

Coupled quasi-harmonic bases

A. Kovnatsky¹, M. M. Bronstein¹, A. M. Bronstein², K. Glashoff^{1,3} and R. Kimmel⁴

¹Institute of Computational Science, Faculty of Informatics, University of Lugano, Switzerland

²School of Engineering, Tel Aviv University, Israel

³Department of Mathematics, University of Hamburg, Germany

⁴Department of Computer Science, Technion, Israel

Abstract

The use of Laplacian eigenbases has been shown to be fruitful in many computer graphics applications. Today, state-of-the-art approaches to shape analysis, synthesis, and correspondence rely on these natural harmonic bases that allow using classical tools from harmonic analysis on manifolds. However, many applications involving multiple shapes are obstructed by the fact that Laplacian eigenbases computed independently on different shapes are often incompatible with each other. In this paper, we propose the construction of common approximate eigenbases for multiple shapes using approximate joint diagonalization algorithms, taking as input a set of corresponding functions (e.g. indicator functions of stable regions) on the two shapes. We illustrate the benefits of the proposed approach on tasks from shape editing, pose transfer, correspondence, and similarity.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

1. Introduction

It is well-established that the eigenfunctions of the Laplace-Beltrami operator (*manifold harmonics*) of a 3D shape modeled as a 2-manifold play the role of the Fourier basis in the Euclidean space [Tau95, Lé06]. Methods based on the Laplace-Beltrami operator have been used in a wide range of applications, including remeshing [Kob97, NISA06], parametrization [FH05], compression [KG00], recognition [RWP05, Rus07], and clustering. Many methods in computer graphics and geometry processing draw inspiration from the world of physics, finding analogies between physical processes such as heat diffusion [CL06] or wave propagation [ASC11] and the geometric properties of the shape [SOG09, BK10]. Several papers have studied consistent discretizations of the Laplace-Beltrami operator [PP93, MDSB03, WMKG08].

Taubin [Tau95] drew the analogy between the classical signal processing theory and manifold harmonics, showing that standard tools in signal processing such as analysis and synthesis of signals can be carried out on manifolds. This idea was extended in [KR05] and later in [Lé06, VL08, LZ09], who showed a practical framework for shape filtering and editing using the manifold harmon-

ics transform. In [OBCS*12], the authors proposed a novel representation of correspondences between shapes as linear maps between functional spaces on manifolds. In this representation, the Laplace-Beltrami eigenbases of the shapes play a crucial role, as they allow to parametrize the linear map as a matrix mapping the Fourier coefficients from one shape to another.

Applications involving multiple shapes rely on the fact that the harmonic bases computed on each shape independently are compatible with each other. However, this assumption is often unrealistic. First, eigenfunctions are only defined up to sign flips for shapes having simple spectra (i.e., multiplicity of eigenvalues equal to one). In the more general case, the eigenfunctions corresponding to an eigenvalue with non-trivial multiplicity span an eigen-subspace in which one can select an arbitrary orthonormal basis. Second, due to numerical instabilities, the ordering of the eigenfunctions, especially those representing higher frequencies, is not repeatable across shapes. Finally, harmonic bases computed independently on different shapes can be expected to be reasonably compatible only when the shapes are approximately isometric, since isometries preserve the eigenfunctions of the Laplace-Beltrami operator. When this assumption is vi-

olated, it is generally impossible to expect that the n -th harmonic of one shape will correspond to the n -th harmonic of another shape. These drawbacks limit the use of harmonic bases in simultaneous shape analysis and processing to approximately isometric shapes, they do not allow to use high frequencies, and usually require some intervention to order the eigenfunctions or solve sign ambiguities.

Contributions. In this paper, we propose a general framework allowing to extend the notion of harmonic bases by finding a common (approximate) eigenbasis of multiple Laplacians. Numerically, this problem is posed as *approximate joint diagonalization*. Such methods have received limited attention in the numerical mathematics community [BGBM93] and have been employed mostly in blind source separation applications [CS96, Yer02]. Most recently, Eyraud et al. [EGBB12] introduced these approaches to machine learning, however, to the best of our knowledge, this is the first time they are applied to problems in shape analysis. The coupling between the joint approximate eigenbases relies on a given set of corresponding functions on two shapes, which in particular settings can be rough sparse point correspondences or stable regions established using standard methods [BBK06, LBB11]. We show a few examples of applications of such coupled quasi-harmonic bases in Section 5. Additional examples are shown in the preliminary version of this paper [KBB*12].

2. Background

Let us be given a compact two-dimensional manifold X . Given a smooth scalar field f on X , the negative divergence of the gradient of a scalar field, $\Delta f = -\text{div} \nabla f$, is called the *Laplace-Beltrami operator* of f and can be considered a generalization of the standard notion of the Laplace operator to manifolds [Tau95, LZ09]. The Laplace-Beltrami operator admits an eigendecomposition with non-negative eigenvalues λ and corresponding orthonormal eigenfunctions ϕ ,

$$\Delta \phi = \lambda \phi \quad (1)$$

where orthonormality is understood in the sense of the inner product $\langle f, g \rangle = \int_X f g da$ on the space $\mathcal{F}(X, \mathbb{R})$ of real functions on X , where da is the area element induced by the Riemannian metric on the manifold.

Furthermore, due to the assumption that our manifold is compact, the spectrum is discrete, $0 = \lambda_1 \leq \lambda_2 \leq \dots$ [BGM71]. In physics, (1) is known as the *Helmholtz equation* representing the spatial component of the wave equation. Thinking of our shape as of a vibrating membrane, the eigenfunctions ϕ_i can be interpreted as natural vibration modes of the membrane, while the λ_i 's assume the meaning of the corresponding vibration frequencies [TS90]. The eigenbasis of the Laplace-Beltrami operator is frequently referred to as the *harmonic basis* of the manifold, and the functions ϕ_i as *manifold harmonics* [VL08].

There exist numerous ways of approximating the Laplace-Beltrami operator and its eigenfunctions on different discrete representations of manifolds [VL08]. In computer graphics, the manifold X is typically represented as a triangular mesh built upon the vertex set $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, and a function $f \in \mathcal{F}(X, \mathbb{R})$ is represented by the vector $\mathbf{f} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^T$ of its samples. A common approach to discretizing manifold harmonics is by first constructing a discrete Laplace-Beltrami operator on the mesh, represented as an $n \times n$ matrix \mathbf{L} , followed by its eigendecomposition, $\mathbf{L}\Phi = \Phi\Lambda$, where $\Phi = (\phi_1, \dots, \phi_n)$ is the matrix of eigenvectors arranged as columns, discretizing the eigenfunctions of the Laplace-Beltrami operator at the sampled points, and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$. A popular discretization is the cotangent scheme [PP93, MDSB03] where $\mathbf{L} = \mathbf{D}^{-1}\mathbf{W}$,

$$\mathbf{D} = \text{diag}(s_1, \dots, s_n)/3, \quad (2)$$

$$w_{ij} = \begin{cases} (\cot(\alpha_{ij}) + \cot(\beta_{ij}))/2 & i \neq j; \\ -\sum_{k \neq i} w_{ik} & i = j, \end{cases} \quad (3)$$

