$

where one takes the minimum over all functions that are L^2 -orthogonal to the constants. Higher eigenvalues can be obtained by successively applying the min-max principle to the orthogonal complements of the eigenspaces of lower eigenvalues.

The first non-trivial eigenvalue already tells a great deal about the geometry of the underlying Riemannian manifold. As an example, consider Cheeger's isoperimetric constant

$$\lambda_C := \inf_N \left\{ \frac{\text{vol}_{n-1}(N)}{\min(\text{vol}_n(M_1), \text{vol}_n(M_2))} \right\},$$

where N runs over all compact codimension-1 submanifolds that partition M into two disjoint open sets M_1 and M_2 with $N = \partial M_1 = \partial M_2$. Intuitively, the optimal N for which λ_C is attained partitions M into two sets that have maximal volume and minimal perimeter. As an example, suppose that M has the shape of the surface of a smooth dumbbell. Then N is a curve going around the axis of the dumbbell at the location where the dumbbell is thinnest.

A relation of Cheeger's constant to the first non-trivial eigenvalue of the Laplacian is provided by the Cheeger inequalities

$$\frac{\lambda_C^2}{4} \leq \lambda_1 \leq c(K\lambda_C + \lambda_C^2),$$

where the constant c only depends on dimension and $K \geq 0$ provides a lower bound on the Ricci curvature of M in the sense that $\text{Ric}(M, g) \geq -K^2(n-1)$; see [11, 14]. Recall that for surfaces, Ricci curvature and Gauß curvature coincide. The first non-trivial eigenvalue of the Laplacian is thus related to the metric problem of minimal cuts—thus providing a relation between an analytical quantity (the first eigenvalue) and a purely geometric quantity (the Cheeger constant). Intuitively, if λ_1 is small, then M must have a small bottleneck; vice-versa, if λ_1 is large, then M is somewhat thick.

Equipped with the full set of eigenfunctions $\{\varphi_i\}$ of the Laplace–Beltrami operator, one can perform Fourier analysis on manifolds by decomposing any square-integrable function u into its modes,

$$u = \sum_i (u, \varphi_i)_0 \varphi_i,$$

provided that one chooses the eigenfunctions such that $(\varphi_i, \varphi_i) = \delta_{ij}$. (Notice that $(\varphi_i, \varphi_j)_0 = 0$ is automatic for eigenfunctions belonging to different eigenvalues.) The Fourier perspective is of great relevance in signal and geometry processing.

Maintaining a spectral eye on geometry, it is natural to ask the inverse question: How much *geometric* information can be reconstructed from information about the Laplacian? If the entire Laplacian is known on a smooth manifold, then one can reconstruct the metric, for example, by using the expression of Δ in local coordinates. If, however, “only” the spectrum is known, then less can be said in general. For example, Kac’s famous question *Can one hear the shape of a drum?* [36], that is, whether the entire geometry can be inferred from the spectrum alone, has a negative answer: There exist isospectral but non-isometric manifolds [31, 46].

5.3 DISCRETE LAPLACIANS

Discrete Laplacians can be defined on simplicial manifolds or, more generally, on graphs. We treat the case of graphs first and discuss simplicial manifolds further below.

5.3.1 Laplacians on graphs

Consider an undirected graph $\Gamma = (V, E)$ with vertex set V and edge set E . For simplicity we only consider finite graphs here. Suppose that every edge $e \in E$ between vertices $i \in V$ and $j \in V$ carries a real-valued weight $\omega_e = \omega_{ij} = \omega_{ji} \in \mathbb{R}$. We discuss below how weights can be chosen; suppose for now that such a choice has been made. A discrete Laplacian acting on a function $u : V \rightarrow \mathbb{R}$ is defined as

$$(Lu)_i := \sum_{j \sim i} \omega_{ij}(u_i - u_j), \quad (5.1)$$

where the sum ranges over all vertices j that are connected by an edge with vertex i .¹ This allows for representing the linear operator L as a matrix by

$$L_{ij} := \begin{cases} -\omega_{ij} & \text{if there is an edge between } i \text{ and } j, \\ \sum_{k \sim i} \omega_{ik} & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

The matrix L is called the discrete Laplace matrix. This definition may seem to come a bit out of the blue. In order to see how it relates to smooth Laplacians, consider again the smooth case and the quantity

$$E_D[u] := \frac{1}{2} \int_{\Omega} \|\nabla u\|^2,$$

¹Some authors include a division by vertex weights in the definition of Laplacians on graphs. Such a division arises naturally when considering strongly defined Laplacian, instead of weakly defined Laplacians. We cover this distinction below.

which is known as the *Dirichlet energy* of u . In the discrete setup, one may discretize the gradient ∇u along an edge $e = (i, j)$ as the finite difference $(u_i - u_j)$. Accordingly, one defines discrete Dirichlet energy as

$$E_D[u] := \frac{1}{2} \sum_{e \in E} \omega_{ij} (u_i - u_j)^2,$$

where the sum ranges over all edges. One then has

$$\text{(discrete) } E_D[u] = \frac{1}{2} u^T L u \quad \text{vs.} \quad \text{(smooth) } E_D[u] = \frac{1}{2} (u, \Delta u)_0,$$

which justifies calling L a Laplace matrix. Due to this representation (and using the language of partial differential equations [25]) we call discrete Laplacians of the form (5.1) *weakly defined* instead of *strongly defined*. We return to this distinction below.

