

An Elastic Basis for Spectral Shape Correspondence

Florine Hartwig
florine.hartwig@uni-bonn.de
University of Bonn
Bonn, Germany

Josua Sassen
josua.sassen@uni-bonn.de
University of Bonn
Bonn, Germany

Omri Azencot
azencot@cs.bgu.ac.il
Ben Gurion University of the Negev
Beer Sheva, Israel

Martin Rumpf
martin.rumpf@uni-bonn.de
University of Bonn
Bonn, Germany

Mirela Ben-Chen
mirela@cs.technion.ac.il
Technion - Israel Institute of
Technology
Haifa, Israel




Figure 1: Correspondence between two animal shapes. We compute correspondences between shapes by computing a functional map using (projected) eigenfunctions of an elastic shell energy’s Hessian as a functional basis. These basis functions are sensitive to extrinsic features such as extremities and creases (left). Hence, using them in a functional map pipeline enables us to accurately align crease lines such as mouth, ears, and toes (right). Here, we visualize the resulting correspondence by a pullback of normals from the ferret (top) to the weasel (bottom). The (transferred) normal directions are mapped to colors as shown on the little sphere.

ABSTRACT

Finding correspondences between shapes is a central task in geometry processing with applications such as texture or deformation transfer and shape interpolation. We develop a spectral method for finding correspondences between non-isometric shapes that aligns *extrinsic* features. For this, we propose a novel *crease aware* spectral basis, that is derived from the Hessian of an elastic thin shell energy. We incorporate this basis in a functional map framework and demonstrate the effectiveness of our approach for mapping non-isometric shapes such that prominent features are put in correspondence. Finally, we describe the necessary adaptations to the functional map framework for working with *non-orthogonal* basis functions, thus considerably widening the scope of future uses of spectral shape correspondence.

CCS CONCEPTS

• Computing methodologies → Shape analysis.

KEYWORDS

shape matching, functional maps, spectral methods

ACM Reference Format:

Florine Hartwig, Josua Sassen, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen. 2023. An Elastic Basis for Spectral Shape Correspondence. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH ’23 Conference Proceedings)*, August 6–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3588432.3591518>

1 INTRODUCTION

Shape correspondence is an important and challenging task in geometry processing and shape analysis, with applications to shape deformation, texture transfer, and shape collection analysis, to name a few. When formulated directly in terms of the input shapes’ geometry, the resulting optimization problems are often computationally intractable. Alternatively, Ovsjanikov et al. [2012] proposed to formulate the matching problem in terms of *function spaces* on the

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGGRAPH ’23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0159-7/23/08.
<https://doi.org/10.1145/3588432.3591518>

shapes, known as the *functional map approach*. In this setting, finding a correspondence translates to finding a *change of basis* between function spaces on the source and target shapes.

Most methods based on functional maps (FM) use the eigenfunctions of the Laplace-Beltrami (LB) operator as the (reduced) functional basis. While highly successful on many datasets [Ovsjanikov et al. 2016], this is, however, mostly effective when the shapes are close to isometric, as the LB operator is invariant under isometric deformations. If, on the other hand, the mapping requires a different objective, e.g., matching *extrinsic* features of the shapes, then the LB eigenvectors are less appropriate. Furthermore, since FM methods use almost exclusively the *orthogonal* LB basis, all the available mathematical machinery and derived methods are built on the orthogonality assumption. This is unfortunate, as it limits the development of alternative basis functions.

To bridge this gap and leverage the functional map approach for crease-driven matching, we propose a novel functional basis derived from the *elastic eigenmodes*, cf. Fig. 1 (left). These are the eigenvectors of the Hessian of an *elastic energy*. Since this basis is not orthonormal, we generalize the functional map approach—including the extraction of a point-to-point map—to this setting. We show that computing a correspondence using our elastic basis leads to a better crease alignment—see Fig. 2 (right)—and a more accurate correspondence compared to using the LB basis. In addition, our approach serves as a stepping stone to using a non-orthogonal basis in the functional maps pipeline, thus enabling the future development of a wide variety of basis functions.

1.1 Related Work

An extensive literature review on the problem of shape correspondence is beyond our scope, please see the recent review by Sahillioglu [2020]. We focus on methods that use functional maps and methods that use *extrinsic information*.

Functional map methods. Following the introduction of functional maps (FM) [Ovsjanikov et al. 2012], a plethora of methods leveraging this approach have been proposed. These added to the FM setup various regularizers [Magnet et al. 2022], incorporated FM in a genetic optimization framework [Edelstein et al. 2020], computed an adapted basis [Azencot and Lai 2021], and improved upon the pointwise map extraction [Ezuz and Ben-Chen 2017; Ren et al. 2021], to mention just a few extensions. Further can be found, for example, in [Ovsjanikov et al. 2016]. Currently, one of the most prominent end-to-end approaches is ZoomOut [Melzi et al. 2019b], that iterates between solving for the functional map, and solving for the point-to-point map, while increasing the size of the functional basis. We also leverage this idea, and formulate the approach for the general case of a non-orthogonal basis.

FM with extrinsic data. The vast majority of FM methods use the eigenvectors of the LB operator as the functional basis, and are therefore *fully intrinsic*. A few methods incorporate extrinsic information *in addition* to the LB basis. For example, SmoothShells [Eisenberger et al. 2020] uses the product space of the LB basis and the Cartesian coordinates (including normals) of the input mesh. The Functional Remeshing approach [Melzi et al. 2020] adds to the basis

3 functions (denoted *Coordinate Manifold Harmonics*) which represent the error of the coordinates in the LB basis. This, however, limits the correspondence to shapes in very similar poses. Finally, Panine et al. [2022] introduce extrinsic information by using a landmark adapted basis, constructed by solving a Dirichlet-Steklov eigenproblem. We compare with these approaches in the results section and show better performance.

Spectral methods on extrinsic operators. Some methods use the extrinsic information encoded in the Steklov operator [Corman et al. 2017; Wang et al. 2018] or the Hessian of the ARAP energy [Huang et al. 2009]. In these approaches the focus was on shape analysis, and the basis was not used for shape correspondence.

Elastic eigenmodes. Note, that while we use the eigenfunctions of an operator related to the elastic energy, our approach *does not* aim to minimize the elastic energy of the correspondence. This approach has been suggested with point-to-point maps [Ezuz et al. 2019a], and, while very effective, is considerably slower than the FM approach and is sensitive to the quality of the initialization. Previously, elastic eigenfunctions have been used, for example, by Hildebrandt et al. [2010, 2012] for shape analysis, i.e. to construct shape signatures, and in [Hildebrandt et al. 2011] to construct a reduced basis for deformation-based surface modeling.

1.2 Contributions

Our main contributions are:

- A novel *crease-aware* functional basis which is highly effective for non-isometric spectral shape correspondence.
- A generalization of the functional map approach to *non-orthogonal* basis functions. This is an important missing ingredient for the flexible design of basis functions and objectives.
- Improved correspondence results compared to the state-of-the-art functional map approaches.

2 THE ELASTIC FUNCTIONAL BASIS

In what follows, we represent a discrete shell as a triangular mesh $\mathcal{S} = (\mathcal{V}, \mathcal{T})$ with n vertices, the vertex set \mathcal{V} , and the set of triangles \mathcal{T} . Let us denote by $\mathcal{F}(\mathcal{S})$ the space of continuous, piecewise affine functions on \mathcal{S} , which is spanned by hat functions and coefficients in this basis correspond to values at vertices. To simplify notation we identify throughout this paper a function $f \in \mathcal{F}(\mathcal{S})$ with the vector $f \in \mathbb{R}^n$ of nodal values. The L_2 scalar product of two functions $f, g \in \mathcal{F}(\mathcal{S})$ given by $\langle f, g \rangle_M = f^T M g$ turns $\mathcal{F}(\mathcal{S})$ into a Hilbert space with norm $\|f\|_M := \sqrt{f^T M f}$. Here, $M \in \mathbb{R}^{n,n}$ is the diagonal lumped mass matrix and for the Euclidean norm $\|f\|_2 = \|f\|_I$.