$s_i > 0$ denotes the sum of the areas of all triangles sharing the vertex i , and α_{ij}, β_{ij} are the two angles opposite to the edge between vertices i and j in the two triangles sharing the edge [WBH*07, RWP06]. The eigendecomposition of such a Laplacian is often more conveniently written as the generalized eigenvalue problem $\mathbf{W}\Phi = \mathbf{D}\Phi\Lambda$. Numerically, this problem can be posed as the minimization

$$\min_{\Phi} \text{off}(\Phi^T \mathbf{W} \Phi) \text{ s.t. } \Phi^T \mathbf{D} \Phi = \mathbf{I}, \quad (4)$$

of the sum of squared off-diagonal elements, $\text{off}(\mathbf{A}) = \sum_{i \neq j} a_{ij}^2$ [CS96]. Several classical algorithms for finding eigenvectors such as the Jacobi method in fact try to reduce the off-diagonal values in an iterative way.

Laplace-Beltrami eigenbases are equivalent to Fourier bases on Euclidean domains, and allow to represent square-integrable functions on the manifold as linear combinations of eigenfunctions, akin to Fourier analysis. In particular, solutions of linear PDEs on non-Euclidean domains can be expressed in the Laplacian eigenbasis, giving rise to numerous efficient methods for computing e.g. local descriptors based on fundamental solutions of heat and wave equations [SOG09, ASC11], isometric embeddings of shapes [BN03, Rus07], diffusion metrics [CL06], shape correspondence and similarity [RWP05, BBK*10, OMMG10, DK10]

Importantly, manifold harmonics depend on the domain on which they are defined. Many applications working with several shapes (such as shape matching or pose transfer) rely on the fact that harmonic bases defined on two or more different shapes are consistent and behave in a similar way [Lév06]. While experimentally it is known that often low-frequency harmonics have similar behavior (finding protrusions in shapes, a fact often employed for shape segmentation [Reu10]), there is no theoretical guarantee whatsoever of such behavior. Theoretically, consistent behavior of eigenfunctions can be guaranteed only in the case of isomet-

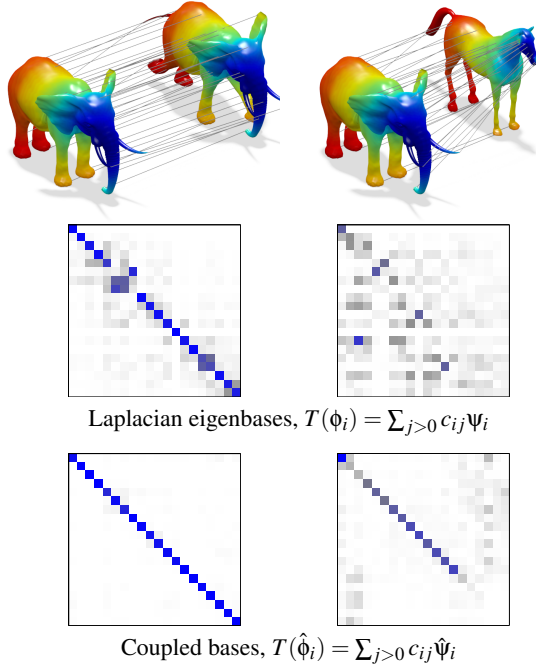


Figure 1: Matrix \mathbf{C} of coefficients expressing a given correspondence between two poses of an elephant (left) and elephant and horse (right) in the Laplacian eigenbases (second row) and coupled bases (third row). First row: correspondence between shapes shown with similar colors.

ric (or more generally conformal) shapes with simple spectrum [OBCS*12], which in practice may be too restrictive. As an illustration, we give here two examples of applications in which the assumption of consistent behavior of basis functions is especially crucial. Additional applications are discussed in Section 5.

Functional correspondence. Ovsjanikov et al. [OBCS*12] proposed an elegant way to avoid direct representation of correspondences as maps between shapes using a functional representation. The authors noted that when two shapes X and Y are related by a bijective correspondence $t : X \rightarrow Y$, then for any real function $f : X \rightarrow \mathbb{R}$, one can construct a corresponding function $g : Y \rightarrow \mathbb{R}$ as $g = f \circ t^{-1}$. In other words, the correspondence t uniquely defines a mapping between two function spaces $T : \mathcal{F}(X, \mathbb{R}) \rightarrow \mathcal{F}(Y, \mathbb{R})$. Furthermore, such a mapping is linear.

Equipping X and Y with harmonic bases, $\{\phi_i\}_{i \geq 1}$ and $\{\psi_j\}_{j \geq 1}$, respectively, one can represent a function $f : X \rightarrow \mathbb{R}$ using the set of (generalized) Fourier coefficients $\{a_i\}_{i \geq 1}$ as $f = \sum_{i \geq 1} a_i \phi_i$. Then, translating the representation into the other harmonic basis, one obtains a simple representation of the correspondence between the shapes

$$T(f) = \sum_{i,j \geq 1} a_i c_{ij} \psi_j, \quad (5)$$




Figure 2: Pose transfer from horse (top left) to camel shape (top right) by substituting the first 6 Fourier coefficients in the decomposition of extrinsic coordinates of the shape in the Laplacian eigenbasis as done in [Lév06] (bottom left) and coupled basis (bottom right).

where c_{ij} are Fourier coefficients of the basis functions of X expressed in the basis of Y , defined as $T(\phi_i) = \sum_{j \geq 1} c_{ij} \psi_j$. In the discrete setting, taking the first k basis elements, this relation can be expressed as $\mathbf{b}^T = \mathbf{a}^T \mathbf{C}$, where \mathbf{a}, \mathbf{b} are k -dimensional vectors of Fourier coefficients of discretized functions \mathbf{f}, \mathbf{g} , respectively, and \mathbf{C} is a $k \times k$ correspondence matrix [OBCS*12]. In this representation, the computation of the shape correspondence $t : X \rightarrow Y$ is translated into a simpler task of finding \mathbf{C} from a set of correspondence constraints. This matrix has a diagonal structure if the harmonic bases are compatible, an assumption crucial for the efficient computation of the correspondence. However, the authors report that in practice the elements of \mathbf{C} spread off the diagonal with the increase of the frequency due to the lack of perfect compatibility of the harmonic bases.