Weakly defined discrete Laplacians of the form (5.1) are always symmetric—they satisfy (SYM) due to the assumption that $\omega_{ij} = \omega_{ji}$. However, whether or not a discrete Laplacian satisfies (at least some of) the other properties of the smooth setting heavily depends on the choice of weights.

The simplest choice of weights is to set $\omega_{ij} = 1$ whenever there is an edge between vertices i and j . This results in the so-called *graph Laplacian*. The diagonal entries of the graph Laplacian are equal to the degree of the respective vertex, that is, the number of edges adjacent to that vertex. The graph Laplacian is just a special case of what we call a *Laplacian on graphs* here.

Positive edge weights are a natural choice if weights resemble transition probabilities of a random walker. Discrete Laplacians with positive weights are always positive semi-definite (PSD) and, just like in the smooth setting, they only have the constant functions in their kernel provided that the graph is connected. As a word of caution we remark that positivity of weights is *not* necessary to guarantee (PSD). Below we discuss Laplacians that allow for (some) negative edge weights but still satisfy (PSD).

Laplacians with positive edge weights always satisfy the mean value property since every harmonic function u (a function for which $Lu = 0$) satisfies

$$u_i = \sum_{j \sim i} l_{ij} u_j \quad \text{with} \quad l_{ij} = \frac{\omega_{ij}}{L_{ii}} > 0.$$

Therefore, discrete Laplacians with positive weights also satisfy the maximum principle (MAX) since $\sum_{j \sim i} l_{ij} = 1$ and thus u_i is a convex combination of its neighbors u_j for discrete harmonic functions. This convex combination property establishes a connection between Laplacians and barycentric coordinates via the partition of unity property.

5.3.2 The spectrum

As in the smooth case, one cannot discuss discrete Laplacians without mentioning their spectrum and their eigenfunctions, which provide a fingerprint of the structure of the underlying graph.

As an example consider again Cheeger’s isoperimetric constant. In order to define this constant in the discrete setting consider a partitioning of Γ into two disjoint subgraphs Γ_1 and $\bar{\Gamma}_1$ such that the vertex set V of Γ is the disjoint union of the vertex sets V_1 of Γ_1 and \bar{V}_1 of $\bar{\Gamma}_1$. Here a subgraph of Γ denotes a graph whose vertex set is a subset of the vertex set of Γ such that two vertices in the subgraph are connected by an edge if and only if they are connected by an edge in Γ . For positive edge weights, the discrete Cheeger constant (sometimes called *conductance* of a weighted graph) is defined as

$$\lambda_C := \min \left\{ \frac{\text{vol}(\Gamma_1, \bar{\Gamma}_1)}{\min(\text{vol}(\Gamma_1), \text{vol}(\bar{\Gamma}_1))} \right\},$$

where

$$\text{vol}(\Gamma_1, \bar{\Gamma}_1) := \sum_{i \in V_1, j \notin V_1} \omega_{ij} \quad \text{and} \quad \text{vol}(\Gamma_1) := \sum_{i \in V_1, j \in V_1} \omega_{ij},$$

and similarly for $\text{vol}(\bar{\Gamma}_1)$. Notice that for the case of the graph Laplacian $\text{vol}(\Gamma_1)$ equals twice the number of edges in Γ_1 and $\text{vol}(\Gamma_1, \bar{\Gamma}_1)$ equals the number of edges with one vertex in Γ_1 and another vertex in its complement.²

Similar to the smooth case, one then obtains the Cheeger inequalities

$$\frac{\lambda_C^2}{2} \leq \tilde{\lambda}_1 \leq 2\lambda_C,$$

where $\tilde{\lambda}_1$ is the first nontrivial eigenvalue of the rescaled Laplace matrix

$$\tilde{L} := SLS,$$

where S is a diagonal matrix with $S_{ii} = 1/\sqrt{\omega_{ii}}$. This rescaling is necessary since λ_C is invariant under a uniform rescaling of edge weights (and so is \tilde{L}), whereas L scales linearly with the edge weights. A proof of the discrete Cheeger inequalities for the case of the graph Laplacians ($\omega_{ij} = 1$) can be found in [16]; the proof for arbitrary positive edge weights is nearly identical.

The Cheeger constant—and alongside the corresponding partitioning of Γ into the two disjoint subgraphs Γ_1 and $\bar{\Gamma}_1$ —has applications in graph clustering, since the edges connecting Γ_1 and $\bar{\Gamma}_1$ tend to “cut” the graph along its bottleneck [37].

