

Localized Manifold Harmonics for Spectral Shape Analysis

S. Melzi¹, E. Rodolà^{2,3}, U. Castellani¹, M. M. Bronstein^{2,4,5,6}

¹University of Verona, Italy

²USI Lugano, Switzerland

³Sapienza University of Rome, Italy

⁴Tel Aviv University, Israel

⁵Intel Perceptual Computing, Israel

⁶IAS, TU Munich, Germany




Figure 1: A few Localized Manifold Harmonics (LMH) on two different regions of the dog shape. By changing the region location on the surface, our model provides an ordered set of localized harmonic functions (i.e., defined on the entire surface, but strongly concentrated on the selected region). In this figure the localized harmonics are clearly visible across different frequencies. The LMH constitute a valid alternative to the classical manifold harmonics and can be used in conjunction with those, or as a drop-in replacement in typical spectral shape analysis tasks.

Abstract

The use of Laplacian eigenfunctions is ubiquitous in a wide range of computer graphics and geometry processing applications. In particular, Laplacian eigenbases allow generalizing the classical Fourier analysis to manifolds. A key drawback of such bases is their inherently global nature, as the Laplacian eigenfunctions carry geometric and topological structure of the entire manifold. In this paper, we introduce a new framework for local spectral shape analysis. We show how to efficiently construct localized orthogonal bases by solving an optimization problem that in turn can be posed as the eigendecomposition of a new operator obtained by a modification of the standard Laplacian. We study the theoretical and computational aspects of the proposed framework and showcase our new construction on the classical problems of shape approximation and correspondence. We obtain significant improvement compared to classical Laplacian eigenbases as well as other alternatives for constructing localized bases.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Shape Analysis, 3D Shape Matching, Geometric Modeling

1. Introduction

Spectral methods are ubiquitously used in 3D shape analysis and geometry processing communities for a wide range of applications ranging from constructing local shape descriptors [SOG09], shape retrieval [JZ07, BBGO11] and correspondence between deformable shapes [OBBS*12, KBB*13] to mesh filtering [VL08, Tau95], remeshing [DBG*06] and compression [KG00]. The centerpiece of such methods is the construction of an orthogonal basis for

the space of functions defined on a manifold, allowing to generalize classical Fourier analysis to non-Euclidean domains. Typically, such bases are constructed by the diagonalization of the Laplace-Beltrami operator [Lév06]. The choice of Laplacian eigenbasis is convenient for several reasons. First, it is intrinsic and thus invariant to manifold parametrization and its isometric deformations [LZ10]. Second, it allows to be agnostic to a specific shape representation, as the Laplace-Beltrami operator can be discretized

on meshes, point clouds, volumes, etc. Third, Laplacian eigenbasis turns to be optimal for approximating functions with bounded variation [ABK15] and in many applications only the first few eigenfunctions are sufficient to achieve a good approximation. Finally, in the discrete setting, the computation of the Laplacian eigenbasis has relatively low complexity due to the sparse structure of the Laplacian matrix.

One of the key disadvantages of the Laplacian eigenbases is their *global* support. Thus, representing local structures requires using (potentially, infinitely) many basis functions. In many applications, one wishes to have a local basis that allows to limit the analysis to specific parts of the shape. The recently proposed *compressed manifold harmonics* [OLCO13, NVT*14, KGB16, BCKS16] attempt to construct local orthogonal bases that approximately diagonalize the Laplace-Beltrami operator. The main disadvantage of this framework is the inability to explicitly control the localization of the basis functions. Moreover, the basis is computed by solving an optimization problem on the Stiefel manifold of orthogonal matrices which does not guarantee a global solution.

Contribution. In this paper, we consider a new type of intrinsic operators whose spectral decomposition provides a local basis. Similarly to related constructions like [CSBK17], the new basis is smooth, local, and orthogonal; it is localized at specified regions of the shape, explicitly controllable; and it is efficiently computed by solving a standard eigendecomposition, thus coming with global optimality guarantees. The key novelty of our approach comes from its capability to integrate the *global* information obtained by the Laplacian eigenfunctions with *local* details given by our new basis. To this end, the localized basis is constructed in an incremental way, such that the new functions are orthogonal to some given set of functions (e.g., standard Laplacian eigenfunctions). Due to the aforementioned properties, we name our new basis *Localized Manifold Harmonics (LMH)*.

The rest of the paper is structured as follows. In Section 2 we review related works. Section 3 discusses the classical notions of spectral analysis on manifolds and the construction of manifold harmonic bases. In Section 4, we introduce our construction of localized manifold harmonics and study their properties in Section 5. In Section 6 we discuss the implementation details, and in Section 7 we show experimental results. We exemplify our construction in several applications including shape approximation and matching, where it shows superior performance compared to standard Laplacian eigenbases, as well as other alternatives for constructing localized bases. Finally, Section 8 concludes the paper discussing the advantages and limitations of our framework and possible future research directions. Proofs of theorems appear as Appendices in the supplementary materials.

2. Related work

The Laplace-Beltrami operator is so ubiquitous in computer graphics and geometry processing that it has earned the title of the “Swiss army knife” of geometry processing [SCV14]. The seminal work of Taubin realized the similarity between the Laplacian spectral decomposition and classical Fourier analysis [Tau95]. Karni and Gotsman used Laplacians for spectral mesh compression [KG00].

An influential paper of Levy [Lév06] showed how a wide range of applications can be addressed in the Laplacian spectral domain, igniting the interest in spectral approaches in the computer graphics field.

Many popular shape descriptors such as heat- [SOG09, GBAL09] and wave- [ASC11] kernel signatures, global point signatures [Rus07], and shape DNAs [RWP06] were constructed in the spectral domain. Coifman et al. introduced the notion of diffusion distances [CL06] on non-Euclidean domains, also constructed considering the spectral decomposition of heat operators. More recently, anisotropic versions of Laplacians and their heat kernels were considered [ARAC14, BMRB16]. Hildebrandt et al. considered Hessians of surface energies as an alternative family of operators [HSvTP12] incorporating extrinsic curvature information.

Ovsjanikov et al. [OBCS*12] introduced the functional maps framework to find correspondences between functions rather than points on the shapes, and used Laplacian eigenfunctions as natural choice for the basis to represent such maps (later, this choice was theoretically justified by Aflalo et al. [ABK15], who showed that Laplacian eigenbases are optimal for representing classes of functions with bounded variation). Kovnatsky et al. proposed the construction of compatible quasi-harmonic bases on collections of shapes using simultaneous diagonalization of Laplacians [KBB*13]. The approach was later extended by Litany et al. [LRBB17] to shapes having missing parts, giving rise to ad-hoc localized harmonic bases for correspondence problems.

A key drawback of Laplacian eigenbases is their global structure, which has adverse effects in numerous applications. In spectral shape deformation, it is hard to concentrate the analysis on local parts of the shape. In shape correspondence, the dependence on the Laplacian eigenfunctions on the global structure of the shape makes it hard to cope with topological noise and missing parts. Recently, to introduce local analysis on non-Euclidean manifolds, the Windowed Fourier Transform (WFT) has been proposed for graphs [SRV16] and shapes [BMM*15, MRCB16]. Although these methods improved the encoding of local parts, the authors still proposed Laplacian eigenfunctions as a basis to compute the spectral components, and therefore it was still hard to perform well in the challenging scenarios mentioned above.

As a possible remedy, Ozoliņš et al. [OLCO13] introduced *compressed modes*, a construction of local orthogonal bases that approximately diagonalize the Laplacian. The key idea of this method is the addition of a sparsity-promoting L_1 -norm to the Dirichlet energy (the combined effect of smoothness and sparsity results in localization of the basis functions). Rustamov [Rus11] previously used a similar regularization to construct local biharmonic kernels for function interpolation. Neumann et al. [NVT*14] applied the approach of [OLCO13] to problems in computer graphics. Kovnatsky et al. [KGB16] showed an efficient way of computing compressed manifold modes, while Bronstein et al. [BCKS16] proposed a more theoretically sound approach for the computation of L_1 -norm on manifolds.

Closely related to our method is the recent approach of Choukroun et al. [CSBK17], who considered the spectral decomposition of an elliptic operator realized as a diagonal update to the standard Laplacian. Differently from [CSBK17], our solutions are

simultaneously *localized* and *orthogonal* to the globally-supported Laplacian eigenbasis, leading to important practical consequences in several applications.

The key idea of this paper is the construction of localized bases by spectral decomposition of a modified Laplacian operator, crafted especially to provide eigenfunctions with local support. Our new operator inherits the important properties of the original Laplacian such as isometry invariance. In particular, it has a clear Fourier-like meaning that makes its use well interpretable. Differently from [NVT*14, KGB16, BCKS16, Rus11] which impose locality through an L_1 constraint, we allow an explicit indication of the local support of each function. This improves the versatility in controlling the local analysis, especially for semantically-guided interventions. Another important difference from other methods is that our new basis is computed by solving a standard eigendecomposition problem avoiding the need for more complex optimization methods.

3. Background

Manifolds. We model shapes as smooth two-dimensional manifolds \mathcal{X} (possibly with a boundary $\partial\mathcal{X}$) embedded into \mathbb{R}^3 . Locally around point x , the manifold is homeomorphic to the *tangent space* (or *plane*) $T_x\mathcal{X}$. The disjoint union of all the tangent spaces is the *tangent bundle* $T\mathcal{X}$. We further equip the manifold with a *Riemannian metric*, defined as an inner product $\langle \cdot, \cdot \rangle_{T_x\mathcal{X}} : T_x\mathcal{X} \times T_x\mathcal{X} \rightarrow \mathbb{R}$ on the tangent space depending smoothly on x . Properties expressed solely in terms of the metric are called *intrinsic*. In particular, *isometric* (metric-preserving) deformations of the embedded manifold preserve all intrinsic structures.