Pose transfer. Lévy [Lév06] proposed a pose transfer approach based on the Fourier decomposition of the manifold embedding coordinates. Given two shapes X and Y embedded in \mathbb{R}^3 with the corresponding harmonic bases $\{\phi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$, respectively, one first obtains the Fourier decompositions of the embeddings

$$\mathbf{X} = \sum_{i \geq 1} \mathbf{a}_i \phi_i, \quad \mathbf{Y} = \sum_{i \geq 1} \mathbf{b}_i \psi_i \quad (6)$$

(we denote by \mathbf{X} and \mathbf{Y} the Euclidean embeddings of manifolds X and Y , and by $\mathbf{a}_i, \mathbf{b}_i$ the three-dimensional vectors of the Fourier coefficients corresponding to each embedding coordinate). Next, a new shape Z is composed according to

$$\mathbf{Z} = \sum_{j=1}^n \mathbf{a}_j \psi_j + \sum_{i>n} \mathbf{b}_i \psi_i, \quad (7)$$

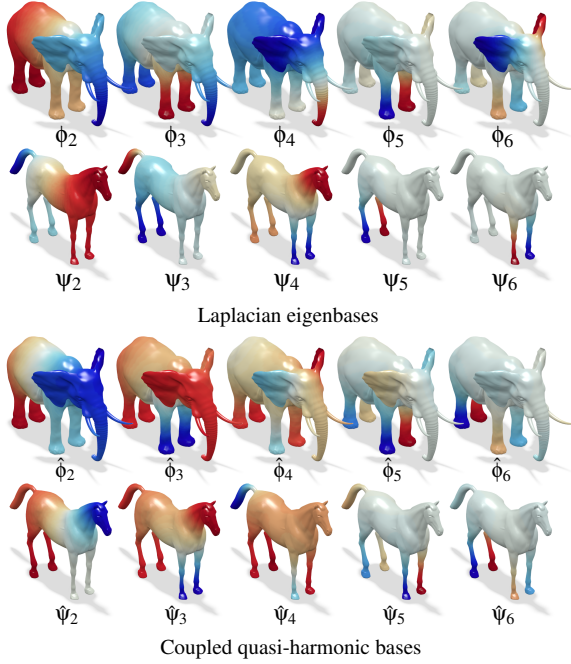


Figure 3: Top: Laplace-Beltrami eigenfunctions $\{\phi_i\}$ and $\{\psi_i\}$ of elephant and horse shapes. Bottom: coupled basis function $\{\hat{\phi}_i\}, \{\hat{\psi}_i\}$. Hot and cold colors represent positive and negative values, respectively.

with the first n low frequency coefficients taken from X , and higher frequencies taken from Y . This transfers the “layout” (pose) of the shape X to the shape Y while preserving the geometric details of Y . This method works when the \mathbf{a}_i ’s and the \mathbf{b}_i ’s are expressed in the same “language”, i.e., when the Laplacian eigenfunctions behave consistently in X and Y .

3. Coupled quasi-harmonic bases

Let us be given two shapes X, Y with the corresponding Laplacians Δ_X, Δ_Y .[†] If X and Y are related by an isometry $t: X \rightarrow Y$ and have simple Laplacian spectrum (no eigenvalues with multiplicity greater than 1), the eigenfunctions are defined up to a sign flip, $\psi_i = \pm \phi_i \circ t^{-1}$. If some eigenvalue $\lambda_i = \dots = \lambda_{i+p}$ has multiplicity $p+1$, the individual eigenvectors are not defined, but rather the subspaces they span: $\text{span}\{\psi_i, \dots, \psi_{i+p}\} = \text{span}\{\phi_i, \dots, \phi_{i+p}\} \circ t^{-1}$. More generally, if the shapes are not isometric, the behavior of their eigenfunctions can differ dramatically (Figure 3, top).

This poses severe limitations on applications we mentioned in the previous section. Figure 1 (middle) shows the coefficient matrix \mathbf{C} defined in (5), representing the functional correspondence between two shapes. For near-isometric shapes (two deformations of an elephant, Figure 1, middle left), since $\psi_i \approx \pm \phi_i \circ t^{-1}$, the coefficients

$c_{ij} \approx \pm \delta_{ij}$, and thus the matrix \mathbf{C} is nearly diagonal. However, when trying to express correspondence between non-isometric shapes (elephant and horse, Figure 1, middle right), the Laplacian eigenfunctions manifest a very different behavior breaking this diagonality.

The same problem is observed when we try to use the technique of Lévy [Lév06] for pose transfer by expressing the embedding coordinates of the shape in the respective Laplacian eigenbasis and substituting the low-frequency coefficients from another shape. Lévy disclaims that his method works “provided that the eigenfunctions that correspond to the lower frequencies match” [Lév06]. However, such a consistent behavior is not guaranteed at all; Figure 2 (bottom right) shows how the pose transfer breaks when the eigenfunction are inconsistent.

Main idea. We try to find bases $\hat{\phi}_i, \hat{\psi}_i$ that approximately diagonalize the respective Laplacians ($\Delta_X \hat{\phi} \approx \lambda_X \hat{\phi}$, $\Delta_Y \hat{\psi} \approx \lambda_Y \hat{\psi}$) and are coupled ($\hat{\psi}_i \approx \hat{\phi}_i \circ t^{-1}$). Such new coupled bases, while being nearly harmonic, make the basis functions consistent across shapes and cure the problems we outlined above (Figure 3, bottom). Figure 1 (bottom) shows the functional correspondence represented in the coupled bases $\{\hat{\phi}_i\}_{i \geq 1}, \{\hat{\psi}_i\}_{i \geq 1}$. Due to the coupling, the basis functions behave consistently resulting in almost perfectly diagonal matrices \mathbf{C} even when the shapes are highly non-isometric (bottom right). Likewise, Figure 2 (rightmost) shows that pose transfer using Fourier coefficients in the coupled bases works correctly for shapes with inconsistent Laplacian eigenfunctions.

Approximate joint diagonalization. We assume that the shapes are sampled at n_X, n_Y points, and their Laplacians are discretized as matrices $\mathbf{W}_X, \mathbf{W}_Y$ and $\mathbf{D}_X, \mathbf{D}_Y$ of size $n_X \times n_X$ and $n_Y \times n_Y$ respectively, as defined in (3). We further assume to be given a set of p corresponding functions $g_i \approx f_i \circ t^{-1}$, represented in the discrete setting by matrices $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_p)$ and $\mathbf{G} = (\mathbf{g}_1, \dots, \mathbf{g}_p)$ of size $n_X \times p$ and $n_Y \times p$, respectively.[‡] The bases behave consistently if the respective Fourier coefficients of the functions f_1, \dots, f_p and g_1, \dots, g_p are approximately equal $\langle f_i, \hat{\phi}_j \rangle \approx \langle g_i, \hat{\psi}_j \rangle$, or in other words, $\mathbf{F}^T \hat{\Phi} \approx \mathbf{G}^T \hat{\Psi}$ (Fourier coupling).