Our goal is to construct a new reduced basis for the space, that, unlike the Laplace-Beltrami basis, is not invariant under isometric deformations of the shape \mathcal{S} but instead captures extrinsic geometric features. To this end, we consider the eigenfunctions of the Hessian of an elastic discrete shell deformation energy. Hildebrandt et al. [2010] showed that eigenmodes of a similar energy are useful for shape analysis due to their sensitivity to the extrinsic curvature of the surface. Here, we take into account sets of elastic eigenmodes to construct a functional map approach to compute non-isometric correspondences that align relevant extrinsic shape details.




Figure 2: Correspondences are computed for two spheres with small and larger cut out, respectively (top right), based on normal components of elastic eigenmodes and LB eigenmodes. The elastic basis nicely respects creases whereas the LB basis does not due to its invariance to isometric deformations (left). For a qualitative visualization the functions have been scaled to $[-1, 1]$. Consequently, the functional map method using the elastic basis properly matches these creases whereas the method using LB eigenmodes falls short in that respect aiming for an isometric matching (bottom right). Colors are based on direction to visualize (transferred) normals as in Figure 1.

2.1 Background: elastic energy

The discrete elastic energy underlying our approach is composed of a membrane contribution \mathcal{W}_{mem} measuring the local stretching of the surface and a bending contribution $\mathcal{W}_{\text{bend}}$ measuring the change of (discrete) curvature, i.e.

$$\mathcal{W}_S[\phi] = \mathcal{W}_{\text{mem}}[\phi] + \mathcal{W}_{\text{bend}}[\phi],$$

for the elastic energy depending on a deformation $\phi \in (\mathcal{F}(S))^3$.

To compute the membrane contribution, one requires the triangle-wise constant first fundamental form $G[S]$ of S . The membrane energy density depends on the Cauchy–Green strain tensor $\mathcal{G}[\phi] := (G[S])^{-1} G[\phi(S)]$, which encodes the change of metric under the deformation. Then, the membrane energy is defined via integration

$$\mathcal{W}_{\text{mem}}[\phi] := \delta \sum_{t \in \mathcal{T}} a_t W_{\text{mem}}(\mathcal{G}[\phi]|_t),$$

where δ is the thickness of the thin shell, and a_t is the area of t . We use the rigid body motion invariant neo-Hookean energy density $W_{\text{mem}}(A) := \frac{\mu}{2} \text{tr}(A) + \frac{\lambda}{4} \det A - \left(\mu + \frac{\lambda}{2}\right) \log \det A - \mu - \frac{\lambda}{4}$ as proposed by Wirth et al. [2011]. The logarithmic term guarantees triangles remain valid for finite-energy deformations. The Hessian of this energy at the identity deformation is the quadratic form of the Navier–Lamé model of linearized elasticity for tangential displacements with corresponding material constants μ and λ .

To define the discrete bending energy of a deformation ϕ of S , we evaluate the difference of the dihedral angles θ in the undeformed and θ^ϕ in the deformed configuration as a discrete measure of the change of curvature, where the dihedral angle on edges is the angle between normals of neighboring triangles. Then, the discrete bending energy [Grinspun et al. 2003] reads as

$$\mathcal{W}_{\text{bend}}[\phi] = \delta^3 \sum_{e \in \mathcal{E}} a_e^{-1} l_e^2 (\theta_e - \theta_e^\phi)^2.$$

Here, l_e is the length of the an edge e , and $a_e := \frac{1}{3}(a_t + a_{t'})$ for the two triangles t and t' adjacent to $e \in \mathcal{E}$.

2.2 Elastic eigenmodes

To compute the vibration modes of a shape S , we consider the Hessian $\text{Hess } \mathcal{W}_S[\text{Id}] \in \mathbb{R}^{3n, 3n}$ of the total elastic energy evaluated at the identity. This is a linearization operator acting on (infinitesimal) displacements $\psi \in (\mathcal{F}(S))^3$ of the shape S identified with vectors in \mathbb{R}^{3n} . The Hessian is symmetric semi positive-definite with a six-dimensional kernel reflecting the infinitesimal rigid body motions, cf. [Heeren et al. 2014]. The eigenmodes are solutions of the generalized eigenfunction problem

$$\text{Hess } \mathcal{W}_S[\text{Id}]\psi_i = \lambda_i \bar{M}\psi_i$$

where we denote by $\bar{M} \in \mathbb{R}^{3n, 3n}$ a block diagonal matrix with stacked lumped mass matrices on the diagonal. We only consider ψ_i with $\lambda_i \neq 0$.

By construction, the elastic energy of the identity $\mathcal{W}_S[\text{Id}]$ vanishes and the identity is a minimizer which implies $D\mathcal{W}_S[\text{Id}] = 0$. Hence, the eigenmodes of the Hessian with increasing eigenvalues represent deformations ordered with respect to their induced elastic energy. For example, on humanoid shapes, the first eigenmodes are usually supported on the arms and legs and represent close to isometric deformations of the limbs as a whole, see Figure 3. More generally, we observe that the eigenmodes are well suited to detect extrinsic geometric features such as crease lines, e.g., in Figure 2.

The eigenmodes appear to combine—not necessarily additively—eigenmodes of the membrane and bending energy. This combination is controlled by the thickness δ of the elastic shell. For smaller δ , the bending energy has an increasing influence on eigenmodes with small eigenvalue and thus these eigenmodes are more sensitive towards curvature features. In all our applications, $\delta = 10^{-2}$ appeared to be a suitable choice with respect to this sensitivity for shapes of approximately unit area. This leads to spectra of approximately the same scale for the Hessian of the membrane and bending energy.

In principal, it is possible to use the vector-valued eigenmodes to compute a functional map $C_{12}: (\mathcal{F}(S_2))^3 \rightarrow (\mathcal{F}(S_1))^3$. However, this increases the degrees of freedom compared to functional maps for scalar-valued functions and in addition requires handling local rotations. To circumvent both of these shortcomings, we project each eigenmode onto the uniform weighted vertex-wise normals of the surface resulting in $3n$ (linear dependent) functions. The particular choice of normals has no significant impact on our output. From these projected functions, we practically select the first k functions as a basis for the functional map approach. In Section 5.4.1, we demonstrate that the reduced basis of projected eigenmodes leads to a more accurate representation of a ground-truth map, compared to the vector-valued eigenmodes.

It is important to notice that the resulting scalar functions are no longer orthogonal in $\mathcal{F}(S)$. This led us to formulate the functional map in a generalized way allowing for non-orthogonal functions with potentially additional applications than those considered here (cf. Section 6).

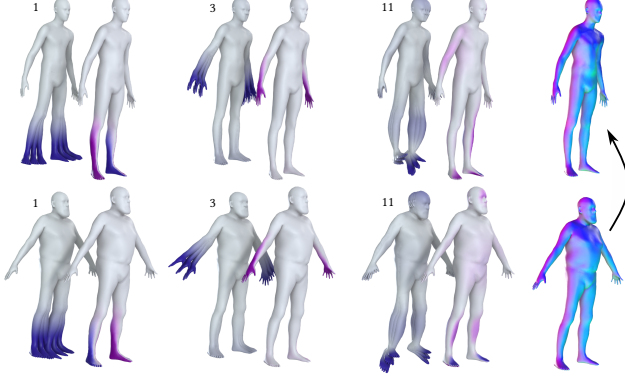


Figure 3: Three elastic eigenmodes for small eigenvalues (order given in Figure) on two human shapes (pairwise: vibration mode (left) and projection on normals (right)) and computed correspondence visualized by normal pullback from bottom to top row. Color visualization as in Figure 2.