Let $f : \mathcal{X} \rightarrow \mathbb{R}$ and $F : \mathcal{X} \rightarrow T\mathcal{X}$ denote real *scalar* and *tangent vector fields* on the manifold, respectively. We can define inner products

$$\langle f, g \rangle_{L^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx; \quad (1)$$

$$\langle F, G \rangle_{L^2(T\mathcal{X})} = \int_{\mathcal{X}} \langle F(x), G(x) \rangle_{T_x\mathcal{X}} dx; \quad (2)$$

and denote them by $L^2(\mathcal{X}) = \{f : \mathcal{X} \rightarrow \mathbb{R} \text{ s.t. } \langle f, f \rangle_{L^2(\mathcal{X})} < \infty\}$ and $L^2(T\mathcal{X}) = \{F : \mathcal{X} \rightarrow T\mathcal{X} \text{ s.t. } \langle F, F \rangle_{L^2(T\mathcal{X})} < \infty\}$ the respective Hilbert spaces of square-integrable functions (here, dx is the area element induced by the metric).

Laplace-Beltrami operator. In classical calculus, the notion of derivative describes how the value of a function f changes with an infinitesimal change of its argument x . Due to the lack of vector space structure on the manifold (meaning that we cannot add two points, $x + dx$), we need to define the *differential* of f as an operator $df : T\mathcal{X} \rightarrow \mathbb{R}$ acting on tangent vector fields. At each point x , the differential is a linear functional $df(x) = \langle \nabla f(x), \cdot \rangle_{T_x\mathcal{X}}$ acting on tangent vectors $F(x) \in T_x\mathcal{X}$, which model a small displacement around x . The change of the function value as the result of this displacement is given by applying the differential to the tangent vector, $df(x)F(x) = \langle \nabla_{\mathcal{X}} f(x), F(x) \rangle_{T_x\mathcal{X}}$, and can be thought of as an extension of the notion of the classical directional derivative. The operator $\nabla_{\mathcal{X}} f : L^2(\mathcal{X}) \rightarrow L^2(T\mathcal{X})$ in the above definition is called the *intrinsic gradient*, and is similar to the classical notion of the gradient defining the direction of the steepest change of the function at a point.

The *intrinsic divergence* $\text{div}_{\mathcal{X}} : L^2(T\mathcal{X}) \rightarrow L^2(\mathcal{X})$ is defined as an operator adjoint to the intrinsic gradient, $\langle F, \nabla_{\mathcal{X}} f \rangle_{L^2(T\mathcal{X})} = \langle -\text{div}_{\mathcal{X}} f, F \rangle_{L^2(\mathcal{X})}$, where $f \in L^2(\mathcal{X})$ and $F \in L^2(T\mathcal{X})$ are some scalar and vector fields, respectively. The positive semi-definite *Laplace-Beltrami operator* (or *manifold Laplacian*) is defined as $\Delta_{\mathcal{X}} f = -\text{div}_{\mathcal{X}}(\nabla_{\mathcal{X}} f)$, generalizing the corresponding notion from Euclidean spaces to manifolds. The Laplacian is self-adjoint,

$$\langle \nabla_{\mathcal{X}} f, \nabla_{\mathcal{X}} g \rangle_{L^2(T\mathcal{X})} = \langle \Delta_{\mathcal{X}} f, g \rangle_{L^2(\mathcal{X})} = \langle f, \Delta_{\mathcal{X}} g \rangle_{L^2(\mathcal{X})}. \quad (3)$$

Geometrically, the Laplace-Beltrami operator can be interpreted as the (normalized) difference between the average of a function on an infinitesimal sphere around a point and the value of the function at the point itself.

Spectral analysis on manifolds. Due to self-adjointness, on a compact manifold \mathcal{X} with boundary $\partial\mathcal{X}$, the Laplace-Beltrami operator admits an orthonormal eigendecomposition [Cha84]

$$\Delta_{\mathcal{X}} \phi_i(x) = \lambda_i \phi_i(x) \quad x \in \text{int}(\mathcal{X}) \quad (4)$$

$$\langle \nabla_{\mathcal{X}} \phi_i(x), \hat{n}(x) \rangle = 0 \quad x \in \partial\mathcal{X}, \quad (5)$$

with Neumann boundary conditions (5), where \hat{n} is the normal vector to the boundary. Here, $0 = \lambda_1 \leq \lambda_2 \leq \dots$ is a countable set of non-negative real eigenvalues and ϕ_1, ϕ_2, \dots are the corresponding orthonormal eigenfunctions satisfying $\langle \phi_i, \phi_j \rangle_{L^2(\mathcal{X})} = \delta_{ij}$.

The Laplacian eigenfunctions form an orthonormal basis for $L^2(\mathcal{X})$ referred to as *manifold harmonics (MH)*. A function $f \in L^2(\mathcal{X})$ can therefore be expressed as the *Fourier series*

$$f(x) = \sum_{i \geq 1} \underbrace{\langle \phi_i, f \rangle_{L^2(\mathcal{X})}}_{\hat{f}_i} \phi_i(x), \quad (6)$$

where \hat{f}_i are the Fourier coefficients (or the forward Fourier transform), and the synthesis $\sum_{i \geq 1} \hat{f}_i \phi_i(x)$ is the inverse Fourier transform. The eigenvalues λ_i can be interpreted as frequencies in the classical harmonic analysis; thus, truncating the series (6) to the first k terms will result in a band-limited (with bandwidth λ_k) representation of f .

Given a scalar field $f \in L^2(\mathcal{X})$, the *Dirichlet energy*

$$\mathcal{E}_S(f) := \langle \nabla_{\mathcal{X}} f, \nabla_{\mathcal{X}} f \rangle_{L^2(T\mathcal{X})} = \langle f, \Delta f \rangle_{L^2(\mathcal{X})} \quad (7)$$

measures how ‘smooth’ the field is. It is possible to show that the Laplacian eigenbasis is the solution to the optimization problem

$$\min_{\Psi_1, \dots, \Psi_k} \sum_{i=1}^k \mathcal{E}_S(\Psi_i) \quad \text{s.t.} \quad \langle \Psi_i, \Psi_j \rangle_{L^2(\mathcal{X})} = \delta_{ij} \quad (8)$$

and thus can be considered as the smoothest possible orthonormal eigenbasis. Furthermore, the eigenvalues can be obtained as the values of the Dirichlet energy, $\mathcal{E}_S(\phi_i) = \lambda_i$.

4. Localized manifold harmonics

A notable drawback of classical spectral analysis on manifolds lies in its inherently ‘global’ nature. In fact, despite the Laplacian itself being a local (differential) operator, its eigenfunctions and eigenvalues carry geometric and topological information about the entire manifold [Cha84, RWP06, Rus07]. As a practical consequence,

operations that should be local by design are often affected by such global effects (see examples in Section 7).

In this paper, we introduce a new framework for spectral shape analysis that is designed to be at the same time *local* and *compatible* with the existing spectral constructions. In practice, our approach boils down to the computation of the eigenfunctions of a new operator, which is realized as a simple update to the classical manifold Laplacian – thus fully retaining the computational efficiency and theoretical guarantees of the resulting optimization process.

Definition. Let us be given a manifold \mathcal{X} , a region $R \subseteq \mathcal{X}$ thereof, a set of orthonormal functions $\phi_1, \dots, \phi_{k'}$ (e.g. the first k' Laplacian eigenfunctions), and an integer k . We seek a new set ψ_1, \dots, ψ_k of functions that are *smooth*, *orthonormal*, and *localized* on R , as the solution to the following optimization problem:

$$\min_{\psi_1, \dots, \psi_k} \sum_{j=1}^k \mathcal{E}_S(\psi_j) + \mu_R \mathcal{E}_R(\psi_j) \quad (9)$$

$$\text{s.t. } \langle \psi_i, \psi_j \rangle_{L^2(\mathcal{X})} = \delta_{ij} \quad i, j = 1, \dots, k \quad (10)$$

$$\langle \psi_i, \phi_j \rangle_{L^2(\mathcal{X})} = 0 \quad i = 1, \dots, k; j = 1, \dots, k' \quad (11)$$

where the constraints (11) demand the basis functions to be *orthogonal* to the subspace $\text{span}\{\phi_1, \dots, \phi_{k'}\}$. As we will see in what follows, it allows constructing an incremental set of functions that are orthogonal to a given set of standard Laplacian eigenfunctions.

The first term \mathcal{E}_S is the Dirichlet functional (7) promoting the *smoothness* of the new basis. The term

$$\mathcal{E}_R(f) := \int_{\mathcal{X}} (f(x)(1 - u(x)))^2 dx, \quad (12)$$

is a quadratic penalty promoting the localization of the basis functions on the given region $R \subseteq \mathcal{X}$. Here $u: \mathcal{X} \rightarrow [0, 1]$ is a membership function such that $u(x) = 1$ for $x \in R$ and $u(x) = 0$ otherwise. Note that we let function u assume a continuum of values in $[0, 1]$, implementing the notion of “soft” membership (the choice between binary and soft u is application-dependent).

We refer to the solutions of problem (9) as *localized manifold harmonics (LMH)*. Figure 2 provides an illustration of LMH in the $[0, 1]$ interval, while Figures 1 and 3 depict a few examples of such bases on 2D manifolds.




Figure 2: Classical (top row) and localized (bottom row) harmonics in 1D under Neumann boundary conditions. Note that the localized harmonics are orthogonal to those in the first row. The selected region $R \subset [0, 1]$ is marked as a black segment.