Any discrete Laplacian (having positive edge weights or not) that satisfies (SYM) and (PSD) can be used for Fourier analysis on graphs. Indeed, let $\{\varphi_i\}$ be the eigenfunctions of L , chosen such that $\varphi_i^T \varphi_j = \delta_{ij}$. Then one has

$$u = \sum_i (u^T \varphi_i) \varphi_i,$$

just like in the smooth setting. This decomposition of discrete functions on graphs

²Some authors use a different version of the respective volumes in the definition of the Cheeger constant, resulting in different versions of the Cheeger inequalities; see [47].

into their Fourier modes has a plethora of applications in geometry processing; see [39] and references therein.

As in the smooth setting, it is natural to ask the inverse question of how much geometric information can be inferred from information about the Laplacian. Recall that in the smooth case, knowing the full Laplacian allows for recovering the Riemannian metric. Similarly, in the discrete case it can be shown that for simplicial surfaces the knowledge of the cotan weights for discrete Laplacian (covered below) allows for reconstructing edge lengths (i.e., the discrete metric) of the underlying mesh up to a global scale factor [53]. If, however, only the spectrum of the Laplacian is known, then there exist isospectral but non-isomorphic graphs [10]. In fact, for the case of the cotan Laplacian (see below), the exact same isospectral domains considered in [31] that were originally proposed for showing that “One cannot hear the shape of a drum” for smooth Laplacians work in the discrete setup.

A curious fact concerning the connection between discrete Laplacians and the underlying geometry is *Rippa’s theorem* [42]: The Delaunay triangulation of a fixed point set in \mathbb{R}^n minimizes the Dirichlet energy of any piecewise linear function over this point set. In [15], this result is taken a step further, where the authors show that the spectrum of the cotan Laplacian obtains its minimum on a Delaunay triangulation in the sense that the i -th eigenvalue of the cotan Laplacian of any triangulation of a fixed point set in the plane is bounded below by the i -th eigenvalue resulting from the cotan Laplacian associated with the Delaunay triangulation of the given point set.

5.3.3 Laplacians on simplicial manifolds

Recall that in the smooth case, Laplacians acting on k -forms take the form

$$\Delta = dd^* + d^*d.$$

In order to mimic this construction in the discrete setting, one requires a bit more structure than just an arbitrary graph. To this end, consider a simplicial manifold, such as a triangulated surface. We keep referring to this manifold as M . As in the smooth case, for simplicity, suppose that M is orientable and has no boundary.

Simplicial manifolds allow for a natural definition of discrete k -forms as duals of k -cells. Indeed, every 0-form is a function defined on vertices, a 1-form α is dual to edges, that is, $\alpha(e)$ is a real number for any oriented 1-cell (oriented edge), a 2-form is dual to oriented 2-cells, and so forth. In the sequel, let the linear space of discrete k -forms (better known as simplicial cochains) be denoted by C^k .

As in the smooth case, the discrete differential δ (better known as the coboundary operator) maps discrete k -forms to discrete $(k + 1)$ -forms, $\delta : C^k \rightarrow C^{k+1}$. Again as in the smooth case, the discrete differential can be defined by *postulating* Stokes’ formula: Let α be a discrete k -form. Then one requires that

$$\delta\alpha(\sigma) = \alpha(\partial\sigma)$$

for all $(k + 1)$ -cells σ , where ∂ denotes the simplicial boundary operator. The

simplicial boundary operator, when applied to a vertex returns zero (since vertices do not have a boundary). When applied to an oriented edge, the boundary operator returns the difference between the edge's vertices. Likewise, ∂ applied to an oriented 2-cell σ returns the sum of oriented edges of σ (with the orientation induced by that of σ). Since the boundary of a boundary is empty ($\partial \circ \partial = 0$) one has $\delta \circ \delta = 0$, just like in the smooth case. Notice that the definition of δ does not require the notion of inner products.

In order to define the discrete codifferential δ^* one additionally requires an inner product $(\cdot, \cdot)_k$ on the linear space of k -forms for each k . Below we discuss the construction of such inner products. Given a fixed choice of inner products on discrete k -forms, the codifferential is defined by requiring that

$$(\delta\alpha, \beta)_{k+1} = (\alpha, \delta^*\beta)_k$$

for all k -forms α and all $(k+1)$ -forms β , and the discrete *strongly defined* Laplacian acting on k -forms takes the form

$$\mathbb{L} := \delta\delta^* + \delta^*\delta. \quad (5.2)$$

This perspective is that of *discrete exterior calculus* (DEC) [17, 19], where—by slight abuse of notation—inner products are referred to as “discrete Hodge stars”.