3 FUNCTIONAL MAP OPERATORS WITH A NON-ORTHOGONAL BASIS

In this section, we formulate the main operators used in functional map approaches for non-orthogonal basis functions. This requires to carefully incorporate the proper mass-weighted scalar products and norms, on functions and on operators, as well as the associated dual calculus for the adjoint operators.

3.1 The classical functional map setup

We arrange a generic—not necessarily orthogonal—basis of a k -dimensional subspace of $\mathcal{F}(S)$ as columns of a matrix

$$\Phi_k = [\phi_1, \dots, \phi_k] \in \mathbb{R}^{n,k}$$

and denote the subspace spanned by the columns of Φ_k as $\langle \Phi_k \rangle$. The source and target shape are distinguished by subscripts S_1, S_2 with respective bases $\Phi_{1,k}$ and $\Phi_{2,k}$. A vertex map $\Pi_{12}: \mathcal{V}_1 \rightarrow \mathcal{V}_2$ is represented as a column stochastic binary matrix. Namely, $\Pi_{12} \in \Pi := \{\Pi \in \{0, 1\}^{m,n} | \Pi^T \mathbf{1}_m = \mathbf{1}_n\}$, where $\mathbf{1}_n \in \mathbb{R}^n$ respectively $\mathbf{1}_m \in \mathbb{R}^m$ are constant one vectors, and $n = |\mathcal{V}_1|, m = |\mathcal{V}_2|$. The pullback of Π_{12} to the function spaces is denoted by $P_{12}: \mathcal{F}(S_2) \rightarrow \mathcal{F}(S_1)$ where $(P_{12}g)^T e_i = g^T \Pi_{12} e_i$ for $g \in \mathcal{F}(S_2)$ and e_i the i -th canonical basis vector in \mathbb{R}^n , hence $P_{12} = \Pi_{12}^T$. A functional map $C_{12}: \mathbb{R}^k \rightarrow \mathbb{R}^k$ is a rank- k -approximation of P_{12} .

3.2 Orthogonal subspace projection

Reduced mass matrix. For functions f, g in $\langle \Phi_k \rangle$, where $x, y \in \mathbb{R}^k$ are the corresponding basis coefficients, namely $f = \Phi_k x$ and $g = \Phi_k y$, we obtain $\langle f, g \rangle_M = x^T M_k y$. Here,

$$M_k = \Phi_k^T M \Phi_k \in \mathbb{R}^{k,k} \quad (1)$$

is the *mass matrix* with respect to the reduced basis Φ_k , which is *symmetric* but not necessary diagonal.

Orthogonal projector. A fundamental tool that is required when working with a restricted basis, is to compute the *best approximation*

of a function $f \in \mathcal{F}(S)$ in the space $\langle \Phi_k \rangle$. For a given function f , this is given by the *orthogonal projection operator* Φ_k^\dagger defined by

$$\Phi_k^\dagger f = \operatorname{argmin}_{x \in \mathbb{R}^k} \|\Phi_k x - f\|_M^2 \quad \forall f \in \mathcal{F}(S).$$

Using the linearity of Φ_k^\dagger , we can condense this to

$$\Phi_k^\dagger = \operatorname{argmin}_{X \in \mathbb{R}^{k,n}} \|\Phi_k X - I\|_M^2 \quad (2)$$

for the weighted Frobenius norm $\|A\|_M = \sqrt{\operatorname{tr}(A^T M A)}$. Computing the first variation for $\|\Phi_k X - I\|_M^2$, we obtain the necessary optimality condition $0 \stackrel{!}{=} 2\Phi_k^T M (\Phi_k X - I)$, leading to

$$\Phi_k^\dagger = \left(\Phi_k^T M \Phi_k\right)^{-1} \Phi_k^T M = M_k^{-1} \Phi_k^T M \in \mathbb{R}^{k,n}, \quad (3)$$

which is also known as the *Moore-Penrose- or pseudo-inverse*. Note that we have $\Phi_k^\dagger \Phi_k = I$, and that for an *orthonormal* basis, this reduces to $\Phi_k^\dagger = \Phi_k^T M$, as expected.

3.3 Functional maps in a non-orthogonal basis

With the orthogonal projector at hand, we compute the rank- k -approximation of P_{12} and obtain the functional map

$$C_{12} = \Phi_{1,k}^\dagger P_{12} \Phi_{2,k}. \quad (4)$$

Using the defining equations

$$\langle C_{12} y, x \rangle_{M_{1,k}} = \langle y, C_{12}^* x \rangle_{M_{2,k}} \quad \text{and} \quad \langle g, P_{12} f \rangle_{M_1} = \langle P_{12}^* g, f \rangle_{M_2} \quad (5)$$

where $f \in \mathcal{F}(S_1), g \in \mathcal{F}(S_2)$ we obtain the adjoint operators

$$C_{12}^* = M_{2,k}^{-1} C_{12}^T M_{1,k}, \quad P_{12}^* = M_2^{-1} P_{12}^T M_1. \quad (6)$$

Note that for an orthonormal basis, the adjoint of the functional map is simply its transpose.

Together with the definition of $\Phi_{1,k}^\dagger$ and $\Phi_{2,k}^\dagger$, we arrive at the following relation of both adjoints

$$\begin{aligned} C_{12}^* &= M_{2,k}^{-1} \Phi_{2,k}^T P_{12}^T (\Phi_{1,k}^\dagger)^T M_{1,k} \\ &= \left(M_{2,k}^{-1} \Phi_{2,k}^T M_2\right) \left(M_2^{-1} P_{12}^T M_1\right) \Phi_{1,k} M_{1,k}^{-1} M_{1,k} = \Phi_{2,k}^\dagger P_{12}^* \Phi_{1,k}. \end{aligned} \quad (7)$$

For a detailed computation see the supplementary material.

3.4 The Hilbert–Schmidt norm

In the derivation of the algorithm, we will make use of an orthogonal splitting which holds true in a proper norm, see Section 4.1. A natural norm on the space of linear operators between finite-dimensional Hilbert spaces is given by the Hilbert–Schmidt (HS) norm which is defined as $\|F\|^2 = \operatorname{tr}(F^* F)$. Hence, in our setting we obtain for the pullback $P_{12}: \mathcal{F}(S_2) \rightarrow \mathcal{F}(S_1)$ of the vertex map

$$\|P_{12}\|^2 = \operatorname{tr}(P_{12}^* P_{12}) = \operatorname{tr}(M_2^{-1} P_{12}^T M_1 P_{12})$$

and for the functional map C_{12}

$$\|C_{12}\|^2 = \operatorname{tr}(M_{2,k}^{-1} C_{12}^T M_{1,k} C_{12}).$$

Note that the HS norm is invariant to multiplication by the basis Φ_k . Specifically, for $Q \in \mathbb{R}^{k,k}$ mapping from $\langle \Phi_k \rangle$ to $\langle \Phi_k \rangle$, we have

$$\begin{aligned} \|\Phi_k Q\| &= \operatorname{tr}((\Phi_k Q)^* \Phi_k Q) = \operatorname{tr}(M_k^{-1} Q^T \Phi_k^T M \Phi_k Q) = \\ &= \operatorname{tr}(M_k^{-1} Q^T M_k Q) = \operatorname{tr}(Q^* Q) = \|Q\|. \end{aligned}$$

In the functional map context, the weighted Frobenius norm $\|F\|_M = \text{tr}(F^T M F)$ is often used for matrices representing operators. However, this formulation neglects the scalar product on the *domain* of the operator. As a motivational toy example, consider two shapes with the same connectivity but different triangle areas where the correspondence is given by the identity. Then $\|P_{12}\|^2 = \text{tr}(M_2^{-1} M_1)$ provides a useful measure for the change of geometry, whereas $\|P_{12}\|_{M_2}^2 = \text{tr}(M_2)$ only depends on the target shape's geometry.