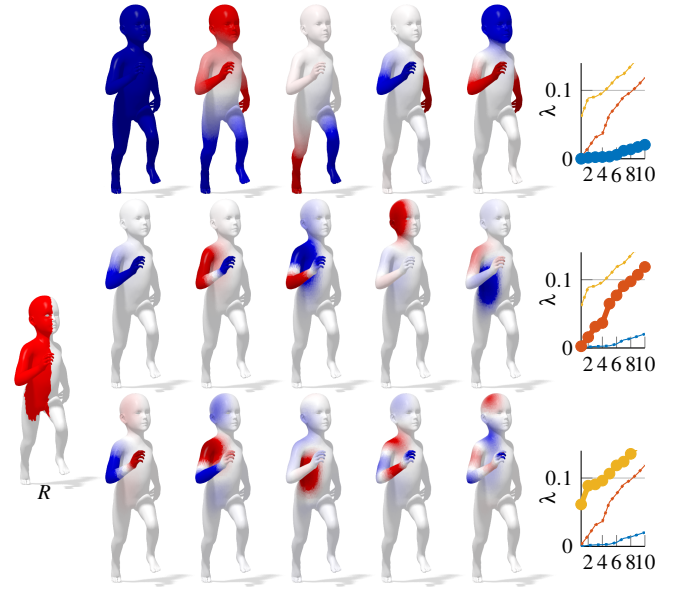


Figure 3: Localized manifold harmonics for the red region shown on the left. We show the first $k' = 5$ standard Laplacian eigenfunctions (top row), the first $k = 5$ LMH with the orthogonality term disabled ($\mu_{\perp} = 0$, middle row), and the first $k = 5$ LMH obtained by optimizing the full energy (14) (bottom row). Note that the latter harmonics are orthogonal to the first k' Laplacian eigenfunctions. We also show the generalized eigenvalues associated with each of the three cases.

Relaxed problem. For practical reasons, in what follows we will consider a relaxed variant of (9), in which we replace the hard constraints (11) by a large penalty:

$$\min_{\psi_1, \dots, \psi_k} \sum_{j=1}^k \mathcal{E}(\psi_j) \quad \text{s.t. } \langle \psi_i, \psi_j \rangle_{L^2(\mathcal{X})} = \delta_{ij}, \quad (13)$$

where

$$\mathcal{E}(\psi_j) = \mathcal{E}_S(\psi_j) + \mu_R \mathcal{E}_R(\psi_j) + \mu_{\perp} \mathcal{E}_{\perp}(\psi_j), \quad (14)$$

$$\mathcal{E}_{\perp}(f) := \sum_{i=1}^{k'} |\langle \phi_i, f \rangle_{L^2(\mathcal{X})}|^2. \quad (15)$$

Note that problems (13) and (9) are equivalent as $\mu_{\perp} \rightarrow \infty$. An empirical evaluation of the equivalence of the two formulations will be provided in Section 6.

Discretization. In the discrete setting, the manifold \mathcal{X} is sampled at n points x_1, \dots, x_n and is approximated by a triangular mesh (V, E, F) constructed upon these points, where $V = \{1, \dots, n\}$, $E = E_i \cup E_b$ and F are the vertices, edges, and faces of the mesh, respectively (E_i and E_b denote the interior and boundary edges, respectively). The discretization of the Laplace-Beltrami operator $\Delta_{\mathcal{X}}$ takes the form of an $n \times n$ sparse matrix $\mathbf{L} = -\mathbf{A}^{-1}\mathbf{W}$ according to the standard lumped linear FEM [Mac49]. The mass matrix \mathbf{A} is a diagonal matrix of area elements $a_i = \frac{1}{3} \sum_{jk:ijk \in F} A_{ijk}$, where A_{ijk} denotes the area of triangle ijk . The stiffness matrix \mathbf{W} contains

the cotangent weights

$$w_{ij} = \begin{cases} (\cot \alpha_{ij} + \cot \beta_{ij})/2 & ij \in E_i; \\ (\cot \alpha_{ij})/2 & ij \in E_b; \\ -\sum_{k \neq i} w_{ik} & i = j; \\ 0 & \text{else;} \end{cases} \quad (16)$$

where α_{ij}, β_{ij} denote the angles $\angle ikj, \angle jhi$ of the triangles sharing the edge ij . Scalar fields $f \in L^2(\mathcal{X})$ are represented as n -dimensional vectors $\mathbf{f} = (f(x_1), \dots, f(x_n))^T$; the inner products $\langle f, g \rangle_{L^2(\mathcal{X})}$ are discretized by area-weighted dot products $\mathbf{f}^T \mathbf{A} \mathbf{g}$.

We now turn to the discretization of problem (13). Let $\Psi \in \mathbb{R}^{n \times k}$ be a matrix containing our discretized basis functions ψ_1, \dots, ψ_k as its columns, and same way, let $\Phi \in \mathbb{R}^{n \times k'}$ be a matrix of the first k' Laplacian eigenfunctions $\phi_1, \dots, \phi_{k'}$. The total energy is discretized as $\sum_{j=1}^k \mathcal{E}(\psi_j) = \mathcal{E}(\Psi)$, comprising purely quadratic terms

$$\mathcal{E}_S(\Psi) = \text{tr}(\Psi^T \mathbf{W} \Psi) \quad (17)$$

$$\mathcal{E}_R(\Psi) = \text{tr}(\Psi^T \mathbf{A} \text{diag}(\mathbf{v}) \Psi) \quad (18)$$

$$\mathcal{E}_\perp(\Psi) = \text{tr}(\Psi^T \mathbf{A} \underbrace{\Phi \Phi^T}_{\mathbf{P}_{k'}} \mathbf{A} \Psi) \quad (19)$$

where \mathbf{v} denotes the discrete version of $v(x) \equiv (1 - u(x))^2$. In other words, the locality penalty (12) is implemented as a diagonal update to the standard Laplacian; while the term promoting orthogonality to Φ (15) is realized as a rank- k' projector $\mathbf{P}_{k'}$.

Due to the linearity of the trace, the discrete version of problem (13) can be expressed as

$$\min_{\Psi \in \mathbb{R}^{n \times k}} \text{tr}(\Psi^T \mathbf{Q}_{v,k'} \Psi) \quad \text{s.t.} \quad \Psi^T \mathbf{A} \Psi = \mathbf{I}, \quad (20)$$

where the matrix

$$\mathbf{Q}_{v,k'} = \mathbf{W} + \mu_R \mathbf{A} \text{diag}(\mathbf{v}) + \mu_\perp \mathbf{A} \mathbf{P}_{k'} \quad (21)$$

is symmetric and positive semi-definite (we make the dependency on v, k' explicit as a subscript). Problem (20) is equivalent to the *generalized eigenvalue problem*

$$\mathbf{Q}_{v,k'} \Psi = \mathbf{A} \Psi \Lambda, \quad (22)$$

(see Theorem 1.2 of [SW82]). We stress that the new operator $\mathbf{Q}_{v,k'}$ is intrinsic, and so are its eigenfunctions.

As shown later in Section 6, a *global* optimum of this problem can be found by classical Arnoldi-like methods. Note that global solutions to the original constrained problem (9) can also be easily computed (see Appendix A), however throughout this paper we favor the relaxed formulation for computational efficiency reasons. We refer to Section 6 for comparisons.

5. Properties of LMH

In this section, we discuss the main theoretical properties and computational aspects of our framework.

Basis functions. As mentioned before, our localized basis functions are orthonormal eigenfunctions of a matrix (22) obtained by modification of the Laplacian. Indeed, by setting $\mu_R = \mu_\perp = 0$, the solution of (20) is attained by the first k standard Laplacian

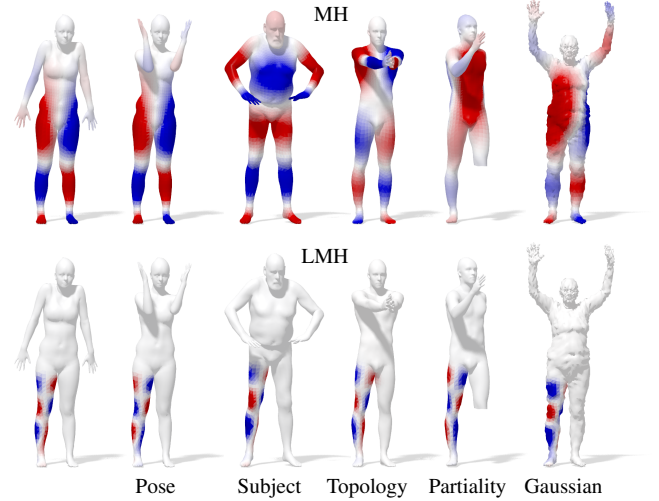


Figure 4: Laplacian eigenfunction ϕ_{12} (top) and localized manifold harmonic ψ_{12} (bottom) under different shape transformations. From left to right: near-isometry (different pose), non-isometric deformation (different subject), topological noise (glued hands), missing part, and geometric noise. LMH is more stable under such deformations compared to the standard MH.

eigenfunctions. Similarly, by setting $\mu_R = 0$ (no locality) and for $\mu_\perp \rightarrow \infty$, problem (20) can be equivalently rewritten as

$$\min_{\Psi \in \mathbb{R}^{n \times k}} \text{tr}(\Psi^T \mathbf{W} \Psi) \quad \text{s.t.} \quad (\Psi \Phi)^T \mathbf{A} (\Psi \Phi) = \mathbf{I}, \quad (23)$$

whose minimizers are the standard Laplacian eigenfunctions $\Psi_1 = \phi_{k'+1}, \dots, \Psi_k = \phi_{k'+k}$.

For $\mu_R > 0$ and $\mu_\perp > 0$, we obtain a new set of k functions Ψ_1, \dots, Ψ_k localized to a given region $R \subseteq \mathcal{X}$. These functions effectively *extend* the Laplacian eigenbasis, in the sense that the new set $\phi_1, \dots, \phi_{k'}, \Psi_1, \dots, \Psi_k$ forms an *orthonormal basis* for a $k + k'$ -dimensional subspace of $L^2(\mathcal{X})$. Importantly, the new basis is still isometry-invariant, and is designed to effectively represent functions with support restricted to the given region (see Figure 4). Compared to only using the ‘global’ Laplacian eigenbasis, the new representation provides a more parsimonious model: fewer localized harmonics are needed to capture the high-frequency content within R , than the number of harmonics that would be needed in the global basis. The localized nature of this construction allows to mitigate considerably the non-local effects associated with the adoption of the global basis (influence of topological noise, etc).