Strongly defined Laplacians are self-adjoint with respect to the inner products $(\cdot, \cdot)_k$ on discrete k -forms, since

$$(\mathbb{L}\alpha, \beta)_k = (\delta\alpha, \delta\beta)_{k+1} + (\delta^*\alpha, \delta^*\beta)_{k-1} = (\alpha, \mathbb{L}\beta)_k.$$

Moreover, strongly defined Laplacians are always positive semi-definite (PSD), since

$$(\mathbb{L}\alpha, \alpha)_k = (\delta\alpha, \delta\alpha)_{k+1} + (\delta^*\alpha, \delta^*\alpha)_{k-1} \geq 0.$$

In particular, a discrete k -form is harmonic ($\mathbb{L}\alpha = 0$) if and only if $\delta\alpha = \delta^*\alpha = 0$, just like in the smooth setting.

5.3.4 Strongly and weakly defined Laplacians

Every strongly defined Laplacian as given by (5.2) has a weakly defined cousin L acting on discrete functions u . The weak version is obtained by requiring that at every vertex i the resulting function Lu is equal to

$$(Lu)_i := (\mathbb{L}u, 1_i)_0 = (\delta^*\delta u, 1_i)_0 = (\delta u, \delta 1_i)_1,$$

where 1_i is the indicator function of vertex i . In particular, let \mathbb{M}_0 and \mathbb{M}_1 be the symmetric positive definite matrices that encode the inner products between 0-forms and 1-forms, respectively,

$$(u, v)_0 = u^T \mathbb{M}_0 v \quad \text{and} \quad (\alpha, \beta)_1 = \alpha^T \mathbb{M}_1 \beta.$$

Then the weakly and strongly defined Laplacians satisfy, respectively,

$$L = \delta^T \mathbb{M}_1 \delta \quad \text{and} \quad \mathbb{L} = \mathbb{M}_0^{-1} L.$$

As an example, consider diagonal inner products on 0-forms and 1-forms,

$$(u, v)_0 = \sum_{i \in V} m_i u_i v_i \quad \text{and} \quad (\alpha, \beta)_1 = \sum_{e \in E} \omega_e \alpha(e) \beta(e),$$

with positive vertex weights $m_i > 0$ and positive edge weights $\omega_e > 0$. The resulting strongly defined Laplacian acting on 0-forms (functions) takes the form

$$(\mathbb{L}u)_i = \frac{1}{m_i} \sum_{j \sim i} \omega_{ij} (u_i - u_j)$$

and its associated weakly defined cousin is the Laplacians on graphs defined in (5.1). In particular, if $\omega_e = 1$, one recovers the graph Laplacian as the weak version.

5.3.5 Hodge decomposition

Given a choice of inner products for k -forms on simplicial manifolds, one always obtains a discrete Hodge decomposition. Indeed, for every discrete k -form α one has

$$\alpha = \delta\mu + \delta^*\nu + h,$$

where μ is a $(k-1)$ -form, ν is a $(k+1)$ -form and h is a harmonic k -form ($\mathbb{L}h = 0$). As in the smooth case, not only is this decomposition unique, it is also orthogonal with respect to the inner products on k -forms,

$$0 = (\delta\mu, \delta^*\nu)_k = (h, \delta\mu)_k = (h, \delta^*\nu)_k,$$

which immediately follows from $\delta \circ \delta = 0$ and the fact that harmonic forms satisfy $\delta h = \delta^* h = 0$.

Akin to the smooth case, the Hodge decomposition establishes relations to global properties of simplicial manifolds, since the linear space of harmonic k -forms is isomorphic to the k -th simplicial cohomology of the simplicial manifold M . Again, as an application of this fact, consider a compact simplicial surface without boundary. Then the dimension of the space of harmonic 1-forms is equal to twice the genus of the surface—independent of the concrete choice of inner products on k -forms.

5.3.6 The cotan Laplacian and beyond

We conclude the discussion of Laplacians on simplicial manifolds by providing an important example of inner products. In [51], Whitney constructs a map from simplicial k -forms (k -cochains) to piecewise linear differential k -forms. In a nutshell, the idea is to linearly interpolate simplicial k -forms across full-dimensional cells. As the simplest example, consider linear interpolation of 0-forms (functions) on vertices. This kind of interpolation can be extended to arbitrary k -forms. The resulting map

$$W : C^k \rightarrow L^2 \Lambda^k$$

takes simplicial k -forms to square-integrable k -forms on the simplicial manifold, where we assume each simplex to carry the standard Euclidean structure. The Whitney map is the right inverse of the so-called de Rham map,

$$\alpha(\sigma) = \int_{\sigma} W(\alpha)$$

for all discrete k -forms α and all k -cells σ . For details we refer to [51]. The Whitney map W is a chain map—it commutes with the differential ($dW = W\delta$) and thus factors to cohomology.

Dodziuk and Patodi [22] use the Whitney map in order to define an inner product on discrete k -forms (k -cochains) by

$$(\alpha, \beta)_k := \int_M g(W\alpha, W\beta) \text{vol}_g,$$

where (in our case) g denotes a piecewise Euclidean metric on the simplicial manifold. From the perspective of the *Finite Element Method* (FEM), Whitney’s construction is a special case of constructing stable finite elements; see [2].