The problem of joint approximate diagonalization (JAD) can be formulated as the coupling of two problems (4),

$$\begin{aligned} \min_{\hat{\Phi}, \hat{\Psi}} \quad & \text{off}(\hat{\Phi}^T \mathbf{W}_X \hat{\Phi}) + \text{off}(\hat{\Psi}^T \mathbf{W}_Y \hat{\Psi}) + \mu \|\mathbf{F}^T \hat{\Phi} - \mathbf{G}^T \hat{\Psi}\|_F^2 \\ \text{s.t.} \quad & \hat{\Phi}^T \mathbf{D}_X \hat{\Phi} = \mathbf{I}, \quad \hat{\Psi}^T \mathbf{D}_Y \hat{\Psi} = \mathbf{I} \end{aligned} \quad (8)$$

where off denotes some off-diagonality penalty, e.g., the sum of the squared off-diagonal elements as defined in Section 2. The simplest choice of \mathbf{F} and \mathbf{G} is having columns containing

[†] We consider the case of a pair of shapes for the mere sake of simplicity. Extension to a collection of more shapes is straightforward.

[‡] We assume that \mathbf{F} and \mathbf{G} are correctly normalized by the local area elements $\mathbf{D}_X, \mathbf{D}_Y$.

a single one at points corresponding in the two shapes; in this case, the coupling term forces the basis functions to be approximately equal at the corresponding points. We call this case *point-wise coupling*. Another possible choice is setting \mathbf{F} and \mathbf{G} as the indicator functions of corresponding stable regions e.g. detected using shape MSER [LBB11].

The parameter μ determines the coupling strength. For $\mu = 0$, the problem becomes uncoupled and boils down to individual diagonalization (4) of the two Laplacians. The joint eigenvectors in this setting coincide with the eigenvectors of the Laplacians: $\hat{\Phi} = \Phi$ and $\hat{\Psi} = \Psi$.

Subspace parametrization. We can parametrize the joint basis functions as linear combinations of the Laplacian eigenvectors, $\hat{\Phi} = \Phi \mathbf{A}$ and $\hat{\Psi} = \Psi \mathbf{B}$, where \mathbf{A} and \mathbf{B} are matrices of combination coefficients of size $n_X \times n_X$ and $n_Y \times n_Y$, respectively. Noticing that $\hat{\Phi}^T \mathbf{W}_X \hat{\Phi} = \mathbf{A}^T \Phi^T \mathbf{W}_X \Phi \mathbf{A} = \mathbf{A}^T \Lambda_X \mathbf{A}$, and similarly, $\hat{\Psi}^T \mathbf{W}_Y \hat{\Psi} = \mathbf{B}^T \Lambda_Y \mathbf{B}$, we transform problem (8) into

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}} \quad & \text{off}(\mathbf{A}^T \Lambda_X \mathbf{A}) + \text{off}(\mathbf{B}^T \Lambda_Y \mathbf{B}) + \mu \|\mathbf{F}^T \Phi \mathbf{A} - \mathbf{G}^T \Psi \mathbf{B}\|_F^2 \\ \text{s.t.} \quad & \mathbf{A}^T \mathbf{A} = \mathbf{I}, \quad \mathbf{B}^T \mathbf{B} = \mathbf{I} \end{aligned} \quad (9)$$

Since \mathbf{A} and \mathbf{B} are orthonormal, they act as isometries in the respective eigenspaces of the Laplacians of X and Y . We can thus think geometrically of problem (9) as an attempt to rotate and reflect the eigenbases Φ and Ψ such that they align in the best way (in the least squares sense), while still approximately diagonalizing the Laplacians. Note that the Laplacians do not appear explicitly in problem (9) but rather their discretized eigenfunctions and eigenvalues; for this reason, we can employ any method for Laplacian discretization.

Since in many applications we are not interested in the entire eigenbasis but in the first k eigenvectors, this formulation is especially convenient, as it allows us to express the first k joint eigenvectors as a linear combination of k' eigenvectors (we provide a justification of this assumption in Appendix A), thus having the matrices \mathbf{A} and \mathbf{B} of size $k' \times k$:

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}} \quad & \text{off}(\mathbf{A}^T \bar{\Lambda}_X \mathbf{A}) + \text{off}(\mathbf{B}^T \bar{\Lambda}_Y \mathbf{B}) + \mu \|\mathbf{F}^T \bar{\Phi} \mathbf{A} - \mathbf{G}^T \bar{\Psi} \mathbf{B}\|_F^2 \\ \text{s.t.} \quad & \mathbf{A}^T \mathbf{A} = \mathbf{I}, \quad \mathbf{B}^T \mathbf{B} = \mathbf{I}, \end{aligned} \quad (10)$$

where $\bar{\Phi} = (\phi_1, \dots, \phi_{k'})$ and $\bar{\Lambda}_X = \text{diag}(\lambda_k^X, \dots, \lambda_{k'}^X)$; matrices $\bar{\Psi}, \bar{\Lambda}_Y$ are defined accordingly. Typically, $k, k' \ll n_X, n_Y$, and thus the problem is much smaller than the full joint diagonalization (8). Note that the coupling term provides pk constraints, so in order not to over-determine the problem, we should have $2k' > p$. Typical values used in our experiments were $k = k' \sim 20$ and $p \sim 20$.

It is important to note that in problem (10) the use of the sum of squared off-diagonal elements as the off-penalty does not produce an ordered set of approximate joint eigenvectors. This issue can be solved by using an alternative penalty,

$$\|\hat{\Phi}^T \mathbf{L}_X \hat{\Phi} - \bar{\Lambda}_X\|_F^2 = \|\mathbf{A}^T \bar{\Lambda}_X \mathbf{A} - \bar{\Lambda}_X\|_F^2, \quad (11)$$

which is similar to the sum of squared off-diagonal elements but also includes the difference of the diagonal elements [Yer02]. The penalty for the shape Y is defined in the same way.

Procrustes problem. In the limit case $\mu \rightarrow \infty$ we can ignore the off-diagonal penalties and problem (10) becomes

$$\min_{\mathbf{A}, \mathbf{B}} \|\mathbf{F}^T \bar{\Phi} \mathbf{A} - \mathbf{G}^T \bar{\Psi} \mathbf{B}\|_F^2 \quad \text{s.t.} \quad \mathbf{A}^T \mathbf{A} = \mathbf{I}, \quad \mathbf{B}^T \mathbf{B} = \mathbf{I}, \quad (12)$$

Using the invariance of the Frobenius norm under orthogonal transformation, we can rewrite problem (12) as an *orthogonal Procrustes problem*

$$\min_{\Omega} \|\mathbf{F}^T \bar{\Phi} - \mathbf{G}^T \bar{\Psi} \Omega\|_F^2 \quad \text{s.t.} \quad \Omega^T \Omega = \mathbf{I}, \quad (13)$$

where $\Omega = \mathbf{B} \mathbf{A}^T$. The problem has an analytic solution $\Omega = \mathbf{S} \mathbf{R}^T$, where $\bar{\Phi}^T \mathbf{F} \mathbf{G}^T \bar{\Psi} = \mathbf{S} \mathbf{\Sigma} \mathbf{R}^T$ is the singular value decomposition of the matrix $\bar{\Phi}^T \mathbf{F} \mathbf{G}^T \bar{\Psi}$ with left- and right singular vectors \mathbf{S}, \mathbf{R} [Sch66]. Then, $\mathbf{A} = \mathbf{R}$ and $\mathbf{B} = \mathbf{S}$.