Finally, disabling the orthogonal penalty ($\mu_R > 0, \mu_\perp = 0$) would lead to the set of (now possibly linearly dependent) $k + k'$ functions spanning a $k'' \leq k + k'$ -dimensional subspace of $L^2(\mathcal{X})$. This may result in a redundant representation of functions supported on $R \subseteq \mathcal{X}$, e.g., whenever a standard Laplacian eigenfunction has also support in R . We refer to the experimental section for a deeper analysis of the effect of orthogonality on the representation quality.

Spectrum. By the interpretation of (20) as a generalized eigenvalue problem (22), we obtain a natural notion of *spectrum* associated with the LMH, namely given by $\lambda_j(\mathbf{Q}) = \Psi_j^T \mathbf{Q} \Psi_j, j =$

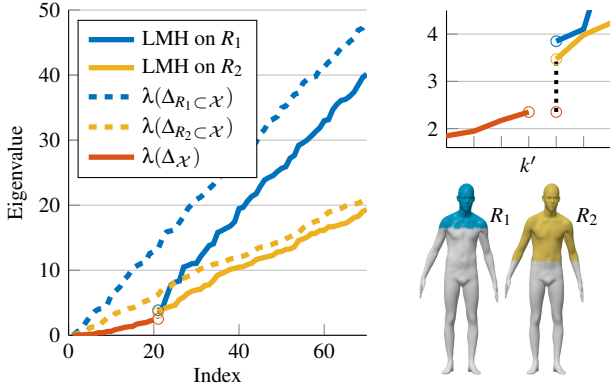


Figure 5: Left: LMH spectra grow linearly with rate inversely proportional to the area of the region, and are bounded from above by the standard Laplacian eigenvalues of submanifold $R \subseteq \mathcal{X}$. Top-right: Enlargement of the left plot around the index $k' = 20$. We illustrate the spectral gap (dotted) between $\lambda_{k'}(\mathbf{W})$ and $\lambda_1(\mathbf{Q}_{v,k'})$; note that the gap is different among the two regions.

$1, \dots, k$. Indeed, since at the optimum $\mathcal{E}_R(\boldsymbol{\Psi}_j) \approx \mathcal{E}_\perp(\boldsymbol{\Psi}_j) \approx 0$ for all j , we have $\boldsymbol{\Psi}_j^\top \mathbf{Q} \boldsymbol{\Psi}_j \approx \boldsymbol{\Psi}_j^\top \mathbf{W} \boldsymbol{\Psi}_j$, i.e., the Dirichlet energy of $\boldsymbol{\Psi}_j$ as in the classical setting (7). This provides us with a natural ordering of the basis functions; it remains to see in what measure do the localized harmonics bring additional (higher frequency) information to the global basis formed by the first k' Laplacian eigenfunctions $\phi_1, \dots, \phi_{k'}$. A first answer is provided by the following

Theorem 1 (spectral gap). Let $(\phi_i, \lambda_i(\mathbf{W}))_{i=1}^n$ be the eigenpairs of the standard Laplacian, and let $\mathbf{Q}_{v,k'}$ be defined as in (21). Then, for large enough μ_\perp , any non-negative v and any choice of $k' \leq n-1$, we have $\lambda_{k'}(\mathbf{W}) \leq \lambda_1(\mathbf{Q}_{v,k'})$, with equality holding iff $\phi_{k'+1}(x) = 0$ whenever $v(x) \neq 0$.

Proof. See supplementary material.

This theorem ensures the existence of a non-negative gap between the two spectra, i.e., the new basis functions do not introduce any redundancy, and the gap is the smallest possible by the global optimality of (20). In other words, the localized basis “picks up” where the global basis “left off” (see Figure 5 for examples). Note that the last condition on v in Theorem 1 is almost never realized in practice: equality is obtained only when the Laplacian eigenfunction $\phi_{k'+1}$ is localized to the same region indicated by v .

Interestingly, for a special class of functions v the spectrum $\lambda_1(\mathbf{Q}_{v,k'}) \leq \lambda_2(\mathbf{Q}_{v,k'}) \leq \dots$ follows a well-defined behavior, as remarked below.

Observation. Let v be a binary indicator function supported on some (possibly disconnected) region $R \subseteq \mathcal{X}$. Then, $\lambda_i(\mathbf{Q}_{v,k'}) - \lambda_1(\mathbf{Q}_{v,k'}) \propto i/\sqrt{\text{Area}(R)}$ as $i \rightarrow \infty$.

The observation above can be thought of as a generalization of Weyl’s asymptotic law [Cha84] to sub-regions of \mathcal{X} (see Figure 5).

Comparison to standard Laplacian on parts. Perhaps the most direct way to achieve locality is to consider the given region $R \subseteq \mathcal{X}$ as a separate manifold with boundary ∂R and Laplacian Δ_R , and




Figure 6: Our model allows to smoothly transition from a localized solution equivalent to a standard Laplacian eigenfunction ϕ_i^R on a partial shape with Neumann boundary conditions (first column), to a globally supported solution equivalent to a standard Laplacian eigenfunction ϕ_k on the full shape (last column). Each $\boldsymbol{\Psi}$ is obtained by solving a different problem with a different μ_R and $\mu_\perp = 0$. Note that the resulting “interpolating” harmonics $\boldsymbol{\Psi}_{j1\dots5}$ do not necessarily correspond to the same eigenvalue.

then compute the eigen-decomposition $\mathbf{W}^R \boldsymbol{\Psi}^R = \mathbf{A}^R \boldsymbol{\Psi}^R \boldsymbol{\Lambda}^R$ of Δ_R (note that $\mathbf{W}^R, \mathbf{A}^R$ can be obtained as submatrices of \mathbf{W}, \mathbf{A} followed by normalization and by fixing the weights along ∂R). The eigenfunctions $\boldsymbol{\Psi}_i^R$ can then be extended to the entire \mathcal{X} by means of zero-padding,

$$\tilde{\boldsymbol{\Psi}}_i^R(x) = \begin{cases} \boldsymbol{\Psi}_i^R(x) & x \in R \\ 0 & \text{else} \end{cases}$$

A first difference between this and our approach lies in the fact that $\langle \tilde{\boldsymbol{\Psi}}_i^R, \phi_j \rangle_{L^2(\mathcal{X})} \neq \delta_{ij}$ in general, i.e., the extended partial eigenfunctions do not “complete” the global basis and there is no separation of spectra (guaranteed in our case by Theorem 1), leading in turn to a redundant representation.

Secondly, our approach is more general in that we allow “soft” regions represented by allowing v to obtain values in the interval $[0, 1]$, which is obviously not achievable by extracting sub-regions. This latter property is especially important in applications where a sharp (binary) selection would lead to undesirable boundary effects around the region of interest (see Section 7 for examples).

Finally, we stress that the standard Laplacian $\Delta_{R \subseteq \mathcal{X}}$ may have an eigenspace in common with our operator with $\mu_\perp = 0$, a binary v on R , and large enough μ_R . In turn, the full Laplacian $\Delta_{\mathcal{X}}$ is always obtained for $\mu_R = 0$. A remarkable manifestation of this fact is given by the “interpolation effect” shown in Figure 6. Note that the observation above is not true in general, since we do not impose any boundary conditions w.r.t. R in our problem (indeed, we allow R to be soft), while all eigenfunctions of Δ_R always satisfy specific boundary conditions such as (5). Despite the loose connection, this observation allows us to complement the lower bound of Theorem 1 by the following

Theorem 2 (upper bound). Let \mathbf{W}^R be the stiffness matrix associated with the submanifold $R \subseteq \mathcal{X}$ and define v as the binary indicator function of $\mathcal{X} \setminus R$. Then, $\lambda_i(\mathbf{Q}_{v,k'}) \leq \lambda_{i+k'}(\mathbf{W}^R)$ for any $k \leq n$.

Proof. See supplementary material.

See Figure 5 for an example.

Comparison to compressed manifold modes. Ozoliš et

al. [OLCO13] proposed computing *compressed manifold modes* (CMM) as solutions to

$$\min_{\Psi \in \mathbb{R}^{n \times k}} \text{tr}(\Psi^\top \mathbf{W} \Psi) + \mu \|\Psi\|_1 \quad \text{s.t.} \quad \Psi^\top \mathbf{A} \Psi = \mathbf{I}. \quad (24)$$

Problem (24) makes use of a sparsity-inducing L_1 prior which, together with the smoothness promoted by the Dirichlet term, leads to the resulting functions having compact support controlled by parameter μ . It is important to note that this model does not allow to explicitly control the modes location. As shown in [NVT*14], these functions tend to concentrate around areas like shape protrusions and ridges. While different in its nature, the CMM model (24) admits a computational procedure which shares some similarities with ours. Assume that the solution of (24) for a given μ is a set of k functions supported on regions R_1, \dots, R_k (obtained *a posteriori*), represented by the soft indicators $1 - v_1, \dots, 1 - v_k$, and let $\mu_\perp = 0$. Then, the application of our framework using the matrices $\mathbf{Q}_{v_i, 0}$ corresponds to one iteration of the iterative reweighting scheme proposed for the efficient computation of CMMs in [BCKS16].

Comparison to elliptic operator. Concurrently with our work, Choukroun et al. [CSBK17] considered a family of elliptic operators of the form $\mathbf{H} = \mathbf{W} + \mathbf{V}$, where the *potential* \mathbf{V} is a diagonal operator akin to our localization term (18). The same approach was recently followed in [LJC17] to obtain localized basis functions around points of interest on the surface.

Differently from these approaches, we seek for localized basis functions that *simultaneously* lie in a subspace orthogonal to $\text{span}\{\phi_1, \dots, \phi_{k'}\}$, where ϕ_i are the first k' standard Laplacian eigenfunctions. In other words, we seek to “augment” the global basis by introducing a local refinement, while the aforementioned works attempt to construct a complete basis in agreement with the input potential. This is a crucial difference that has noticeable effects in practice, as we will demonstrate in Section 7.