For triangle meshes the resulting strongly defined Laplacian acting on 0-forms (functions) takes the form

$$\mathbb{L} = \mathbb{M}_0^{-1}L,$$

where \mathbb{M}_0 is the *mass matrix* given by

$$(\mathbb{M}_0)_{ij} := \begin{cases} \frac{A_{ij}}{12} & \text{if } i \sim j, \\ \frac{A_i}{6} & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Here A_{ij} denotes the combined area of the two triangles incident to edge (i, j) and A_i is the combined area of all triangles incident to vertex i . The corresponding weakly defined Laplacian L is the *stiffness matrix* with entries

$$L_{ij} := \begin{cases} -\frac{1}{2}(\cot \alpha_{ij} + \cot \beta_{ij}) & \text{if } i \sim j, \\ -\sum_{j \sim i} L_{ij} & \text{if } i = j, \\ 0 & \text{otherwise,} \end{cases}$$

where α_{ij} and β_{ij} are the two angles opposite to edge (i, j) . The matrix L is often referred to as the *cotan Laplacian*.

The cotan Laplacian has been rediscovered many times in different contexts [23, 24, 41]; the earliest explicit mention seems to go back to MacNeal [40], but perhaps it was already known at the time of Courant. The cotan Laplacian has been enjoying a wide range of applications in geometry processing (see [17, 39, 44] and references therein), including barycentric coordinates, mesh parameterization, mesh compression, fairing, denoising, spectral fingerprints, shape clustering, shape matching, physical simulation of thin structures, and geodesic distance computation.

The construction of discrete Laplacians based on inner products can be extended from simplicial surfaces to meshes with (not necessarily planar) polygonal faces; see [1] for such an extension that is similar to the approach considered in [8, 9] for planar polygons. The cotan Laplacian has furthermore been extended to semi-discrete surfaces [13] as well as to subdivision surfaces [18].

5.3.7 Discrete versus smooth Laplacians

Structural properties of discrete Laplacians play an important role for computations, e.g., when solving the Poisson problem

$$\mathbb{L}u = f$$

for a given right hand side f and an unknown function u with Dirichlet boundary conditions. For solving this problem, one often prefers to work with the weak formulation of the problem ($Lu = \mathbb{M}_0 f$) instead of the strong one ($\mathbb{L}u = f$) since the weakly defined cousins of strongly defined Laplacians satisfy properties (SYM) and (PSD), which allows for efficient linear solvers for this problem.

We now turn to the question of which properties are desirable for discrete Laplacians on top of (SYM) and (PSD) and if it is furthermore possible to recover *all* properties of smooth Laplacians in the discrete case. We follow [50], on which this section is largely based.

Smooth Laplacians are differential operators that act *locally*. Locality can be represented in the discrete case by working with (weakly defined) discrete Laplacians based on edge weights:

$$\text{vertices } i \text{ and } j \text{ do not share an edge} \Rightarrow \omega_{ij} = 0. \quad (\text{LOC})$$

This property reflects locality of action by ensuring that if vertices i and j are not connected by an edge, then changing the function value u_j at a vertex j does not alter the value $(Lu)_i$ at vertex i . Property (LOC) results in sparse matrices, which can be treated efficiently in computations.

Keeping an eye on the relation between discrete Laplacians and barycentric coordinates, it is desirable (or even mandatory) to require *linear reproduction*. In the smooth setting, linear reproduction corresponds to the fact that linear functions on \mathbb{R}^n are in the kernel of the standard Laplacian on Euclidean domains. For discrete Laplacians on a graph Γ , linear reproduction means that $(Lu)_i = 0$ at each interior vertex whenever Γ is embedded into the plane with straight edges and u is a *linear* function on the plane, point-sampled at the vertices of Γ ,

$$\Gamma \subset \mathbb{R}^2 \text{ embedded and } u : \mathbb{R}^2 \rightarrow \mathbb{R} \text{ linear} \Rightarrow (Lu)_i = 0 \text{ at interior vertices.} \quad (\text{LIN})$$

In applications this property is desirable for de-noising of surface meshes [20] (where one expects to remove normal noise only but not to introduce tangential vertex drift), mesh parameterization [27] (where one expects planar regions to remain invariant under parameterization), and plate bending energies [44] (which must vanish for flat configurations).

Furthermore, it is often natural and desirable to require positive edge weights:

$$\text{vertices } i \text{ and } j \text{ share an edge} \Rightarrow \omega_{ij} > 0. \quad (\text{POS})$$

This requirement implies (PSD) and is a sufficient (but by no means necessary) condition for a *discrete maximum principle* (MAX). In diffusion problems corresponding to $u_t = -\Delta u$, (POS) ensures that flow travels from regions of higher potential to regions of lower potential. Additionally, as discussed above, positive edge weights establish a connection to barycentric coordinates.