4. Numerical computation

Problem (10) is a non-linear optimization with orthogonality constraints. In our experiments, we used the first-order constrained minimization algorithm implemented in MATLAB Optimization Toolbox. We provide below the gradients of our cost function.

Gradient of the off-diagonality penalty is given by

$$\nabla_{\mathbf{A}} \sum_i (\mathbf{A}^T \bar{\Lambda}_X \mathbf{A})_{ii}^2 = 4(\bar{\Lambda}_X \mathbf{A} \mathbf{A}^T \bar{\Lambda}_X \mathbf{A} - \mathbf{O} \circ \mathbf{A} \bar{\Lambda}_X), \quad (14)$$

where \mathbf{O} is a matrix of equal columns $\mathbf{o} = \text{diag}(\mathbf{A}^T \bar{\Lambda}_X \mathbf{A})$ and \circ denotes element-wise product of matrices (see derivation in Appendix B). Gradient of the alternative penalty (11) is derived similarly as

$$\nabla_{\mathbf{A}} \|\mathbf{A}^T \bar{\Lambda}_X \mathbf{A} - \bar{\Lambda}_X\|_F^2 = 4(\bar{\Lambda}_X \mathbf{A} \mathbf{A}^T \bar{\Lambda}_X \mathbf{A} - \bar{\Lambda}_X \mathbf{A} \bar{\Lambda}_X). \quad (15)$$

Gradient of the coupling term w.r.t. to \mathbf{A} is given by

$$\nabla_{\mathbf{A}} \|\mathbf{F}^T \bar{\Phi} \mathbf{A} - \mathbf{G}^T \bar{\Psi} \mathbf{B}\|_F^2 = 2\bar{\Phi}^T \mathbf{F} (\mathbf{F}^T \bar{\Phi} \mathbf{A} - \mathbf{G}^T \bar{\Psi} \mathbf{B}); \quad (16)$$

the gradient w.r.t. \mathbf{B} is obtained in the same way.

Initialization. Assuming the Laplacians are nearly jointly diagonalizable and have simple spectrum, their joint eigenvectors will be equal to the harmonic basis functions up to sign flips. Thus, in this case \mathbf{A} and \mathbf{B} are diagonal matrices of ± 1 . This was found to be a reasonable initialization to the joint diagonalization procedure. We set $\mathbf{A} = \mathbf{I}$, and then solve sign flip by setting the elements of \mathbf{B} to

$$b_{ij} = \begin{cases} +1 & i = j, \quad \|\mathbf{F}^T \phi_i - \mathbf{G}^T \psi_i\| \leq \|\mathbf{F}^T \phi_i + \mathbf{G}^T \psi_i\|; \\ -1 & i = j, \quad \|\mathbf{F}^T \phi_i - \mathbf{G}^T \psi_i\| > \|\mathbf{F}^T \phi_i + \mathbf{G}^T \psi_i\|; \\ 0 & \text{else.} \end{cases} \quad (17)$$

Initialized this way, the problem starts with decoupled but diagonalizing bases, and the optimization tries to improve the coupling.




Figure 4: Examples of joint diagonalization of Laplacians of near-isometric shapes (two poses of an elephant, left) and non-isometric shapes (elephant and horse, right). Point-wise coupling was done using 40 points. Numbers show the ratio of the norm of the off-diagonal and diagonal values.




Figure 5: Sensitivity of joint diagonalization to errors in correspondence (in % of geodesic diameter of the shape) used in the coupling term. Shapes are shown with similar colors representing corresponding points. Correspondences between 10 points used for coupling are shown with lines.

5. Results and Applications

In this section, we show several examples of coupled bases construction, as well as some potential applications of the proposed approach (for additional examples, see [KBB*12]). We used shapes from publicly available datasets [BBK08, SP04, SMKF04]. Mesh sizes varied widely between 600 - 25K vertices. Discretization of the Laplace-Beltrami operator was done using the cotangent formula [MDSB03]. In all our examples, we constructed coupled bases solving the JAD problem (10) with off-diagonality penalty (11), as described in Section 4; the value of $\mu = 0.132$ was used in all experiments. Optimization was performed using MATLAB `fmincon` with default settings. Typical time to compute 15 joint eigenvectors was about 1 minute.

Isometric vs non-isometric. Figure 4 (top) shows examples of joint diagonalization of Laplacians of different shapes: near isometric (two poses of an elephant) and non-

isometric (elephant and horse). We computed the first $k = 20$ joint approximate eigenvectors. The Laplacians are almost perfectly diagonalized by the obtained coupled bases in the case of near-isometric shapes; for non-isometric shapes, off-diagonal elements are more prominent. Nevertheless, a clear diagonally-dominant structure is present.

Sensitivity to correspondence error. In this experiment, we computed the coupled bases for two near-isometric human shapes using point-wise coupling at $p = 10$ points with noisy correspondence that deviated from groundtruth correspondence by up to 15% of the geodesic diameter of the shape. Table 1 shows the obtained diagonalization quality (measured as the ratio of the norm of the off-diagonal and diagonal values averaged on two shapes), and Figure 5 depicts the approximately diagonalized Laplacians. This experiment illustrates that very few roughly corresponding points are required for the coupling term in our problem, and that the proposed procedure is robust to correspondence noise.

Sensitivity to representation and sampling. Taking the shapes from the previous experiment, we downsampled and re-meshed one of them from 8.5K vertices to 850 vertices. In addition, we removed the triangulation and computed a Laplacian on the point cloud, using 6 nearest neighbors and Gaussian weight $w_{ij} = e^{-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / 2 \times 10^{-4}}$. Table 2 shows the quality of the joint diagonalization of the Laplacians on the original mesh and subsampled mesh (first row) and the point cloud (second row). Figures 6-7 show the obtained coupled bases.

exact	3%	6%	15%
0.0417	0.0534	0.0892	0.1439

Table 1: Sensitivity of JAD quality $\|\text{off}\|/\|\text{diag}\|$ to correspondence error (measured as % of geodesic diameter).

	full	90%	75%	50%	10%
mesh-mesh	0.0417	0.0414	0.0384	0.0429	0.0625
mesh-cloud	0.0934	0.1399	0.1265	0.1303	0.1424

Table 2: Sensitivity of JAD quality $\|\text{off}\|/\|\text{diag}\|$ to sampling density for two meshes and mesh/point cloud.

Shape correspondence. Ovsjanikov et al. [OBCS*12] computed correspondence between near-isometric shapes from a set of constraints on the $k \times k$ matrix \mathbf{C} (encoding the correspondence represented using the first k harmonic functions as described in Section 2). Given a set of p corresponding functions on X and Y represented using the matrices \mathbf{F} and \mathbf{G} as defined in Section 4, \mathbf{C} is recovered from a system of pk equations with k^2 variables,

$$\mathbf{F}^T \Phi = \mathbf{G}^T \Psi \mathbf{C} \quad (18)$$

Additional constraints stemming from the properties of the matrix \mathbf{C} are also added [OBCS*12].