6. Implementation

Optimization. As shown in Section 4, computing our localized basis functions boils down to solving a generalized eigenvalue problem $\mathbf{Q}\Psi = \mathbf{A}\Psi\Lambda$, with $\mathbf{Q} = \mathbf{W} + \mu_R \text{Adiag}(\mathbf{v}) + \mu_\perp \mathbf{A}\Phi\Phi^\top \mathbf{A}$. We note that computing \mathbf{Q} explicitly involves the construction of a dense $n \times n$ matrix $\mathbf{A}\Phi\Phi^\top \mathbf{A}$, which may become prohibitive for large meshes. However, we avoid this computation altogether by noticing that the μ_\perp -term has very low rank $k' \ll n$. This condition allows the application of exact update formulas throughout the optimization of problem (20), which can be solved efficiently and globally as detailed in Appendix B.

Timing. In Table 1 we report the runtime (in seconds) required by our method as executed on an Intel 3.6 GHz Core i7 cpu with 16GB ram. We compare the execution time for the exact problem with hard orthogonality constraints (9) and the relaxed problem (13). Note that while the latter relaxation is significantly more efficient, the two formulations yielded numerically close solutions in all our experiments (see Figure 7).

Choice of parameters. Parameters μ_R and k' control the locality and the number of global harmonics to use, respectively, and are application-dependent (see Section 7). Parameter μ_\perp enforces orthogonality w.r.t. the standard Laplacian eigenfunctions, and should

	k = 100		k = 200		k = 300	
	hard	soft	hard	soft	hard	soft
~ 120K	-	112.3s	-	200.5s	-	304.9s
~ 12K	40.8s	7.6s	66.7s	16.4s	79.3s	25.0s
~ 1.2K	1.1s	0.7s	2.2s	1.0s	4.6s	1.4s

Table 1: Runtime comparison for global optimization of our problem under hard (9) and soft constraints (13) across different mesh resolutions (number of faces) and basis size k . Tests denoted by ‘-’ could not run due to memory limitations (see Appendix A).

be chosen large enough so that the orthogonality constraints are satisfied. In our experiments we used $\mu_R \approx 10^2$ and $\mu_\perp \approx 10^5$; see Figure 7 for a quantitative evaluation on the choice of μ_\perp .




Figure 7: Left: Matrix of inner products $\langle b_i, b_j \rangle_{L^2(\mathcal{X})}$, where $b_i, b_j \in \{\phi_1, \dots, \phi_{k'}, \psi_1, \dots, \psi_k\}$. Here ψ_i are the optimal localized basis functions computed with $\mu_\perp = 10^{-1}$ (first column) and $\mu_\perp = 10^5$ (second column). Right: We plot the discrepancy between solutions to the exact (9) and relaxed (13) problems as a function of μ_\perp , measured as the L_2 distance between the resulting spectra.

7. Applications

Localized manifold harmonics are a general tool that can be employed as a drop-in replacement for, or in conjunction with the classical manifold harmonics ubiquitous in spectral shape analysis. In this Section we showcase their application in two broad tasks in graphics: spectral shape processing and shape correspondence.

Spectral shape processing. In this context, the surface \mathcal{X} is represented as a vector-valued function $\mathbf{x}: \mathcal{X} \rightarrow \mathbb{R}^3$, encoding the spatial coordinates of its embedding in \mathbb{R}^3 ; transformations to the surface geometry are then phrased as filtering operations applied to the coordinate functions. Vallet and Lévy [VL08] proposed to perform such filtering in the Fourier domain, where the coordinates \mathbf{x} are expressed as linear combinations of Laplacian eigenfunctions,

$$\mathbf{x} = \sum_{i \geq 1} \langle \phi_i, \mathbf{x} \rangle_{L^2(\mathcal{X})} \phi_i, \quad (25)$$

where, with some abuse of notation, $\langle \phi_i, \mathbf{x} \rangle_{L^2(\mathcal{X})} = (\langle \phi_i, x_1 \rangle_{L^2(\mathcal{X})}, \dots, \langle \phi_i, x_3 \rangle_{L^2(\mathcal{X})})$. By truncating the summation to the first k' terms, one obtains a band-limited representation of the surface. The representation is coarse for small k' , while finer details are captured for large values of k' ; see Figure 8 (top row) for an example. The expression in (25) provides an effective way for representing and manipulating simple shapes with smoothly varying coordinate functions, which can be compactly represented

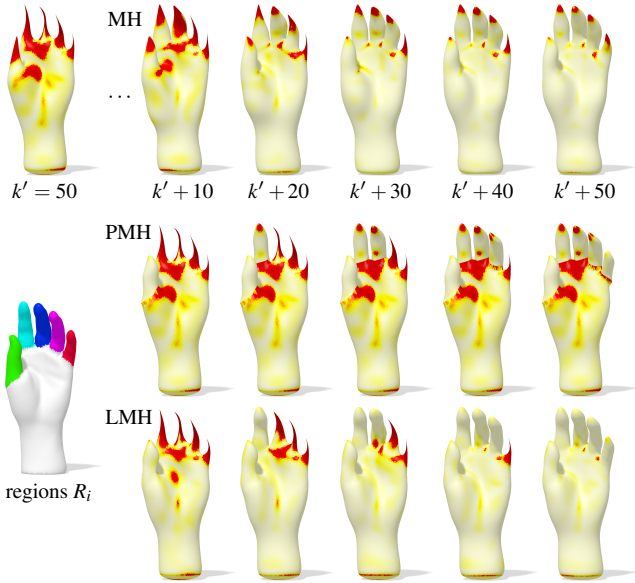


Figure 8: First row: Surface reconstruction via (25) using the first k' to $k' + 50$ Laplacian eigenfunctions. Second row: Each finger $R_i \subset \mathcal{X}$ is treated as a separate sub-manifold, and the eigenfunctions of the “partial” Laplacians Δ_{R_i} are used to update the initial reconstruction by adding 10 harmonics per finger. Third row: Reconstruction via (26) with 10 localized harmonics per finger. Note the significantly higher accuracy of LMH despite using the same number of harmonics as MH and PMH. The heatmap encodes reconstruction error, growing from white to dark red.

in the first few harmonics. Conversely, this representation is much less efficient for surfaces having details at smaller scales.

Assume a given set of regions, identifying areas of the shape with geometric detail. By computing localized harmonics $\{\psi_j\}_j$ on the given regions, we obtain a representation of the surface geometry:

$$\mathbf{x} \approx \sum_{i=1}^{k'} \langle \phi_i, \mathbf{x} \rangle_{L^2(\mathcal{X})} \phi_i + \sum_{j=1}^k \langle \psi_j, \mathbf{x} \rangle_{L^2(\mathcal{X})} \psi_j, \quad (26)$$

where $\phi_1, \dots, \phi_{k'}$ are the standard Laplacian eigenfunctions and $\langle \phi_i, \psi_j \rangle_{L^2(\mathcal{X})} \approx 0$ for all $i = 1, \dots, k'$ and $j = 1, \dots, k$. For a fixed number of terms in the series, the expression (26) yields a more accurate approximation of the original surface than (25), since the localized basis functions capture the high-frequency content more quickly (as also manifested in the rapid growth of the spectrum, see Figure 5). We refer to Figure 8 for a detailed illustration of this behavior. In Figure 9, we demonstrate the effect of the lack of orthogonality ($\mu_{\perp} = 0$) on the reconstruction quality, and in Figure 11 provide a quantitative evaluation on pre-segmented meshes from the Princeton segmentation benchmark [CGF09].

In these tests, we compare with standard manifold harmonics (MH) and “partial” manifold harmonics (PMH). The latter approach consists in reconstructing the surface indicated by each region $R_i \subset \mathcal{X}$ separately by using the eigenfunctions of the Laplacian Δ_{R_i} ; the reconstructed part is then “glued” back to the full shape. We measure the point-wise reconstruction error by the Eu-

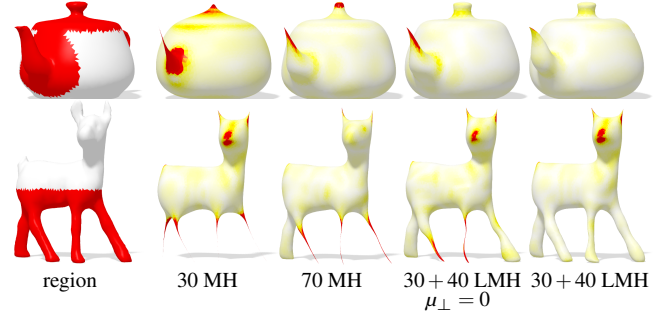


Figure 9: Comparison between manifold harmonics (MH) and localized manifold harmonics (LMH) without and with the orthogonality term (15) using $k' = 30$ and $k = 40$.




Figure 10: Comparison between different pipelines using a fixed number of basis functions (equal to 100 for all methods – LMH uses 50 global and 50 localized harmonics in both experiments). We report the reconstruction error below each method.

clidean distance between each reconstructed vertex and its corresponding point in the original surface.

Comparison to other pipelines. Differently from our method, both CMM [NVT*14] and elliptic operator (EO) [CSBK17] do not allow to build upon and enrich a given set of basis functions. In particular, CMM relies on the assumption that the set of localized basis functions covers the entire surface, while EO employs a soft potential with global support. In Figure 10 we compare the reconstruction error of standard MH, EO using the potential defined in [CSBK17], our method using the latter potential as a soft region v_{EO} , our method using the binary regions of Figure 8, and CMM using a covering set of compressed modes.

For completeness, we also test the performance of EO when fed with a sequence of binary potentials (the finger regions in Figure 8). We emphasize that EO is not designed to operate in conjunction with an existing basis, hence there is no natural way to implement an incremental update as the one shown in Figure 8: the harmonics computed on each binary region would not have an underlying global structure to attach to. For this reason, we provide an extra region (the palm) where 50 EO basis functions are computed, and this basis is incrementally updated with 10 EO basis functions per finger. The result is shown in Figure 12.