The combination (LOC)+(SYM)+(POS) is related to Tutte's embedding theorem for planar graphs [32, 48]: Positive weights associated to edges yield a straight-line embedding of an abstract planar graph for a fixed convex boundary polygon. Tutte's embedding is unique for a given set of positive edge weights, and it satisfies (LIN) by construction since each interior vertex (and therefore its x - and y -coordinate) is a convex combination of its adjacent vertices with respect to the given edge weights.

It is natural to ask, which of the properties (LOC), (SYM), (POS), and (LIN) are satisfied by the discrete Laplacians considered so far.

As the simplest example consider again the graph Laplacian ($\omega_{ij} = 1$). This Laplacian clearly satisfies (LOC)+(SYM)+(POS), but in general fails to satisfy (LIN). Next, consider the cotan Laplacian. The edge weights of the cotan Laplacian turn out to be a special case of weights arising from *orthogonal duals*. To define those, consider a graph embedded into the plane with straight edges that do not cross. An orthogonal dual is a realization of the dual graph in the plane, with straight edges orthogonal to primal edges (viewed as vectors in the plane). Edge weights can then be constructed as the ratio between the *signed* lengths of dual edges and the unsigned lengths of primal edges,

$$\omega_e = \frac{|\star e|}{|e|}.$$

Here, $|e|$ denotes the usual Euclidean length, whereas $|\star e|$ denotes the *signed* Euclidean length of the dual edge. The sign is obtained as follows. The dual edge $\star e$ connects two dual vertices $\star f_1$ and $\star f_2$, corresponding to the primal faces f_1 and f_2 , respectively. The sign of $|\star e|$ is positive if along the direction of the ray from $\star f_1$ to $\star f_2$, the primal face f_1 lies before f_2 . The sign is negative otherwise.

For the cotan Laplacian, the dual graph is obtained by connecting circumcentres of (primal) triangles by straight edges; see Figure 5.1 (left). More generally, discrete Laplacians derived from orthogonal duals on arbitrary (including non-planar) triangular surfaces were considered in [30]. These Laplacians always satisfy (LOC)+(SYM)+(LIN) for the case of planar primal graphs, but fail to satisfy (POS) in general. Indeed, for the cotan Laplacian one has $(\cot \alpha_{ij} + \cot \beta_{ij}) > 0$ if and only if $(\alpha_{ij} + \beta_{ij}) < \pi$. This is the case for all edges (i, j) if and only if the triangulation is Delaunay. One may restore (POS) by successive edge flips (thereby changing the combinatorics) until one arrives at a Delaunay triangulation [6]. Unfortunately, the number of required edge flips to obtain a Delaunay triangulation from an arbitrary given triangulation cannot be bounded a priori. Therefore, this approach fails to satisfy locality in general.




Figure 5.1 *Left:* Primal graph (solid lines) and orthogonal circumcentric dual graph (dashed lines) defining the cotan Laplacian. *Middle:* Mean value weights correspond to dual edges tangent to the unit circle around a primal vertex. *Right:* The projection of the Schönhardt polytope does not allow for a discrete Laplacian satisfying (SYM)+(LOC)+(LIN)+(POS).

More generally, so-called *weighted Delaunay triangulations* turn out to be the only triangulations that possess positive orthogonal duals and thus admit discrete Laplacians that satisfy (LOC)+(SYM)+(LIN)+(POS). Like Rippa's theorem (see above) this fact provides an instance of the intricate connection between properties of discrete differential operators (the Laplacian) and purely geometric properties (weighted Delaunay triangulations).

Finally, if one drops the requirement of symmetric edge weights, then one enters the realm of barycentric coordinates. In this case one may still obtain an orthogonal *dual face per primal vertex*, but these dual faces no longer fit into a consistent dual graph; see Figure 5.1 (middle). Hence, for the case dual edges with positive lengths, one obtains edge weights satisfying (LOC)+(LIN)+(POS) but not (SYM). Floater et al. [28] explored a subspace of this case: a one-parameter family of linear precision barycentric coordinates, including mean value and Wachspress coordinates. Langer et al. [38] showed that each member of this family corresponds to a specific choice of orthogonal dual face per primal vertex.

The fact that no discrete Laplace operator satisfies all of the desired properties simultaneously is not a coincidence. It can be shown that general simplicial meshes do not allow for discrete Laplacians that satisfy (LOC)+(SYM)+(LIN)+(POS); see [50] for more on this topic and Figure 5.1 (right) for a simple example. This limitation provides a taxonomy on existing literature and explains the plethora of existing discrete Laplacians: Since not all desired properties can be fulfilled simultaneously, it depends on the application at hand to design discrete Laplacians that are tailored towards the specific needs of a concrete problem.

Another important desideratum is *convergence*: In the limit of refinement of simplicial manifolds that approximate a smooth manifold, one seeks to approximate the smooth Laplacian by a sequence of discrete ones. For applications this is important in terms of obtaining discrete operators that are as mesh independent as possible—re-meshing a given shape should not result in a drastically different Laplacian.