The use of our coupled bases in place of standard Laplace-Beltrami eigenbases allows to exploit the sparse structure of \mathbf{C} , which is mentioned by Ovsjanikov et al., but



Figure 6: Coupled bases computed on a full (top) and 10-times subsampled (bottom) mesh.




Figure 7: Coupled bases computed on a full mesh (top) and 10-times subsampled point cloud (bottom).




Figure 8: Coupled bases elements of the human and gorilla shapes obtained with Fourier coupling using 25 MSER regions from Figure 9.

never used explicitly in their paper. Approximating $\mathbf{C} \approx \text{diag}(c_{11}, \dots, c_{kk})$, we can rewrite (18) as a system of pk equations with only k variables corresponding to the diagonal of \mathbf{C} ,

$$\begin{bmatrix} \text{diag}(\mathbf{g}_1^T \hat{\Psi}) \\ \vdots \\ \text{diag}(\mathbf{g}_p^T \hat{\Psi}) \end{bmatrix} \begin{bmatrix} c_{11} \\ \vdots \\ c_{kk} \end{bmatrix} = \begin{bmatrix} \hat{\Phi}^T \mathbf{f}_1 \\ \vdots \\ \hat{\Phi}^T \mathbf{f}_p \end{bmatrix}. \quad (19)$$

Equation (19) allows to use significantly less data to fully determine the correspondence, and is also more computationally efficient.

Figure 10 shows an example of finding functional correspondence between non-isometric shapes of human and gorilla. As columns of \mathbf{F} and \mathbf{G} we used binary indicator functions of $p = 25$ regions detected using the MSER algorithm [LBB11]. The ordering of the regions was assumed to be given (corresponding regions are denoted by similar color in Figure 9). We compared the method described in [OBCS*12] for computing \mathbf{C} by solving the system (18) in the standard Laplace-Beltrami eigenbases (Figure 10, left) and the diagonal-only approximation (19) in the coupled bases, computed using the same data (\mathbf{F}, \mathbf{G}) for the coupling term (see Figure 8). In both cases, we used $k = 20$ first basis vectors. Then, the ICP-like refinement approach [OBCS*12] was applied to \mathbf{C} to obtain point-wise correspondence.

Simultaneous mesh editing. Rong et al. [RCG08] proposed an approach for mesh editing based on elastic energy minimization. Given a shape with embedding coordinates \mathbf{X} , the method attempts to find a deformation field \mathbf{d} producing a new shape $\mathbf{X}' = \mathbf{X} + \mathbf{d}$, providing a set of user-defined n' anchor points for which the displacement is known (w.l.o.g. assuming to be the first n' points, $\mathbf{d}_i = \mathbf{d}'_i$ for $i = 1, \dots, n'$), as a solution of the system of equations

$$\begin{bmatrix} k_b \mathbf{L}_X^2 - k_c \mathbf{L}_X & \mathbf{M} \\ \mathbf{M}^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{d} \\ \boldsymbol{\gamma} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{d}' \end{bmatrix}, \quad (20)$$

where $\mathbf{M} = (\mathbf{I}, \mathbf{0})^T$ is an $n \times n'$ identity matrix, $\boldsymbol{\gamma}$ are unknown Lagrange multipliers corresponding to the constraints on anchor points, and k_b, k_c are parameters trading off between resistance to bending and stretching, respectively [RCG08]. The system of equations can be expressed in the frequency domain using $k \ll n$ first harmonic basis functions,

$$\begin{bmatrix} \bar{\Phi}^T (k_b \mathbf{L}_X^2 - k_c \mathbf{L}_X) \bar{\Phi} & \bar{\Phi}^T \mathbf{M} \\ \mathbf{M}^T \bar{\Phi} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha} \\ \boldsymbol{\gamma} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{d}' \end{bmatrix} \quad (21)$$

where $\boldsymbol{\alpha} = \bar{\Phi}^T \mathbf{d}$ are the k Fourier coefficients. The desired deformation field is obtained by solving the system of equations for $\boldsymbol{\alpha}$ and transforming it to the spatial domain $\mathbf{d} = \bar{\Phi} \boldsymbol{\alpha}$ (for details, the reader is referred to [RCG08]).

Using coupled bases, it is possible to easily extend this approach to simultaneous editing of multiple shapes, solving the system (21) with the coupled basis $\hat{\Phi}$ in place of $\bar{\Phi}$, and applying the deformation to the second mesh using $\mathbf{d} = \hat{\Psi} \boldsymbol{\alpha}$. Figure 12 exemplifies this idea, showing how a deformation of the cat shape is automatically transferred to the lion shape, which accurately and naturally repeats the cat pose.

Shape similarity. The diagonalization quality of the Laplacians can be used as a criterion for shape similarity, with isometric shapes having ideal diagonalization. With this approach, it is possible to compare two shapes from a small



Figure 9: Some of the MSER regions detected in the human and gorilla shapes.

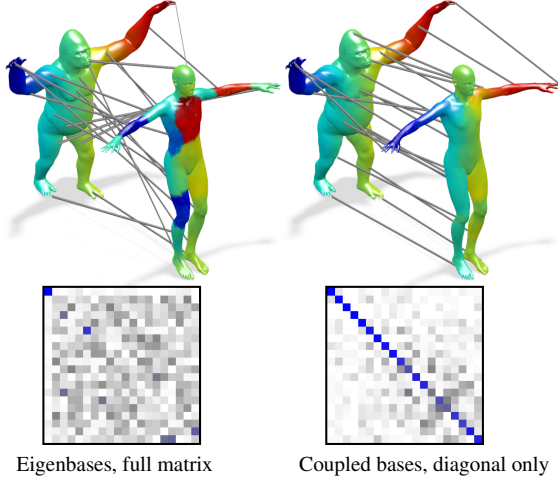


Figure 10: Functional correspondence matrix \mathbf{C} computed according to (18) using Laplace-Beltrami eigenbases (left) and diagonal-only approach (19) in coupled bases (right). Note that correspondences and matrices \mathbf{C} are shown after ICP-like refinement, altering the diagonal-only structure.

number of inaccurate correspondences provided for coupling in the joint diagonalization problem. Figure 11 shows the similarity matrix between 25 shapes belonging to 8 different classes. Each shape is present with 3-4 near-isometric deformation. We used 25 point correspondences for point-wise coupling; dissimilarity of a pair of shapes was computed by jointly diagonalizing the respective Laplacians and then computing the average ratio of the norms of the diagonal and off-diagonal elements of both matrices.

6. Conclusions

We showed several formulations of numerically efficient joint approximate diagonalization algorithms for the construction of coupled bases of the Laplacians of multiple shapes. Such quasi-harmonic bases allow to extend many shape analysis and synthesis tasks to cases where the standard harmonic bases computed on each shape separately cease being compatible. The proposed construction can be used as an alternative to the standard harmonic bases. One straightforward generalization of our method is to get rid of the need to know the correspondence between the set of

functions f_1, \dots, f_p and g_1, \dots, g_p by introducing a permutation into the coupling term as done in [PBB*12].