Shape correspondence. Ovsjanikov et al. [OBGS*12] proposed to represent correspondences between shapes by a linear operator (called *functional map*) $T : L^2(\mathcal{X}) \rightarrow L^2(\mathcal{Y})$ mapping functions on

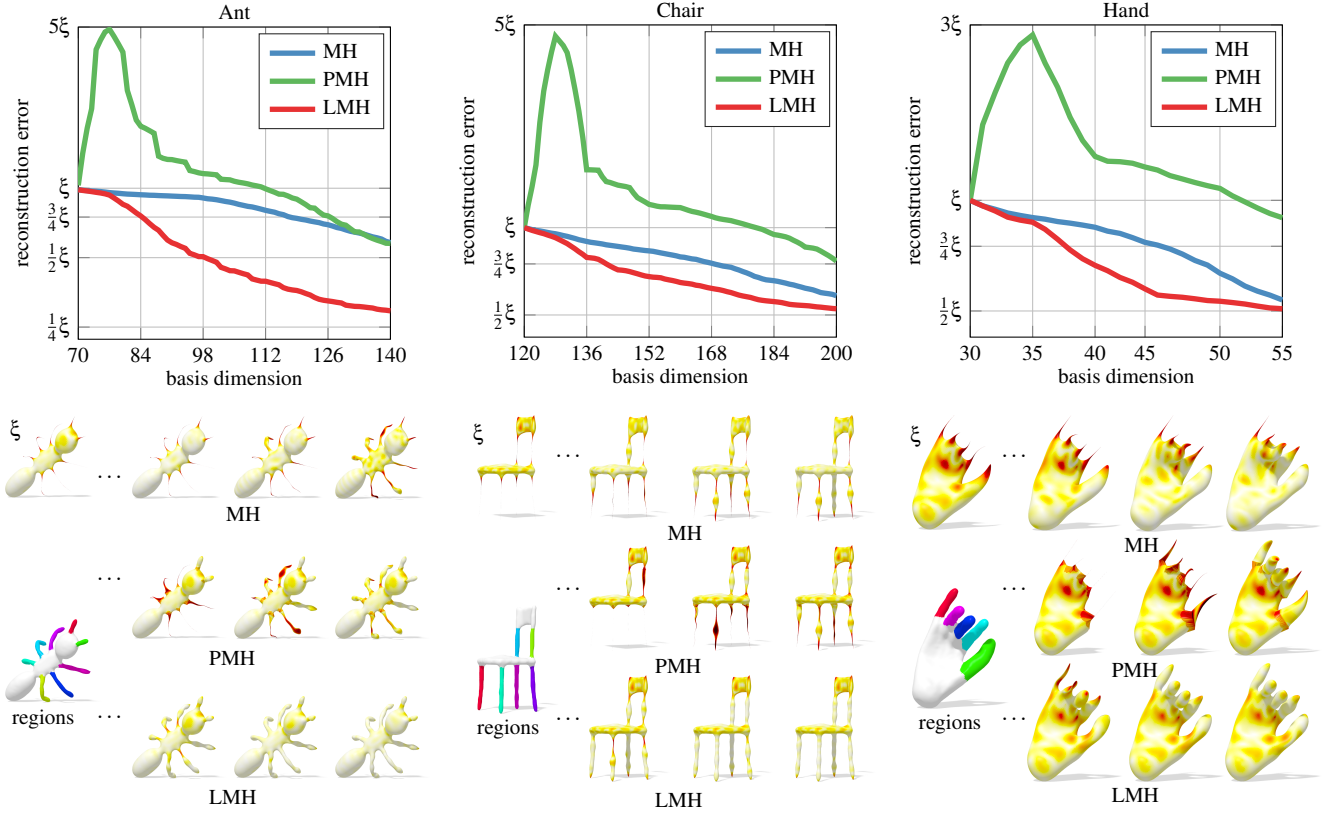


Figure 11: Comparisons among MH, PMH and LMH in surface representation. For each class we report the average reconstruction error at increasing number of basis functions (x-axis). Here, ξ denotes the error obtained by MH with $k' = 50$. The heatmaps encode reconstruction error, growing from white to dark red.




Figure 12: Incremental hand reconstruction using EO basis functions [CSBK17] and the binary regions of Figure 8.

\mathcal{X} to functions on \mathcal{Y} (one can see that point-to-point mappings constitute a special case in which delta-functions are mapped to delta-functions). As a linear operator, the functional map T admits a matrix representation $\mathbf{C} = (c_{ij})$ w.r.t. bases $\{\phi_i^{\mathcal{X}}\}$ and $\{\phi_j^{\mathcal{Y}}\}$ on $L^2(\mathcal{X})$ and $L^2(\mathcal{Y})$ respectively,

$$Tf = \sum_{ij \geq 1} \langle \phi_i^{\mathcal{X}}, f \rangle_{L^2(\mathcal{X})} \underbrace{\langle T\phi_i^{\mathcal{X}}, \phi_j^{\mathcal{Y}} \rangle_{L^2(\mathcal{Y})}}_{c_{ji}} \phi_j^{\mathcal{Y}}, \quad (27)$$

for an arbitrary $f \in L^2(\mathcal{X})$. By choosing the Laplacian eigenfunctions on \mathcal{X} and \mathcal{Y} as the bases $\{\phi_i^{\mathcal{X}}\}$ and $\{\phi_j^{\mathcal{Y}}\}$, one can truncate the

series (27) to the first k' terms – hence obtaining a compact representation which can be interpreted as a band-limited approximation of the full map. Correspondence problems can then be phrased as searching for a $k' \times k'$ matrix \mathbf{C} minimizing simple data fidelity criteria [OBCS*12, NO17] or exhibiting a particular structure depending on the correspondence setting [PBB*13, KBBV15, RCB*17].

Similar to the previous experiments, the standard Laplacian eigenbasis may not be the best choice in the presence of fine details: the low-pass nature of the spectral representation of the map, embodied in matrix \mathbf{C} , negatively affects the quality of the representation at a point-wise level. Indeed, recovering a point-to-point map from a functional map is considered a difficult problem in itself [RMC15, VLR*17], and is at the heart of several applications dealing with maps.

LMH can be directly employed for representing functional correspondence *in conjunction* with the Laplacian eigenbasis:

$$Tf = \sum_{\ell, m=1}^{k+k'} \langle \omega_{\ell}^{\mathcal{X}}, f \rangle_{L^2(\mathcal{X})} \langle T\omega_{\ell}^{\mathcal{X}}, \omega_m^{\mathcal{Y}} \rangle_{L^2(\mathcal{Y})} \omega_m^{\mathcal{Y}}, \quad (28)$$

where $\omega \in \Omega$ and $\Omega = \{\phi_i\}_{i=1}^{k'} \cup \{\psi_j\}_{j=1}^k$ is the union of the standard and localized manifold harmonics. Note that the formula above allows for ‘cross-talk’ between the MH and LMH bases, and




Figure 13: Functional map matrices w.r.t. the standard Laplacian eigenbasis (left) and w.r.t. a “mixed” basis composed of k' Laplacian eigenfunctions and $k+k$ localized harmonics (middle). The maps encode the ground-truth correspondence between the two shapes shown on the right; the regions used for the computation of LMH are highlighted in red and blue. Note the block-diagonal structure of the second matrix, a manifestation of the capability of LMH to encode local information compactly.




Figure 14: Mean geodesic error vs. number of basis functions used in the functional map representation. Note that LMH lead to an increase in accuracy, resulting in turn in a more compact representation of the correspondence. Here ξ is defined as the mean geodesic error of MH at $k' = 50$.

it can be seen as a localized *refinement* to some initial correspondence represented in the (global) Laplacian eigenbasis. An example of the resulting correspondence matrix \mathbf{C} is shown in Figure 13.

In Figure 14 we show a quantitative comparison between the two representations (27) and (28). For this experiment we use near-isometric shapes from the TOSCA dataset [BBK08]. For each pair of shapes, we use their ground-truth point-to-point correspondence to construct functional maps of increasing size in the Laplacian eigenbasis and in the LMH basis. For the latter, we use Eq. (28) with $k' = 50$ and k increasing from 1 to 50. The harmonics $\{\Psi_\ell^\mathcal{X}, \Psi_m^\mathcal{Y}\}_{\ell,m=1}^k$ are localized to the regions having large reconstruction error, computed as in the previous experiments. Note that even though these regions can be arbitrarily disconnected and irregular (see, e.g., Figure 15), our framework can be applied without modifications. A point-to-point map is recovered from each functional map using the nearest neighbor approach [OBBCS*12]. We measure the correspondence quality via its *geodesic error*. Assume a point-to-point match $(x,y) \in \mathcal{X} \times \mathcal{Y}$ is recovered, whereas the

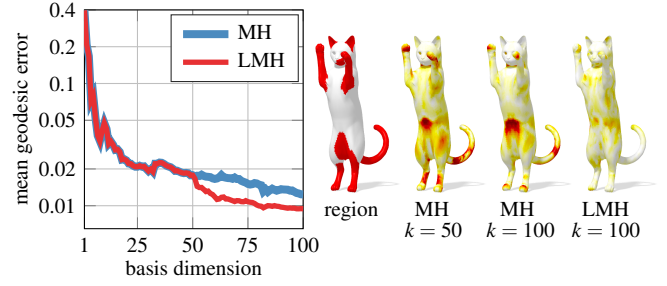


Figure 15: Accuracy improvement in the functional map representation, obtained by introducing LMH after $k' = 50$ Laplacian eigenfunctions. On the right we show the regions used for LMH and the geodesic error (encoded as hot colors, growing from white to dark red) obtained for different configurations. Note the higher accuracy attained by LMH for the same amount of basis functions.

ground-truth correspondence is (x,y^*) ; we compute the quantity

$$\varepsilon(x) = \frac{d_{\mathcal{Y}}(y,y^*)}{\sqrt{\text{Area}(\mathcal{Y})}}, \quad (29)$$

where $d_{\mathcal{Y}}$ is the geodesic distance on \mathcal{Y} .