A closely related concept to convergence is *consistency*. A sequence of discrete Laplacians $(\Delta_n)_{n \in \mathbb{N}}$ is called consistent, if $\Delta_n u \rightarrow \Delta u$ for all appropriately chosen

functions u . For example, it can be shown that Laplacians on point clouds, such as those considered in [4] are consistent; see [21].

Convergence is more difficult to show than consistency since it additionally requires that the solutions u_n to the Poisson problems $\Delta_n u_n = f$ converge (in an appropriate norm) to the solution u of $\Delta u = f$. Discussing convergence in detail is beyond the scope of this short survey. Roughly speaking, Laplacians on simplicial manifolds converge to their smooth counterparts (in an appropriate operator norm) if the inner products on discrete k -forms used for defining simplicial Laplacians converge to the inner products on smooth k -forms. In this case, one obtains convergence of solutions to the Poisson problem, convergence of the components of the Hodge decomposition, convergence of eigenvalues and eigenfunctions [22, 24, 33, 49, 52], and (using different techniques) convergence of Cheeger cuts [7].

Bibliography

- [1] M. Alexa and M. Wardetzky. Discrete Laplacians on general polygonal meshes. *ACM Transactions on Graphics*, 30(4):102:1–102:10, 2011.
- [2] D. N. Arnold, R. S. Falk, and R. Winther. Finite element exterior calculus, homological techniques, and applications. *Acta Numerica*, 15:1–155, 2006.
- [3] V. Arnold and B. Khesin. *Topological Methods in Hydrodynamics*, volume 125 of *Applied Mathematical Sciences*. Springer, New York, 1998.
- [4] M. Belkin, J. Sun, and Y. Wang. Constructing Laplace operator from point clouds in \mathbb{R}^d . In *Proceedings of SODA*, pages 1031–1040, 2009.
- [5] M. Berger. *A Panoramic View of Riemannian Geometry*. Springer, Berlin-Heidelberg, 2003.
- [6] A. I. Bobenko and B. A. Springborn. A discrete Laplace–Beltrami operator for simplicial surfaces. *Discrete and Computational Geometry*, 38(4):740–756, 2007.
- [7] X. Bresson, N. G. Trillos, D. Slepčev, J. von Brecht, and T. Laurent. Consistency of Cheeger and ratio graph cuts. 2016. To appear in JMLR.
- [8] F. Brezzi, K. Lipnikov, and M. Shashkov. Convergence of mimetic finite difference methods for diffusion problems on polyhedral meshes. *SIAM J. Num. Anal.*, 43:1872–1896, 2005.
- [9] F. Brezzi, K. Lipnikov, and V. Simoncini. A family of mimetic finite difference methods on polygonal and polyhedral meshes. *Math. Models Methods Appl. Sci.*, 15:1533–1553, 2005.
- [10] R. Brooks. Isospectral graphs and isospectral surfaces. *Séminaire de théorie spectrale et géométrie*, 15:105–113, 1997.
- [11] P. Buser. A note on the isoperimetric constant. *Annales scientifiques de l’École Normale Supérieure*, 15(2):213–230, 1982.
- [12] Y. Canzani. Analysis on manifolds via the Laplacian, 2013. Lecture notes, Harvard University.
- [13] W. Carl. A Laplace operator on semi-discrete surfaces. *Foundations of Computational Mathematics*, 2015. doi:10.1007/s10208-015-9271-y.

96 ■ Bibliography

- [14] I. Chavel. *Eigenvalues in Riemannian Geometry*. Academic Press, Orlando, San Diego, New York, London, Toronto, Montreal, Sydney, Tokyo, 1984.
- [15] R. Chen, Y. Xu, C. Gotsman, and L. Liu. A spectral characterization of the Delaunay triangulation. *Computer Aided Geometric Design*, 27(4):295–300, 2010.
- [16] F. Chung. Four proofs for the Cheeger inequality and graph partition algorithms. In *Proceedings of the ICCM*, 2007.
- [17] K. Crane, F. de Goes, M. Desbrun, and P. Schröder. Digital geometry processing with discrete exterior calculus. In *ACM SIGGRAPH 2013 Courses*, SIGGRAPH '13. ACM, 2013.
- [18] F. de Goes, M. Desbrun, M. Meyer, and T. DeRose. Subdivision exterior calculus for geometry processing. *ACM Transactions on Graphics*, 35(4):133:1–133:11, 2016.
- [19] M. Desbrun, A. Hirani, M. Leok, and J. E. Marsden. Discrete Exterior Calculus. 2005. arXiv:math.DG/0508341.
- [20] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of ACM SIGGRAPH*, pages 317–324, 1999.
- [21] T. Dey, P. Ranjan, and Y. Wang. Convergence, stability, and discrete approximation of Laplace spectra. In *Proceedings of SODA*, pages 650–663, 2010.
- [22] J. Dodziuk and V. K. Patodi. Riemannian structures and triangulations of manifolds. *Journal of Indian Math. Soc.*, 40:1–52, 1976.
- [23] R. Duffin. Distributed and Lumped Networks. *Journal of Mathematics and Mechanics*, 8:793–825, 1959.
- [24] G. Dziuk. Finite elements for the Beltrami operator on arbitrary surfaces. In S. Hildebrandt and R. Leis, editors, *Partial Differential Equations and Calculus of Variations*, volume 1357 of *Lec. Notes Math.*, pages 142–155. Springer, Berlin Heidelberg, 1988.
- [25] L. C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. AMS, Providence, Rhode Island, 1998.
- [26] M. S. Floater. Mean value coordinates. *Comput. Aided Geom. Des.*, 20(1):19–27, 2003.
- [27] M. S. Floater and K. Hormann. Surface parameterization: A tutorial and survey. In N. A. Dodgson, M. S. Floater, and M. A. Sabin, editors, *Advances in Multiresolution for Geometric Modeling*, pages 157–186. Springer, Berlin, Heidelberg, 2005.