A particularly promising direction is non-rigid shape matching in the functional correspondence representation. The sparse structure of the matrix \mathbf{C} has not been yet taken advantage of for the computation of correspondence in a proper way, and naturally calls for sparse modeling methods successfully used in signal processing. However, the correctness of this model largely depends on the basis in which \mathbf{C} is represented, and standard Laplace-Beltrami eigenbasis performs quite poorly when one deviates from the isometry assumptions. In follow-up studies, we intend to explore the relation between sparse models and joint approximate diagonalization in shape correspondence problems. Finally, we note that the idea of joint diagonalization extends beyond Laplacians and is applicable to other differential operators such as those proposed in [HSvTP12].

Acknowledgement

A.K., M.B. and K.G. are supported by the ERC Starting grant No. 307047. A.B. is supported by the GIF and ISF. R.K. is supported by the ERC Advanced grant No. 267414.

Appendix A - Perturbation analysis of JAD

We analyze here the solution of the basic problem (8). For simplicity, we assume that $n_X = n_Y = l$, the points are ordered, and $\mu \rightarrow \infty$, in which case a single basis ($\hat{\Phi} = \hat{\Psi}$) is searched (this setting of the JAD problem has been considered in [BGBM93, CS96]). We further assume that $\mathbf{L}_X = \Phi \Lambda \Phi^T$ has a simple τ -separated spectrum (i.e., $|\lambda_i - \lambda_j| \geq \tau$). If Y is a near-isometric deformation of X , its Laplacian can be described as a perturbation $\mathbf{L}_Y = \Phi \Lambda \Phi^T + \epsilon \mathbf{R}$ of \mathbf{L}_X . Ignoring permutation of eigenfunctions and sign flips, the joint approximate eigenbasis can be written as the first-order perturbation [Car95] $\hat{\Phi}_i \approx \Phi_i + \epsilon \sum_{j \neq i} \alpha_{ij} \Phi_j$, where $\alpha_{ij} = \Phi_i^T \mathbf{R} \Phi_j / 2(\lambda_j - \lambda_i)$. We can bound these coefficients by the spectral norm $|\alpha_{ij}| \leq \|\mathbf{R}\|_2 / 2\tau = \alpha_{\max}$. We can conclude that the first k joint approximate eigenvectors $\hat{\Phi}_1, \dots, \hat{\Phi}_k$ can be well represented as linear combinations of $\Phi_1, \dots, \Phi_{k'}$, with square error bounded by $\epsilon \sum_{j > k'} |\alpha_{ij}|^2 \leq \epsilon(n_X - k' - 1)\alpha_{\max}^2$. This result also justifies the use of band-wise computation discussed in Section 4.

To relate this bound to the geometry of the shape, let us assume that X and Y have the same connectivity and that the angles $\beta, \bar{\beta}$ of triangles in the two meshes satisfy $\theta_0 \leq \beta, \bar{\beta} \leq \pi - \theta_0$ (all triangles are at least θ_0 fat), the area elements are at least $s, \bar{s} \geq s_0$, and each vertex is connected to at most v vertices. We assume that the mesh Y is obtained as a deformation of mesh X changing the angles by $|\bar{\beta} - \beta| \leq \Delta\theta$ and the area elements by $1 - \delta \leq \bar{s}/s \leq 1 + \delta$. We are interested in a bound on the spectral norm $\|\mathbf{L}_X - \mathbf{L}_Y\|_2 = \epsilon \|\mathbf{R}\|_2$




Figure 11: Shape similarity using joint diagonalization. Darker colors represent more similar shapes. One can clearly distinguish blocks of isometric shapes. Also, two classes of two- and four-legged shapes (marked with green and blue) are visible. Small figures show representative shapes from each class.

expressed in terms of parameters $\Delta\theta, \delta$ (strength of non-isometric deformation) and constants θ_0, s_0 . Using trigonometric identities in (3), we have for $i \neq j$

$$|\bar{w}_{ij} - w_{ij}| \leq \left| \frac{\sin(\bar{\beta} - \beta)}{\sin\beta\sin\bar{\beta}} \right| \leq \frac{\sin\Delta\theta}{\sin^2\theta_0} \leq \frac{\Delta\theta}{\sin^2\theta_0}, \quad (22)$$

and $|\bar{w}_{ii} - w_{ii}| \leq \frac{\nu\Delta\theta}{\sin^2\theta_0}$. Applying the triangle inequality, for $i \neq j$ we get $|\bar{l}_{ij} - l_{ij}| = 3|\bar{w}_{ij}\bar{s}_i^{-1} - w_{ij}s_i^{-1}| \leq 3|\bar{w}_{ij} - w_{ij}|s_i^{-1} + 3|\bar{w}_{ij}||\bar{s}_i^{-1} - s_i^{-1}|$. Plugging in the bounds on $|\bar{w}_{ij} - w_{ij}|$ and \bar{s}_i , we get

$$\begin{aligned} |\bar{l}_{ij} - l_{ij}| &\leq \frac{3}{s_0} \left(\frac{\Delta\theta}{\sin^2\theta_0} + \cot(\theta_0)\delta \right) \leq \frac{3}{s_0} \left(\frac{\Delta\theta}{\sin^2\theta_0} + \frac{\delta}{\theta_0} \right) \\ &\leq \frac{3(\Delta\theta + \delta)}{s_0 \sin^2\theta_0}, \end{aligned}$$

from which it follows by norm inequality

$$\|\mathbf{L}_X - \mathbf{L}_Y\|_2 \leq n_X^{1/2} \|\mathbf{L}_X - \mathbf{L}_Y\|_1 \leq \frac{6\nu n_X^{3/2}}{s_0 \sin^2\theta_0} (\Delta\theta + \delta),$$

where $\varepsilon = \Delta\theta + \delta$ is the degree of shape deformation (“lack of isometry”).

Appendix B - Off-diagonality penalty gradient

Let us rewrite the penalty as $\text{off}(\mathbf{A}) = \|\mathbf{A}^T \bar{\mathbf{L}}_X \mathbf{A}\|_F^2 - \sum_i (\mathbf{A}^T \bar{\mathbf{L}}_X \mathbf{A})_{ii}^2$. The first term is bi-quadratic and its gradient derivation is trivial, so here we derive the gradient of the

second term only. We have $\sum_i (\mathbf{A}^T \bar{\mathbf{L}}_X \mathbf{A})_{ii}^2 = \sum_i \left(\sum_k a_{ki}^2 \lambda_k \right)^2$; differentiating in a coordinate-wise manner,

$$\begin{aligned} \frac{\partial}{\partial a_{pq}} \sum_i \left(\sum_k a_{ki}^2 \lambda_k \right)^2 &= 2 \sum_i \left(\sum_k a_{ki}^2 \lambda_k \right) 2a_{pq} \lambda_p \delta_{iq} \\ &= 4 \sum_k a_{kq}^2 \lambda_k a_{pq} \lambda_p. \end{aligned} \quad (23)$$

Observing that $\mathbf{A}\mathbf{A} = (a_{pq} \lambda_p)$ and denoting by $o_{pq} = (\sum_k a_{kq}^2 \lambda_k)$ the elements of the equal-columns matrix \mathbf{O} , we get the expression $4\mathbf{O} \circ \mathbf{A}\mathbf{A}$.