3.5 Enforcing regularity of the functional map

To define a suitable regularity of a functional map, it is natural to relate its properties to the properties of the vertex map. Here, we aim for a regularization that works for a general basis choice. We will consider

$$\|C_{12} C_{12}^* - I\| \quad (8)$$

as one component of the objective. Previously, Rustamov et al. [2013] referred to $C_{12} C_{12}^*$ as the shape difference operator and observed that its deviation from the identity captures area distortion based disparities between the shapes. Indeed, in the full spectral basis and on continuous surfaces $C_{12} C_{12}^* = I$ relates to area preserving point maps [Rustamov et al. 2013]. In addition, a functional penalizing the discrepancy of $C_{12} C_{12}^*$ from the identity promotes the invertibility of the functional map.

For the LB basis this term simplifies to the commonly used orthonormality promoting term $C_{12} C_{12}^T = I$. Melzi et al. [2019b] showed that minimizing this term incrementally in k in the LB basis translates to the computation of isometric point maps. However, this uses explicit properties of the LB basis and does not hold for a general basis. In particular, in our approach we are *not* regularizing for isometric vertex maps.

4 THE OPTIMIZATION PROBLEM

To compute a correspondence we consider the general procedure presented in [Melzi et al. 2019b] and follow the pipeline of Ren et al. [2021]. This means that we decouple the minimization of the energy by solving separately for the functional map and the vertex correspondence while spectrally upsampling the basis. However, we generalize the underlying theoretical framework to develop this procedure for an *arbitrary* basis. To this end, we introduce the generalized objective and show how to compute the functional and vertex map in an alternating algorithm. We show that this reduces to the algorithm of Melzi et al. [2019b] in case of orthonormal bases. Hence, it is consistent with previous work.

4.1 The objective

As one part of the objective we consider (8) due to its regularization properties described in Section 3.5. i.e. capturing shape differences and improving invertibility. We replace $C_{12} = \Phi_{1,k}^\dagger P_{12} \Phi_{2,k}$ in (8) and obtain

$$\|\Phi_{1,k}^\dagger P_{12} \Phi_{2,k} C_{12}^* - I\|. \quad (9)$$

This term now depends on the functional map and the pullback of the vertex correspondence. Hence, it can be minimized in an

alternating scheme. It quantifies the lack of commutativity in the green corner of the diagram in the space \mathbb{R}^k

$$\begin{array}{ccc} \mathcal{F}(S_2) & \xrightarrow{P_{12}} & \mathcal{F}(S_1) \\ \Phi_{2,k} \uparrow & & \Phi_{1,k}^\dagger \downarrow \\ \langle \Phi_{2,k} \rangle \cong \mathbb{R}^k & \xleftarrow{C_{12}^*} & \mathbb{R}^k \cong \langle \Phi_{1,k} \rangle \end{array} \quad (10)$$

Hence, minimizing it with respect to P_{12} is not controlling the image of P_{12} lying outside $\langle \Phi_{1,k} \rangle$. Thus, we aim at minimizing the objective

$$\mathcal{J}(C_{12}, P_{12}) = \|P_{12} \Phi_{2,k} C_{12}^* - \Phi_{1,k}\|^2 \quad (11)$$

which is the commutativity gap measured in the red corner of the diagram in the larger space $\mathcal{F}(S_1)$. This splits orthogonally into

$$\begin{aligned} \|P_{12} \Phi_{2,k} C_{12}^* - \Phi_{1,k}\|^2 &= \|\Phi_{1,k} \Phi_{1,k}^\dagger (P_{12} \Phi_{2,k} C_{12}^* - \Phi_{1,k})\|^2 \\ &+ \left\| \left(I - \Phi_{1,k} \Phi_{1,k}^\dagger \right) (P_{12} \Phi_{2,k} C_{12}^* - \Phi_{1,k}) \right\|^2. \end{aligned}$$

The first term coincides with (9), using that the HS norm is invariant under multiplication by the basis $\Phi_{1,k}$ and $\Phi_{1,k}^\dagger \Phi_{1,k} = I$. The second term is the corresponding norm on the orthogonal complement of $\langle \Phi_{1,k} \rangle$. Hence, it incorporates the desired control of the full image of P_{12} . The splitting follows from

LEMMA 4.1. *Let $F \in \mathbb{R}^{n,l}$ be a linear operator between two finite dimensional Hilbert Spaces, then*

$$\|F\|^2 = \|\Phi_k \Phi_k^\dagger F\|^2 + \|(I - \Phi_k \Phi_k^\dagger) F\|^2.$$

For the proof we refer to Lemma 1 a) in the supplementary material. A similar splitting idea was originally proposed by Ezuz and Ben-Chen [2017] and since then used in several methods, e.g., [Melzi et al. 2019b; Ren et al. 2021]. In (8) we could alternatively replace the adjoint C_{12}^* by $\Phi_{2,k}^\dagger P_{12}^* \Phi_{1,k}$, see Section 4.3.

4.2 Computing the vertex map

The resulting norm (11) is minimized row-wise to optimize for the vertex map $\Pi_{12} = P_{12}^T$, which is enabled by the following

LEMMA 4.2. *The minimization of (11) in terms of Π_{12}^T is row separable and its minimizer for all $j = 1, \dots, n$ is given by $\Pi_{12}(i, j) = 1$ iff*

$$i = \underset{r \in \{1, \dots, m\}}{\text{argmin}} \left| M_{1,k}^{-1} \left((\Phi_{2,k} C_{12}^*)^T e_r - \Phi_{1,k}^T e_j \right) \right|_{M_{1,k}}^2. \quad (12)$$

Similar results concerning the row separability and using weighted Frobenius norms have been proven for example in [Panine et al. 2022]. For the proof we refer to Lemma 2 a) in the supplementary material. A key point of the argument is that the mass matrices on $\mathcal{F}(S_i)$ are lumped and thus diagonal, whereas the mass matrices on $\langle \Phi_{i,k} \rangle$ might not be diagonal. Using

$$(\Phi_{2,k} C_{12}^*)^T = M_{1,k} C_{12} M_{2,k}^{-1} \Phi_{2,k}^T = M_{1,k} C_{12} \Phi_{2,k}^\dagger M_2^{-1}$$

and $(\Phi_{1,k})^\top = M_{1,k}\Phi_{1,k}^\dagger M_1$, the minimization in (12) can be reformulated in the euclidean norm to

$$i = \operatorname{argmin}_{r \in \{1, \dots, m\}} \left| \sqrt{M_{1,k}} \left(C_{12}\Phi_{2,k}^\dagger M_2^{-1}e_r - \Phi_{1,k}^\dagger M_1^{-1}e_j \right) \right|_2^2 \quad (13)$$

For an orthonormal basis—using $\Phi_{i,k}^\dagger M_i^{-1} = \Phi_{i,k}^\top$ and $M_{1,k} = I$ —the right-hand side of (13) reduces to

$$\operatorname{argmin}_{r \in \{1, \dots, m\}} \left| C_{12}\Phi_{2,k}^\top e_r - \Phi_{1,k}^\top e_j \right|_2^2 \quad (14)$$

which is the same minimization as in [Melzi et al. 2019b].