- [28] M. S. Floater, K. Hormann, and G. Kós. A general construction of barycentric coordinates over convex polygons. *Adv. Comp. Math.*, 24(1–4):311–331, 2006.
- [29] T. Frankel. *The Geometry of Physics: An Introduction*. Cambridge University Press, 2004.
- [30] D. Glickenstein. Geometric triangulations and discrete Laplacians on manifolds. 2005. arxiv:math.MG/0508188.
- [31] C. Gordon, D. L. Webb, and S. Wolpert. One cannot hear the shape of a drum. *Bull. Amer. Math. Soc*, 27:134–138, 1992.
- [32] S. J. Gortler, C. Gotsman, and D. Thurtson. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Computer Aided Geometric Design*, 23(2):83–112, 2006.
- [33] K. Hildebrandt, K. Polthier, and M. Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometricae Dedicata*, 123:89–112, 2006.
- [34] P. Joshi, M. Meyer, T. DeRose, B. Green, and T. Sanocki. Harmonic coordinates for character articulation. *ACM Trans. Graph.*, 26(3):71:1–71:10, 2007.
- [35] T. Ju, S. Schaefer, and J. Warren. Mean value coordinates for closed triangular meshes. *ACM Trans. Graph.*, 24(3):561–566, 2005.
- [36] M. Kac. Can one hear the shape of a drum? *The American Mathematical Monthly*, 73(4):1–23, 1966.
- [37] R. Kannan, S. Vempala, and A. Vetta. On clusterings: Good, bad and spectral. *Journal of the ACM*, 51(3):497–515, 2004.
- [38] T. Langer, A. Belyaev, and H.-P. Seidel. Spherical barycentric coordinates. In *Siggraph/Eurographics Sympos. Geom. Processing*, pages 81–88, 2006.
- [39] B. Lévy and H. R. Zhang. Spectral mesh processing. In *ACM SIGGRAPH 2010 Courses, SIGGRAPH '10*. ACM, 2010.
- [40] R. MacNeal. The Solution of Partial Differential Equations by Means of Electrical Networks, 1949. PhD thesis, California Institute of Technology.
- [41] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experim. Math.*, 2:15–36, 1993.
- [42] S. Rippa. Minimal roughness property of the Delaunay triangulation. *Computer Aided Geometric Design*, 7(6):489–497, 1990.
- [43] S. Rosenberg. *The Laplacian on a Riemannian manifold*. Number 31 in Student Texts. London Math. Soc., 1997.

98 ■ Bibliography

- [44] M. Rumpf and M. Wardetzky. Geometry processing from an elastic perspective. *GAMM Mitteilungen*, 37(2):184–216, 2014.
- [45] G. Schwarz. *Hodge Decomposition - A Method for Solving Boundary Value Problems*, volume 1607 of *Lecture Notes in Mathematics*. Springer, Berlin Heidelberg, 1995.
- [46] T. Sunada. Riemannian coverings and isospectral manifolds. *Ann. of Math.*, 121(1):169–186, 1985.
- [47] T. Sunada. Discrete geometric analysis. In *Proc. Symposia in Pure Mathematics*, volume 77, pages 51–83, 2007.
- [48] W. T. Tutte. How to draw a graph. *Proc. London Math. Soc.*, 13(3):743–767, 1963.
- [49] M. Wardetzky. Discrete Differential Operators on Polyhedral Surfaces—Convergence and Approximation. Ph.D. Thesis, Freie Universität Berlin, 2006.
- [50] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete Laplace operators: No free lunch. In *Siggraph/Eurographics Sympos. Geom. Processing*, pages 33–37, 2007.
- [51] H. Whitney. *Geometric Integration Theory*. Princeton Univ. Press, 1957.
- [52] S. O. Wilson. Cochain algebra on manifolds and convergence under refinement. *Topology and its Applications*, 154:1898–1920, 2007.
- [53] W. Zeng, R. Guo, F. Luo, and X. Gu. Discrete heat kernel determines discrete Riemannian metric. *Graphics Models*, 74(4):121–129, July 2012.