References

- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: a quantum mechanical approach to shape analysis. In *Proc. Dynamic Shape Capture and Analysis* (2011). 1, 2
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Science (PNAS)* 103, 5 (2006), 1168–1172. 2
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical geometry of non-rigid shapes*. Springer, 2008. 6
- [BBK*10] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. 266–286. 2
- [BGBM93] BUNSE-GERSTNER A., BYERS R., MEHRMANN V.: Numerical methods for simultaneous diagonalization. *SIAM J. Matrix Anal. Appl.* 14, 4 (1993), 927–949. 2, 8
- [BGM71] BERGER M., GAUDUCHON P., MAZET E.: *Le spectre d’une variété riemannienne*. Springer, 1971. 2
- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR* (2010). 1
- [BN03] BELKIN M., NIYOGI P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* 13 (2003), 1373–1396. 2
- [Car95] CARDOSO J. F.: Perturbation of joint diagonalizers. 8
- [CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *App. Comp. Harmonic Analysis* 21 (2006), 5–30. 1, 2
- [CS96] CARDOSO J.-F., SOULOUMIAC A.: Jacobi angles for simultaneous diagonalization. *SIAM J. Matrix Anal. Appl.* 17 (1996), 161–164. 2, 8
- [DK10] DUBROVINA A., KIMMEL R.: Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *Proc. 3DPVT* (2010). 2
- [EGBB12] EYNARD D., GLASHOFF K., BRONSTEIN M. M., BRONSTEIN A. M.: Multimodal diffusion geometry by joint diagonalization of laplacians. *arXiv:1209.2295* (2012). 2
- [FH05] FLOATER M., HORMANN K.: Surface parameterization: a tutorial and survey. *Advances in Multiresolution for Geometric Modelling 1* (2005). 1
- [HSVTP12] HILDEBRANDT K., SCHULZ C., VON TYCOWICZ C., POLTHIER K.: Modal shape analysis beyond Laplacian. *Computer Aided Geometric Design* 29 (2012), 204–218. 8




Figure 12: Simultaneous shape editing in the frequency domain using the approach of [RCG08]. Top: editing the cat shape (anchor vertices in problem (20) shown in green). Bottom: the same pose is transferred to the lion shape using coupled basis.

- [KBB*12] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *arXiv:1210.0026* (2012). 2, 6
- [KG00] KARNI Z., GOTSMAN C.: Spectral compression of mesh geometry. In *Proc. Conf. Computer Graphics and Interactive Techniques* (2000), pp. 279–286. 1
- [Kob97] KOBBELT L.: Discrete fairing. In *Proc. IMA Conf. Mathematics of Surfaces* (1997). 1
- [KR05] KIM B., ROSSIGNAC J.: Geofilter: Geometric selection of mesh filter parameters. *Comput. Graph. Forum* 24, 3 (2005), 295–302. 1
- [LBB11] LITMAN R., BRONSTEIN A. M., BRONSTEIN M. M.: Diffusion-geometric maximally stable component detection in deformable shapes. *CG* 35, 3 (2011), 549–560. 2, 5, 7
- [Lév06] LÉVY B.: Laplace-Beltrami eigenfunctions towards an algorithm that ‘understands’ geometry. In *Proc. SMI* (2006). 1, 2, 3, 4
- [LZ09] LÉVY B., ZHANG R. H.: Spectral geometry processing. In *ACM SIGGRAPH ASIA Course Notes* (2009). 1, 2
- [MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization & Mathematics* (2003), 35–57. 1, 2, 6
- [NISA06] NEALEN A., IGARASHI T., SORKINE O., ALEXA M.: Laplacian mesh optimization. In *Proc. Conf. Computer Graphics and Interactive Techniques* (2006), pp. 381–389. 1
- [OBBS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *TOG* 31, 4 (2012). 1, 3, 6, 7
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *Computer Graphics Forum* 29, 5 (2010), 1555–1564. 2
- [PBB*12] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *arXiv:1209.6560* (2012). 8
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2 (1993), 15–36. 1, 2
- [RCG08] RONG G., CAO Y., GUO X.: Spectral mesh deformation. *The Visual Computer* 24, 7 (2008), 787–796. 7, 10
- [Reu10] REUTER M.: Hierarchical shape segmentation and registration via topological features of Laplace-Beltrami eigenfunctions. *IJCV* 89, 2 (2010), 287–308. 2
- [Rus07] RUSTAMOV R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. of SGP* (2007), pp. 225–233. 1, 2
- [RWP05] REUTER M., WOLTER F. E., PEINECKE N.: Laplace-spectra as fingerprints for shape matching. In *Proc. ACM Symp. Solid and Physical Modeling* (2005), pp. 101–106. 1, 2
- [RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-Beltrami spectra as “shape-DNA” of surfaces and solids. *Computer Aided Design* 38 (2006), 342–366. 2
- [Sch66] SCHÖNEMANN P.: A generalized solution of the orthogonal procrustes problem. *Psychometrika* 31, 1 (1966), 1–10. 5
- [SMKF04] SHILANE P., MIN P., KAZHDAN M., FUNKHOUSER T.: The Princeton shape benchmark. In *Proc. SMI* (2004). 6
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP* (2009). 1, 2
- [SP04] SUMNER R., POPOVIĆ J.: Deformation transfer for triangle meshes. In *TOG* (2004), vol. 23, pp. 399–405. 6
- [Tau95] TAUBIN G.: *A signal processing approach to fair surface design*. ACM, 1995. 1, 2
- [TS90] TIKHONOV A. N., SAMARSKI A. A.: *Equations of Mathematical Physics*. Dover, 1990. 2
- [VL08] VALLET B., LÉVY B.: Spectral geometry processing with manifold harmonics. *Computer Graphics Forum* (2008). 1, 2
- [WBH*07] WARDETZKY M., BERGOU M., HARMON D., ZORIN D., GRINSUN E.: Discrete quadratic curvature energies. *CAGD* 24 (2007), 499–518. 2
- [WMKG08] WARDETZKY M., MATHUR S., KÄLBERER F., GRINSUN E.: Discrete Laplace operators: no free lunch. In *Proc. SGP* (2008). 1
- [Yer02] YEREDOR A.: Non-orthogonal joint diagonalization in the least-squares sense with application in blind source separation. *Trans. Signal Proc.* 50, 7 (2002), 1545–1553. 2, 5