For the second set of experiments we consider a challenging setting of shape correspondence known as *deformable object-in-clutter* [CRM*16]. In this scenario, the task is to match a given model to a scene where the model appears in a different pose, and in the simultaneous presence of clutter (extra objects) and missing parts. The problem was recently tackled in [CRM*16] using the functional map representation; to our knowledge, this method represents the current state of the art for this class of problems.

As data for these tests we use the entire dataset adopted for the comparisons in [CRM*16]. The dataset consists of 3 TOSCA models (*cat, centaur, dog*) and 150 synthetic scenes in which the models appear. Sparse point-to-point matches (around 10) between models and scenes, obtained using the approach of [RBA*12], are also provided. Given m input matches, we construct a mixture of m Gaussians with equal variance (set to 1% of the shape diameter) to define a *soft* region u on both model and scene (see Figure 16). We then construct a functional map \mathbf{C} upon the input sparse correspondence, and represent it w.r.t. $k = 15$ localized manifold harmonics computed on the soft regions (note that here we do *not* use any global eigenfunction, i.e., we set $k' = 0$ in (28)). Finally, we recover from




Figure 16: Localized manifold harmonics on a soft region encoded by function u . No thresholding is required in order to obtain a valid set of functions minimizing the energy (14).




Figure 17: Comparisons with the state of the art in deformable object-in-clutter. All methods use the same input data.

C a dense point-to-point map localized on u using the intrinsic ICP approach of [OBCS*12].

The results of this experiment are reported in Figure 17 quantitatively and in Figure 19 qualitatively. Despite the simple approach, our method gains a significant improvement in accuracy, of up to 25% upon the state of the art on this benchmark, highlighting the inherent robustness of LMH to missing parts and topological artifacts. Finally, in Figure 18 we compare (on a single pair of shapes) our pipeline with the counterparts obtained by replacing LMH with MH and PMH.

8. Discussion and conclusions

We introduced a new framework for spectral shape analysis and processing, allowing to perform operations which are localized to a given (possibly soft or disconnected) region of interest on the surface. Our framework is flexible, in that it can naturally enrich or fully replace the standard manifold harmonics in several tasks in graphics. We demonstrated its applicability in applications of geometry processing and shape correspondence, demonstrating a significant advantage if compared with the standard ‘global’ constructions based on the eigendecomposition of the Laplace operator.

Limitations. Perhaps the biggest limitation of our approach lies in the availability of regions (or soft counterparts thereof) upon which to carry out the localized spectral analysis. Such information may not be available in certain unsupervised applications, where it is often difficult to define a meaningful segmentation – indeed, an inherently task-specific notion suggesting the use of data-driven approaches. Further, similarly to the classical setting, it is not obvious how to choose the number of harmonics to employ for a given task.

Future directions. Despite the theoretical and empirical results provided within this paper, we feel that our study is still just ‘scratching the surface’ of a much broader area of research, with potentially extensive applications in geometry processing and graphics. Even though quite elusive at this stage, we foresee connections with results in *localization theory* (see, e.g., [FMP12]),

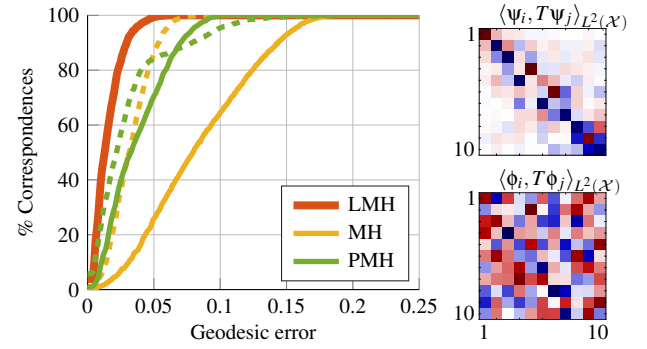


Figure 18: Left: Correspondence accuracy in the object-in-clutter setting. Each curve corresponds to a functional map expressed in a different basis, using the same input data. We show the performance when using $k = 10$ (solid curves) and $k = 50$ (dashed curves) basis functions. In this example, MH could not reach the quality of LMH for any choice of k . Right: Functional map matrices in the LMH basis $\{\psi\}$ (top) and in the MH basis $\{\phi\}$ (bottom).

leading to a promising direction of research which we believe deserves a deeper exploration.

Acknowledgments. The authors wish to thank Dorian Nogneng, Zorah Löhner, Haggai Maron and Nadav Dym for the technical support, and Klaus Glashoff, Or Litany, Christopher Brandt and Maks Ovsjanikov for useful discussions. Research was partially completed while ER was visiting the Institute for Mathematical Sciences, National University of Singapore in 2017. ER and MB are supported by ERC StG grant no. 307048 (COMET).

Appendix A. We show how to compute a global solution to problem (9). We start by observing that the hard constraints (11) require the desired basis functions Ψ to lie in the null space of the linear map $\mathbf{P}_{k'} := \Phi\Phi^\top \mathbf{A}$ (i.e., the projector onto $\text{Im}(\Phi)$), or equivalently to lie in the range of $\mathbf{I} - \mathbf{P}_{k'}$ (i.e. the projector onto the orthogonal subspace). This is easily achieved by letting $\Psi = (\mathbf{I} - \mathbf{P}_{k'})\mathbf{Y}$, and solving the generalized eigenvalue problem:

$$\tilde{\mathbf{Q}}\mathbf{Y} = \tilde{\mathbf{A}}\mathbf{Y}\Lambda, \quad (30)$$

where $\tilde{\mathbf{Q}} = (\mathbf{I} - \mathbf{P}_{k'})^\top (\mathbf{W} + \mu_R \mathbf{A} \text{diag}(\mathbf{v})) (\mathbf{I} - \mathbf{P}_{k'})$ and $\tilde{\mathbf{A}} = (\mathbf{I} - \mathbf{P}_{k'})^\top \mathbf{A} (\mathbf{I} - \mathbf{P}_{k'}) = \mathbf{A} (\mathbf{I} - \mathbf{P}_{k'})$. A similar trick was recently used in [DML17] for computing maximum magnitude eigenvalues of a large matrix. Note that solving problem (30) involves the explicit construction of a dense $n \times n$ matrix $\mathbf{P}_{k'}$, becoming prohibitive for large meshes.

Appendix B. We show how to *efficiently* compute a global solution to the generalized eigenvalue problem:

$$\mathbf{Q}\Psi = \mathbf{A}\Psi\Lambda, \quad (31)$$

with $\mathbf{Q} = \mathbf{W} + \mu_R \mathbf{A} \text{diag}(\mathbf{v}) + \mu_\perp (\mathbf{A}\Phi)(\mathbf{A}\Phi)^\top$.

Since the operator \mathbf{Q} is real and symmetric w.r.t. the positive semi-definite mass matrix \mathbf{A} , we employ the (globally optimal) implicitly restarted Arnoldi method (IRAM) [LS96] (as implemented in the ARPACK suite [LSY97]) for computing its first k eigenpairs.




Figure 19: Qualitative comparisons between our LMH-based approach for deformable shape correspondence in clutter and the state of the art [CRM*16]. For each experiment we show the dense correspondence (corresponding points have same color) and the geodesic error (hot colors growing from white to dark red).

The application of IRAM involves iteratively solving linear systems of the form:

$$\mathbf{Q}\mathbf{x}^{(t)} = \mathbf{A}\mathbf{b}^{(t)} \quad (32)$$

for some given $\mathbf{b}^{(t)} \in \mathbb{R}^n$. Expressing \mathbf{Q} in terms of the matrices $\mathbf{Z} = \mathbf{W} + \mu_{\perp} \mathbf{A} \text{diag}(\mathbf{v})$ and $\mathbf{B} = \mathbf{A}\mathbf{\Phi} \in \mathbb{R}^{n \times k'}$, we come to:

$$(\mathbf{Z} + \mu_{\perp} \mathbf{B}\mathbf{B}^{\top})\mathbf{x}^{(t)} = \mathbf{A}\mathbf{b}^{(t)}. \quad (33)$$

Note that matrix $\mathbf{B}\mathbf{B}^{\top}$ can be interpreted as a rank- k' update to \mathbf{Z} (with $k' \ll n$), allowing us to apply the Sherman-Morrison-Woodbury identity [Woo50]:

$$(\mathbf{Z} + \mu_{\perp} \mathbf{B}\mathbf{B}^{\top})^{-1} = \mathbf{Z}^{-1} - \underbrace{\mu_{\perp} \mathbf{Z}^{-1} \mathbf{B} (\mathbf{I} + \mu_{\perp} \mathbf{B}^{\top} \mathbf{Z}^{-1} \mathbf{B})^{-1} \mathbf{B}^{\top}}_{\mathbf{Y}} \mathbf{Z}^{-1}.$$

It is important to notice that the rhs does not involve the computation of $\mathbf{B}\mathbf{B}^{\top}$, and only involves efficient operations with a sparse $n \times n$ matrix \mathbf{Z} and a dense $k' \times k'$ matrix \mathbf{Y} . The application of this formula for the solution of problem (33), and in turn (31) via IRAM, is illustrated for clarity in Algorithm 1. A similar procedure was followed in [BCKS16] for the computation of CMM.

Algorithm 1 Efficient solution of problem (33).

Solve sparse linear system $\mathbf{Z}\boldsymbol{\xi} = \mathbf{A}\mathbf{b}^{(t)}$ for $\boldsymbol{\xi} \in \mathbb{R}^n$;
 Solve sparse linear system $\mathbf{Z}\boldsymbol{\Gamma} = \mu_{\perp} \mathbf{B}$ for $\boldsymbol{\Gamma} \in \mathbb{R}^{n \times k'}$;
 Solve dense linear system $(\mathbf{I}_{k'} + \mathbf{B}^{\top} \boldsymbol{\Gamma})\boldsymbol{\eta} = \mathbf{B}^{\top} \boldsymbol{\xi}$ for $\boldsymbol{\eta} \in \mathbb{R}^{k'}$;
 Compute final solution $\mathbf{x}^{(t)} = \boldsymbol{\xi} - \boldsymbol{\Gamma}\boldsymbol{\eta}$.