4.3 Dual perspective

In a dual ansatz, we replace in $\|C_{12}C_{12}^* - I\|$ the adjoint C_{12}^* by $\Phi_{2,k}^\dagger P_{12}^* \Phi_{1,k}$. This leads to the *same minimization problem* demonstrating the consistency of our framework. Indeed, we now have

$$\begin{aligned} \left\| C_{12}\Phi_{2,k}^\dagger P_{12}^* - \Phi_{1,k}^\dagger \right\|^2 &= \left\| C_{12}\Phi_{2,k}^\dagger P_{12}^* \Phi_{1,k} - I \right\|^2 \\ &\quad + \left\| C_{12}\Phi_{2,k}^\dagger P_{12}^* \left(I - \Phi_{1,k}\Phi_{1,k}^\dagger \right) \right\|^2 \end{aligned}$$

analogously to the splitting in (11), see Lemma 1 b) in the supplementary material. Let us emphasize that this splitting holds *only* for the Hilbert–Schmidt norm. Finally, for this dual approach, we obtain that the minimization in terms of Π_{12} is column separable and leads to the row-wise minimization for the vertex map described in (13), see Lemma 2 b) in the supplementary material.

4.4 The alternating algorithm

Finally, we have all the ingredients in place to formulate the alternating Algorithm 1 to compute correspondences using our elastic basis. As mentioned, we follow the idea originally proposed by Melzi et al. [2019b] to combine this alternating algorithm with a spectral upsampling, i.e. we increase the number of basis functions considered in each iteration. See Algorithm 1 for the details. This approach has been proven to be an efficient method for computing functional maps and high quality correspondences in recent publications, e.g. [Ren et al. 2021]. It is also possible to solve for rectangular functional maps but we did not observe any benefits in our experiments.

The procedure can be initialized by a small functional map derived, for example, by visual alignment of the first basis functions or by matching the basis functions based on the values at a few given landmark points. Alternatively, one can also start with a noisy vertex map. For more details, see Section 5. We demonstrate an example result in Figure 4 showing the initial functional map (small matrices) and the final correspondence visualized using texture transfer. Note, that for this non-isometric deformation of the models, ZoomOut creates visible artifacts in the map (see e.g. the wings and the tail), whereas our map nicely aligns the two shapes.

4.5 Limitations

As our basis is constructed via spectral analysis of the elastic energy, which is more sensitive to bad meshing, the input meshes should be relatively regular, which is easily addressed through remeshing if required. In addition, our method assumes similar features (identifiable by the elastic energy) exist in the source and target

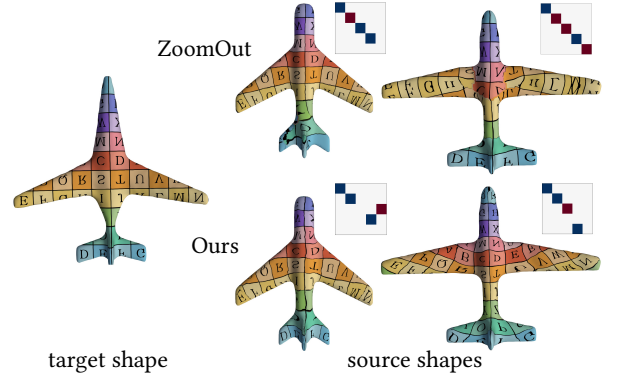


Figure 4: Comparison of ZoomOut and our method on shapes from SHREC'07 [Giorgi et al. 2007]. Initialized by $C_{12,5} \in \mathbb{R}^{5,5}$ created through the visual alignment of basis functions (small matrices) and upsampled to $k_{\max} = 100$.

ALGORITHM 1:

Data: $C_{12, k_{\min}} \in \mathbb{R}^{k_{\min}, k_{\min}}$
for $k = k_{\min} + 1, \dots, k_{\max}$ **do**
 Compute the vertex map $\Pi_{12,k}^\top = \operatorname{argmin}_{P_{12}} \mathcal{J}(C, P_{12})$ by setting
 $\Pi_{12,k}(i, j) = 1$ if

$$i = \operatorname{argmin}_{r \in \{1, \dots, m\}} \left| \sqrt{M_{1,k}} \left(C_{12,k}\Phi_{2,k}^\dagger M_2^{-1}e_r - \Phi_{1,k}^\dagger M_1^{-1}e_j \right) \right|_2^2$$

 for all $j = 1, \dots, n$.
 Compute the functional map

$$C_{12,k} = \Phi_{1,k}^\dagger P_{12,k} \Phi_{2,k}$$
 with $P_{12,k} = \Pi_{12,k}^\top$
end

shape. Note that this is a weaker requirement than isometry, and indeed one would expect at least some similar large scale structures between shapes that are to be put in correspondence.

5 RESULTS

We evaluate our method by comparing it to state-of-the-art approaches. We compare to ZoomOut [Melzi et al. 2019b], which is closest to our approach, and to recent methods for non-isometric correspondence: RHM [Ezuz et al. 2019b], SmoothShells [Eisenberger et al. 2020], and Landmark-Adapted Bases (LAB) [Panine et al. 2022].

If a ground-truth correspondence exists we use the protocol described in [Kim et al. 2011] to evaluate the geodesic error of the correspondence, and the errors are normalized by the square root of the surface area. For qualitative results, we transfer normals using the computed map, and visualize them using a colormap on the unit sphere. Specifically, we show on the source shape the pullback through the computed map of the surface normals of the target shape. This visualization has been described, for example, by Melzi et al. [2019a]. It shows more details of the correspondence than the commonly used colormaps, which are often based on the embedding coordinates. In particular, the normal transfer demonstrates the matching of crease lines. In order to visualize the maps

we compute a vertex-to-point map using the approach described in [Ezuz and Ben-Chen 2017] but adapted to our map conversion (implemented using PyFM [Magnet 2022]). See runtime values in the supplementary. We use $\delta = 10^{-2}$ for all experiments, unless noted otherwise.

5.1 Aligning extrinsic features

Our method is especially effective for aligning extrinsic features. To demonstrate this, we use Mesh Caricaturization [Sela et al. 2015] to deform a source mesh to a non-isometric target mesh with exaggerated features. The deformation defines a ground-truth correspondence, which we use to compute geodesic errors. We remesh the target mesh to avoid any bias from the shared mesh connectivity.

We initialize our method by computing a functional map $C_{12,5} \in \mathbb{R}^{5,5}$ naively aligning our basis functions on five landmarks. This simple initialization is computed using

$$C_{12,k} = \operatorname{argmin}_{C \in \mathbb{R}^{k,k}} \|C\Phi_{2,k}[l_2] - \Phi_{1,k}[l_1]\|^2, \quad (15)$$

where we denote by $l_1, l_2 \in \mathbb{R}^d$ vectors containing the landmark indices and by $\Phi_{1,k}[l_1], \Phi_{2,k}[l_2] \in \mathbb{R}^{k,d}$ the evaluation of the basis functions at the respective landmarks. We use the same approach for initializing ZoomOut, using the LB basis functions instead of ours. For the initialization of RHM and LAB we use the same five landmarks. For all the methods that use eigenfunctions, we take $k_{\max} = 100$. Here, we used $\delta = 10^{-2.5}$ for the bending weight.

Figure 5 shows the comparisons. As the deformations are small all the methods achieve relatively low errors. Still, our method outperforms the others. Note in the qualitative results that our method correctly aligns the features and creases (e.g., the nose and chin of the MaxPlank model, and the mouth and hands of the Homer model), whereas the other methods have a visible “slide” of features. Unlike LAB, that uses the landmarks during the correspondence, we only use the landmarks for *initialization*. Still, our method aligns the features more accurately.

5.2 Accuracy on thin structures

Thin parts of shapes, such as legs, arms, or tails, can be deformed with low elastic energy. Hence, they are already well represented by the first elastic eigenfunctions. This is in stark contrast to the LB eigenfunctions, where a large number of basis functions is required to reconstruct thin structures. We demonstrate this in Figure 6.

Here we compute an LB functional map $C_{12,20} \in \mathbb{R}^{20,20}$ with a standard approach using WKS descriptors, 5 landmarks, and promoting commutativity with the LB operator. We convert it to a vertex map and then upsample from $k_{\min} = 35$ to $k_{\max} = 200$ with our method and ZoomOut. We use the same 5 landmarks for all methods except for SmoothShells which is automatic. Note that our method is the only one that accurately maps the legs and tail. See also Figure 8.

5.3 Comparisons

For the following datasets we use $k_{\max} = 150$ for our approach, for ZoomOut, for LAB, and for SmoothShells. In addition, for SmoothShells we use the provided parameter set which yielded the best results.

5.3.1 FAUST Dataset. We use a remeshed version of the FAUST Dataset [Bogo et al. 2014] (the BCICP dataset [Ren et al. 2018], created using the method of Yan et al. [2014]). The dataset has ten human subjects with different body types, each in ten different poses. We compute new ground-truth correspondences by projecting all the inputs to the high resolution template meshes from the original FAUST dataset. For initializing our method and ZoomOut we use the provided initial maps computed with WKS descriptors and an orientation preserving term [Ren et al. 2018], for generating an initial functional map with $k_{\min} = 20$. For LAB and RHM we choose five landmarks. For reference we also include results for BIM [Kim et al. 2011] given in the dataset. We compute the geodesic error of all the methods over 300 test pairs (the list is included in the dataset), including both isometric and non-isometric matchings.

We plot the percentage of vertex-wise correspondences below a geodesic error threshold in Figure 7 (left). Our method is comparable to SmoothShells and is more accurate than the other methods. While the commonly used vertex-wise percentage plot gives a good overview of the general performance over *all* the test pairs it cannot directly imply an error measure *per test pair*. Therefore, we additionally compute the *median error per test pair*, and show the percentage of *shape pairs* below this error, see Figure 7 (right). Interestingly, we observe here a different behavior of the methods. Our method has a very low median geodesic error in 95% of the test pairs and a rather high one on the other 5%. SmoothShells shows similar results. RHM and LAB, on the other hand, have fewer “failure cases” but only achieve a very high accuracy on a small number of test pairs.

5.3.2 SHREC'20 Dataset. We additionally evaluate our method on the test-sets 1-4 of the SHREC'20 lores dataset [Dyke et al. 2020]. This dataset has different four-legged animals and the twenty test pairs include very challenging shape pairs. We create an initial map for our method and ZoomOut by converting a LB functional map of size $k_{\min} = 35$ (computed using eight landmarks points) to a vertex map. We use the same eight landmarks for LAB. The qualitative and quantitative results are shown in Figure 8. For quantitative evaluation we use the sparse ground-truth (50 points) from the dataset. All the methods lead to comparable accuracy quantitatively. For shapes which have corresponding features we observe qualitatively better results for normal transfer with our method. See, for example, the legs and head of the cow and the giraffe. In the supplementary we provide an additional colormap visualization of the results and show the correspondence with median error of our method.

5.4 Ablation study

5.4.1 Vector-valued vs projected eigenmodes. In our approach we project the vector-valued eigenmodes $\psi_i \in (\mathcal{F}(\mathcal{S}))^3$ on the normal components and construct a (possibly non-orthogonal) basis for $\mathcal{F}(\mathcal{S})$. Indeed, using the vector-valued functions leads to problems with local rotations of the shape as shown in Figure 9. Here, we show the resulting correspondence after converting a ground-truth map between two SCAPE shapes [Anguelov et al. 2005] to a functional map with $k = 100$ basis functions, and reconstructing a vertex map using our map conversion step. We accurately reconstruct the ground-truth map using the projected basis functions, whereas the vector-valued functions lead to large errors on the rotated arms.

5.4.2 Effectiveness of the regularization for vertex map conversion.

To validate the efficiency of the regularizer $C_{12} C_{12}^* \approx I$ we replace our vertex map conversion with one that does not promote this regularity term. For this, we consider

$$i = \operatorname{argmin}_{r \in \{1, \dots, m\}} \left\| \sqrt{M_{2,k}} \left(C_{12,k}^* \Phi_{1,k}^\dagger M_1^{-1} e_j - \Phi_{2,k}^\dagger M_2^{-1} e_r \right) \right\|_2^2 \quad (16)$$

instead of (13). This objective has a similar structure as the map conversion that we use (Eq. (13)) but is formulated using the *adjoint* functional map. As described, e.g., in the work of Pai et al. [2021], this map conversion leverages the fact that the adjoint maps Dirac distributions on the source to Dirac distributions on the target.

We compare the vertex map conversion in two settings: (1) We project a ground-truth vertex map to the reduced basis, and then reconstruct it with and without regularization (Eq. (13) and Eq. (16), respectively), for varying basis size k (Figure 10 (left)). (2) We project a ground-truth vertex map to a reduced basis $k_{\min} = 10$, and then iteratively upsample it to k_{\max} (Figure 10 (right)).

Similarly to the observations of Pai et al. [2021], we notice that while the unregularized reconstruction is well suited to transform a *high dimensional* functional map into a vertex map (Figure 10 (left)), it is not effective in an *upsampling* procedure (Figure 10 (right)). With the regularization, reconstruction in an upsampling framework works well, and the elastic basis outperforms the LB basis in reconstruction error. Explicitly, note that using our approach to upsample a ground-truth representation with 10 elastic basis functions to $k_{\max} = 150$ is nearly as accurate as directly representing the ground-truth with 150 elastic basis functions and reconstructing. As this is the highest accuracy that can be achieved with 150 elastic eigenfunctions this validates the efficiency of our regularization.

6 CONCLUSIONS AND FUTURE WORK

We proposed a crease-aware spectral basis, that is derived from the eigenmodes of the elastic energy. We additionally showed how to incorporate it in the functional map framework, by generalizing it to handle non-orthogonal basis functions. We have used this basis within the iterative upscaling (i.e., ZoomOut) framework, and it would be interesting to combine our approach with other spectral matching techniques that have been proposed in recent years. Additionally, as we have shown how to handle non-orthogonal basis functions, it is interesting to explore other operators, as well as combining different types of basis functions e.g. resulting from machine learning methods as for example in [Marin et al. 2020].

ACKNOWLEDGMENTS

This work was partially supported by the Deutsche Forschungsgemeinschaft (German Research Foundation) via project 211504053 – Collaborative Research Center 1060, Germany’s Excellence Strategy project 390685813 – Hausdorff Center for Mathematics, by the Israel Science Foundation (grant No. 1073/21, No. 668/21 and an equipment grant) as well as the European Research Council (ERC starting grant no. 714776 OPREP), the German-Israeli Foundation for Scientific Research and Development (grant number I-1339-407.6/2016), and by the Israeli Council for Higher Education (CHE) via the Data Science Research Center, Ben-Gurion University of the Negev, Israel.