References

[ABK15] AFLALO Y., BREZIS H., KIMMEL R.: On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences* 8, 2 (2015), 1141–1160. 2

[ARAC14] ANDREUX M., RODOLÀ E., AUBRY M., CREMERS D.: Anisotropic Laplace-Beltrami operators for shape analysis. In *Proc. ECCV Workshops* (Cham, Switzerland, 2014), Springer International Publishing, pp. 299–312. 2

[ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV* (Barcelona, Spain, 2011), IEEE, pp. 1626–1633. 2

[BBG011] BRONSTEIN A. M., BRONSTEIN M. M., GUIBAS L. J., OVSJANIKOV M.: ShapeGoogle: Geometric words and expressions for invariant shape retrieval. *ACM Trans. Graphics* 30, 1 (2011), 1–20. 1

[BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, New York, NY, 2008. 10

[BCKS16] BRONSTEIN A., CHOUKROUN Y., KIMMEL R., SELA M.: Consistent discretization and minimization of the l_1 norm on manifolds. In *Proc. 3DV* (Stanford, CA, 2016), IEEE, pp. 435–440. 2, 3, 7, 12

[BMM*15] BOSCAINI D., MASCI J., MELZI S., BRONSTEIN M. M., CASTELLANI U., VANDERGHEYNST P.: Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *Computer Graphics Forum* 34, 5 (2015), 13–23. 2

[BMRB16] BOSCAINI D., MASCI J., RODOLÀ E., BRONSTEIN M. M.: Learning shape correspondence with anisotropic convolutional neural networks. In *Proc. NIPS* (Barcelona, Spain, 2016). 2

[CGF09] CHEN X., GOLOVINSKIY A., FUNKHOUSER T.: A benchmark for 3D mesh segmentation. *ACM Trans. Graphics* 28, 3 (2009), 73:1–73:12. 8

[Cha84] CHAVEL I.: *Eigenvalues in Riemannian geometry*, vol. 115. Academic press, Orlando, FL, 1984. 3, 6

[CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *Applied and Computational Harmonic Analysis* 21, 1 (2006), 5–30. 2

[CRM*16] COSMO L., RODOLÀ E., MASCI J., TORSELLO A., BRONSTEIN M. M.: Matching deformable objects in clutter. In *Proc. 3DV* (Stanford, CA, 2016), IEEE, pp. 1–10. 10, 11, 12

[CSBK17] CHOUKROUN Y., SHTERN A., BRONSTEIN A., KIMMEL R.: Hamiltonian operator for spectral shape analysis. *arXiv:1611.01990* (2017). 2, 7, 8, 9

[DBG*06] DONG S., BREMER P.-T., GARLAND M., PASCUCCI V.,

- HART J. C.: Spectral surface quadrangulation. *ACM Trans. Graphics* 25, 3 (2006), 1057–1066. 1
- [DML17] DYM N., MARON H., LIPMAN Y.: DS++: A flexible, scalable and provably tight relaxation for matching problems. *arXiv:1705.06148* (2017). 11
- [FMP12] FILOCHE M., MAYBORODA S., PATTERSON B.: Localization of eigenfunctions of a one-dimensional elliptic operator. *Contemporary Mathematics* 581 (2012), 99–116. 11
- [GBAL09] GEBAL K., BÆRENTZEN J. A., ANÆS H., LARSEN R.: Shape analysis using the auto diffusion function. *Computer Graphics Forum* 28, 5 (2009), 1405–1413. 2
- [HSvTP12] HILDEBRANDT K., SCHULZ C., VON TYCOWICZ C., POLTHIER K.: Modal shape analysis beyond Laplacian. *Computer Aided Geometric Design* 29, 5 (2012), 204–218. 2
- [JZ07] JAIN V., ZHANG H.: A spectral approach to shape-based retrieval of articulated 3D models. *Computer-Aided Design* 39, 5 (2007), 398–407. 1
- [KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Computer Graphics Forum* 32, 2pt4 (2013), 439–448. 1, 2
- [KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDEREGHEYNST P.: Functional correspondence by matrix completion. In *Proc. CVPR* (Boston, MA, 2015), IEEE, pp. 905–914. 9
- [KG00] KARNI Z., GOTSMAN C.: Spectral compression of mesh geometry. In *Proc. CGIT* (New York, NY, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 279–286. 1, 2
- [KGB16] KOVNATSKY A., GLASHOFF K., BRONSTEIN M. M.: MADMM: a generic algorithm for non-smooth optimization on manifolds. In *Proc. ECCV* (Cham, Switzerland, 2016), Springer International Publishing, pp. 680–696. 2, 3
- [Lév06] LÉVY B.: Laplace-Beltrami eigenfunctions towards an algorithm that understands geometry. In *Proc. SMI* (Washington, DC, 2006), IEEE, pp. 13–25. 1, 2
- [LJC17] LIU D., JACOBSON A., CRANE K.: A Dirac operator for extrinsic shape analysis. *Computer Graphics Forum* 36, 5 (2017). 7
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Fully spectral partial shape matching. *Computer Graphics Forum* 36, 2 (2017), 247–258. 2
- [LS96] LEHOUCQ R. B., SORENSEN D. C.: Deflation techniques for an implicitly restarted Arnoldi iteration. *SIAM Journal on Matrix Analysis and Applications* 17, 4 (1996), 789–821. 11
- [LSY97] LEHOUCQ R., SORENSEN D., YANG C.: Arpack users' guide: Solution of large scale eigenvalue problems with implicitly restarted Arnoldi methods, 1997. 11
- [LZ10] LÉVY B., ZHANG H. R.: Spectral mesh processing. In *SIG-GRAPH 2010 Courses* (New York, NY, 2010), ACM, pp. 8:1–8:312. 1
- [Mac49] MACNEAL R. H.: *The solution of partial differential equations by means of electrical networks*. PhD thesis, California Institute of Technology, 1949. 4
- [MRCB16] MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M. M.: Shape analysis with anisotropic windowed Fourier transform. In *Proc. 3DV* (Stanford, CA, 2016), IEEE, pp. 470–478. 2
- [MS10] MYRONENKO A., SONG X.: Point set registration: Coherent point drift. *TPAMI* 32, 12 (Dec 2010), 2262–2275. 11
- [NO17] NOGNENG D., OVSJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum* 36, 2 (2017), 259–267. 9
- [NVT*14] NEUMANN T., VARANASI K., THEOBALT C., MAGNOR M., WACKER M.: Compressed manifold modes for mesh processing. *Computer Graphics Forum* 33, 5 (2014), 35–44. 2, 3, 7, 8
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graphics* 31, 4 (2012), 30:1–30:11. 1, 2, 8, 9, 10, 11
- [OLCO13] OZOLIŅŠ V., LAI R., CAFLISCH R., OSHER S.: Compressed modes for variational problems in mathematics and physics. *Proceedings of the National Academy of Sciences* 110, 46 (2013), 18368–18373. 2, 7
- [PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum* 32, 2pt4 (2013), 459–468. 9
- [RABT13] RODOLÀ E., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A scale independent selection process for 3D object recognition in cluttered scenes. *International Journal of Computer Vision* 102, 1–3 (2013), 129–145. 11
- [RBA*12] RODOLÀ E., BRONSTEIN A., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A game-theoretic approach to deformable shape matching. In *Proc. CVPR* (Washington, DC, June 2012), IEEE, pp. 182–189. 10
- [RCB*17] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum* 36, 1 (2017), 222–236. 9
- [RMC15] RODOLÀ E., MOELLER M., CREMERS D.: Point-wise map recovery and refinement from functional correspondence. In *Proc. VMV* (Aachen, Germany, 2015), Eurographics Association, pp. 25–32. 9
- [Rus07] RUSTAMOV R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP* (Aire-la-Ville, Switzerland, 2007), Eurographics Association, pp. 225–233. 2, 3
- [Rus11] RUSTAMOV R. M.: Multiscale biharmonic kernels. *Computer Graphics Forum* 30, 5 (2011), 1521–1531. 2, 3
- [RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-Beltrami spectra as 'Shape-DNA' of surfaces and solids. *Computer-Aided Design* 38, 4 (2006), 342–366. 2, 3
- [SCV14] SOLOMON J., CRANE K., VOUGA E.: Laplace-Beltrami: The Swiss army knife of geometry processing. In *Symposium on Geometry Processing graduate school* (Cardiff, UK, 2014). 2
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum* 28, 5 (2009), 1383–1392. 1, 2
- [SRV16] SHUMAN D. I., RICAUD B., VANDEREGHEYNST P.: Vertex-frequency analysis on graphs. *Applied and Computational Harmonic Analysis* 40, 2 (2016), 260–291. 2
- [SW82] SAMEH A. H., WISNIEWSKI J. A.: A trace minimization algorithm for the generalized eigenvalue problem. *SIAM Journal on Numerical Analysis* 19, 6 (1982), 1243–1259. 5
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *Proc. CGIT* (New York, NY, 1995), ACM, pp. 351–358. 1, 2
- [VL08] VALLET B., LÉVY B.: Spectral geometry processing with manifold harmonics. *Computer Graphics Forum* 27, 2 (2008), 251–260. 1, 7
- [VLR*17] VESTNER M., LITMAN R., RODOLÀ E., BRONSTEIN A. M., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proc. CVPR* (Honolulu, HI, 2017), IEEE. 9
- [Woo50] WOODBURY M. A.: *Inverting modified matrices*. No. 42 in Statistical Research Group Memorandum Reports. Princeton University, Princeton, NJ, 1950. 12