REFERENCES

- Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. 2005. Shape Completion and Animation of People. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 408.
- Omri Azencot and Rongjie Lai. 2021. A Data-Driven Approach to Functional Map Construction and Bases Pursuit. *Computer Graphics Forum* 40, 5 (aug 2021), 97–110. <https://doi.org/10.1111/cgf.14360>
- Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. 2014. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, 3794–3801.
- Etienne Corman, Justin Solomon, Mirela Ben-Chen, Leonidas Guibas, and Maks Ovsjanikov. 2017. Functional characterization of intrinsic and extrinsic geometry. *ACM Transactions on Graphics (TOG)* 36, 2 (2017), 1–17.
- Roberto M. Dyke, Yu-Kun Lai, Paul L. Rosin, Stefano Zappalà, Seana Dykes, Daoliang Guo, Kun Li, Riccardo Marin, Simone Melzi, and Jingyu Yang. 2020. SHREC’20: Shape correspondence with non-isometric deformations. *Computers & Graphics* 92 (2020), 28–43. <https://doi.org/10.1016/j.cag.2020.08.008>
- Michal Edelstein, Danielle Ezuz, and Mirela Ben-Chen. 2020. ENIGMA. *ACM Transactions on Graphics (TOG)* 39, 4 (jul 2020), 16 pages. <https://doi.org/10.1145/3386569.3392447>
- Marvin Eisenberger, Zorah Lahner, and Daniel Cremers. 2020. Smooth shells: Multi-scale shape registration with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 12265–12274.
- Danielle Ezuz and Mirela Ben-Chen. 2017. Deblurring and denoising of maps between shapes. In *Computer Graphics Forum*, Vol. 36. The Eurographics Association & John Wiley & Sons, Ltd. Chichester, UK, 165–174.
- D. Ezuz, B. Heeren, O. Azencot, M. Rumpf, and M. Ben-Chen. 2019a. Elastic Correspondence between Triangle Meshes. *Computer Graphics Forum* 38, 2 (may 2019), 121–134. <https://doi.org/10.1111/cgf.13624>
- Danielle Ezuz, Justin Solomon, and Mirela Ben-Chen. 2019b. Reversible harmonic maps between discrete surfaces. *ACM Transactions on Graphics (TOG)* 38, 2 (2019), 1–12.
- Daniela Giorgi, Silvia Biasotti, and Laura Paraboschi. 2007. Shape retrieval contest 2007: Watertight models track. *SHREC competition* 8, 7 (2007), 7.
- Eitan Grinspun, Anil N Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*. Citeseer, 62–67.
- Behrend Heeren, Martin Rumpf, Peter Schröder, Max Wardetzky, and Benedikt Wirth. 2014. Exploring the geometry of the space of shells. In *Computer Graphics Forum*, Vol. 33. 247–256.
- Klaus Hildebrandt, Christian Schulz, Christoph von Tycowicz, and Konrad Polthier. 2011. Interactive surface modeling using modal analysis. *ACM Transactions on Graphics (TOG)* 30, 5 (2011), 1–11.
- Klaus Hildebrandt, Christian Schulz, Christoph von Tycowicz, and Konrad Polthier. 2010. Eigenmodes of surface energies for shape analysis. In *International Conference on Geometric Modeling and Processing*. Springer, 296–314.
- Klaus Hildebrandt, Christian Schulz, Christoph von Tycowicz, and Konrad Polthier. 2012. Modal shape analysis beyond Laplacian. *Computer Aided Geometric Design* 29, 5 (2012), 204–218.
- Qi-Xing Huang, Martin Wicke, Bart Adams, and Leonidas Guibas. 2009. Shape decomposition using modal analysis. In *Computer Graphics Forum*, Vol. 28. Wiley Online Library, 407–416.
- Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. 2011. Blended intrinsic maps. *ACM transactions on graphics (TOG)* 30, 4 (2011), 1–12.
- Robin Magnet. 2022. pyFM: Python bindings for functional maps. <https://github.com/RobinMagnet/pyFM>.
- Robin Magnet, Jing Ren, Olga Sorkine-Hornung, and Maks Ovsjanikov. 2022. Smooth Non-Rigid Shape Matching via Effective Dirichlet Energy Optimization. In *International Conference on 3D Vision*.
- Riccardo Marin, Marie-Julie Rakotosaona, Simone Melzi, and Maks Ovsjanikov. 2020. Correspondence learning via linearly-invariant embedding. *Advances in Neural Information Processing Systems* 33 (2020), 1608–1620.
- Simone Melzi, Riccardo Marin, Pietro Musoni, Filippo Bardon, Marco Tarini, and Umberto Castellani. 2020. Intrinsic/extrinsic embedding for functional remeshing of 3D shapes. *Computers & Graphics* 88 (2020), 1–12.
- S Melzi, R Marin, P Musoni, U Castellani, M Tarini, et al. 2019a. Visual assessments of functional maps. In *Symposium on Geometry Processing 2019-Posters*. The Eurographics Association, 5–6.
- Simone Melzi, Jing Ren, Kaust Rodolà, Abhishek Sharma, Maks Ovsjanikov, and Peter Wonka. 2019b. ZoomOut: Spectral Upsampling for Efficient Shape Correspondence. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–14.
- Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–11.
- Maks Ovsjanikov, Etienne Corman, Michael Bronstein, Emanuele Rodolà, Mirela Ben-Chen, Leonidas Guibas, Frederic Chazal, and Alex Bronstein. 2016. Computing

- and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses*. 1–60.
- Gautam Pai, Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. 2021. Fast Sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 384–393.
- Mikhail Panine, Maxime Kirgo, and Maks Ovsjanikov. 2022. Non-Isometric Shape Matching via Functional Maps on Landmark-Adapted Bases. In *Computer Graphics Forum*. Wiley Online Library.
- Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. 2021. Discrete optimization for shape matching. In *Computer Graphics Forum*, Vol. 40. Wiley, 81–96.
- Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. 2018. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–16.
- Raif M Rustamov, Maks Ovsjanikov, Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Leonidas Guibas. 2013. Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–12.
- Yusuf Sahillioglu. 2020. Recent advances in shape correspondence. *The Visual Computer* 36, 8 (2020), 1705–1721.
- Matan Sela, Yonathan Aflalo, and Ron Kimmel. 2015. Computational Caricaturization of Surfaces. *Comput. Vis. Image Underst.* 141, C (dec 2015), 1–17. <https://doi.org/10.1016/j.cviu.2015.05.013>
- Yu Wang, Mirela Ben-Chen, Iosif Polterovich, and Justin Solomon. 2018. Steklov spectral geometry for extrinsic shape analysis. *ACM Transactions on Graphics (TOG)* 38, 1 (2018), 1–21.
- Benedikt Wirth, Leah Bar, Martin Rumpf, and Guillermo Sapiro. 2011. A continuum mechanical approach to geodesics in shape space. *International Journal of Computer Vision* 93, 3 (2011), 293–318.
- Dong-Ming Yan, Guanbo Bao, Xiaopeng Zhang, and Peter Wonka. 2014. Low-Resolution Remeshing Using the Localized Restricted Voronoi Diagram. *IEEE Transactions on Visualization and Computer Graphics* 20, 10 (2014), 1418–1427. <https://doi.org/10.1109/TVCG.2014.2330574>

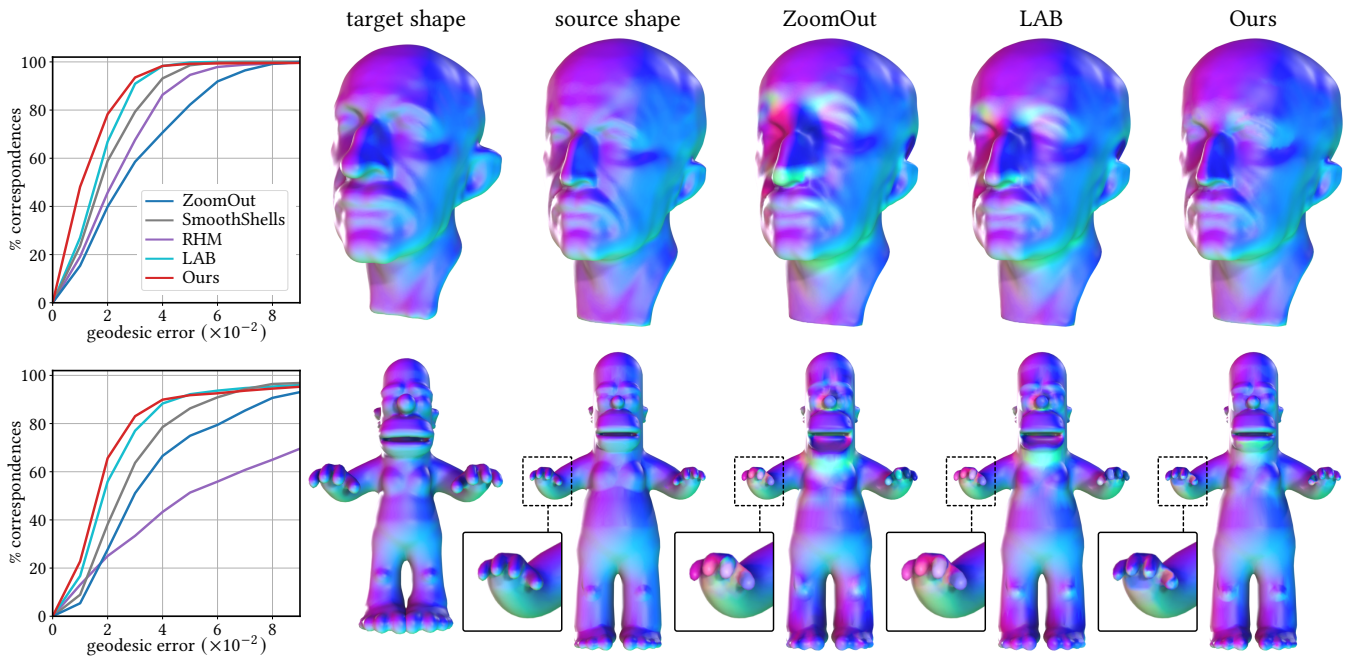


Figure 5: Quantitative and qualitative evaluation of correspondences between two shape pairs created by the caricaturization tool of Sela et al. [2015]. As before, the (transferred) normal direction are visualized by mapping them to colors. Our method accurately aligns extrinsic features preserved by the caricaturization, whereas Laplace–Beltrami-based methods prefer isometric mappings which leads, for example, to wrongly mapped facial features. For qualitative results of the remaining methods see the supplementary material.




Figure 6: Computed correspondence between a cat and lion mesh visualized via normal transfer. Our method is able to generate accurate mappings on thin structures of the shapes—such as the tail—because the elastic basis accentuates them (cf. Figure 1).




Figure 7: Quantitative evaluation on the FAUST dataset. All refinement methods were run with $k_{\min} = 20, k_{\max} = 150$. See Section 5.3.1 for details on the initialization and parameters. Percentage plots showing vertex-wise error (left) and shape pair-wise median error (right).

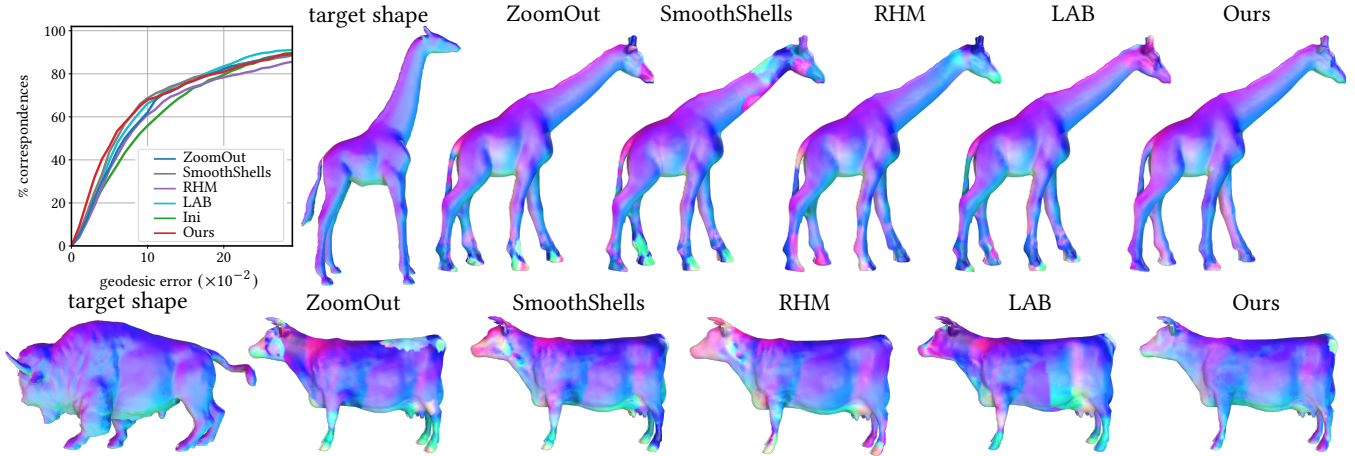


Figure 8: Quantitative and qualitative results for the SHREC'20 low-resolution dataset [Dyke et al. 2020] on 20 correspondences (testsets 1–4). See Section 5.3.2 for more details on the initialization and parameters.

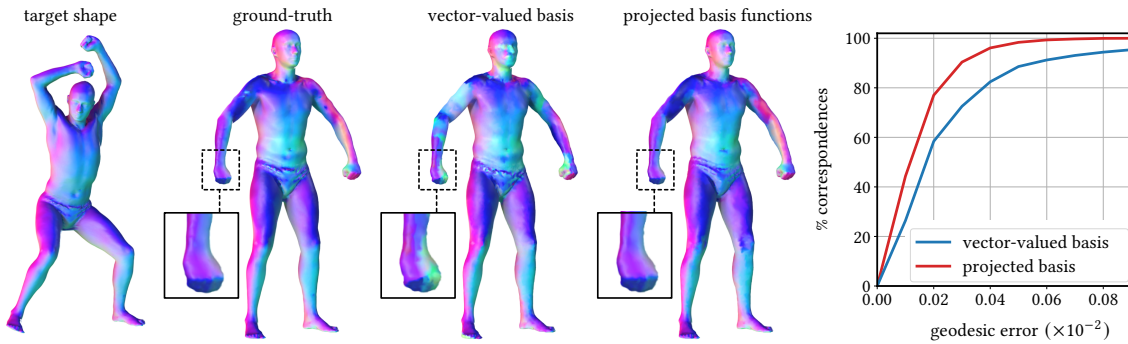


Figure 9: Results for ground-truth reconstruction after converting to a functional map $C_{12,100} \in \mathbb{R}^{100,100}$ computed with the vector-valued eigenmodes and our used projected basis functions. See Section 5.4.1 for details.

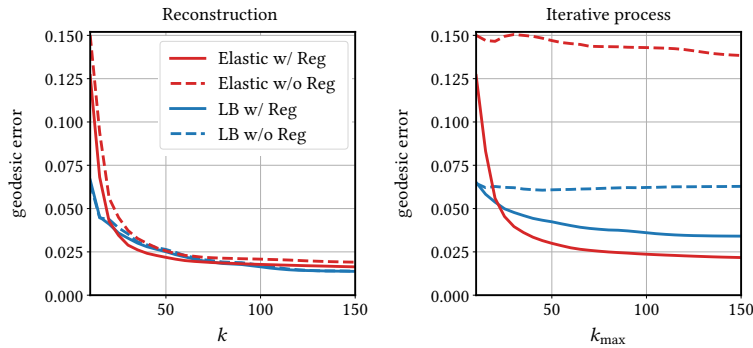


Figure 10: Ablation study on FAUST data set (300 correspondences). Left: mean geodesic error of groundtruth representation with a functional map in the respective basis for $k = 10, \dots, 150$ after map conversion with the regularized and unregularized approach. Right: Iterative process initialized by groundtruth correspondence and upsampled with the regularized resp. unregularized approach. See Section 5.4.2 for details. We show qualitative results for different values of k in the supplementary material.