

A Data-Driven Approach to Functional Map Construction and Bases Pursuit

Omri Azencot¹ and Rongjie Lai²¹Ben-Gurion University, Israel²Rensselaer Polytechnic Institute, New York, USA

Abstract

We propose a method to simultaneously compute scalar basis functions with an associated functional map for a given pair of triangle meshes. Unlike previous techniques that put emphasis on smoothness with respect to the Laplace–Beltrami operator and thus favor low-frequency eigenfunctions, we aim for a basis that allows for better feature matching. This change of perspective introduces many degrees of freedom into the problem allowing to better exploit non-smooth descriptors. To effectively search in this high-dimensional space of solutions, we incorporate into our minimization state-of-the-art regularizers. We solve the resulting highly non-linear and non-convex problem using an iterative scheme via the Alternating Direction Method of Multipliers. At each step, our optimization involves simple to solve linear or Sylvester-type equations. In practice, our method performs well in terms of convergence, and we additionally show that it is similar to a provably convergent problem. We show the advantages of our approach by extensively testing it on multiple datasets in a few applications including shape matching, consistent quadrangulation and scalar function transfer.

CCS Concepts

- **Computing methodologies** → *Shape analysis*;

1. Introduction

Functional maps (FM) [OBCS*12] were introduced in geometry processing in the context of shape matching. During the last few years, FM were quickly adopted by many, serving as the key building block in a range of shape analysis frameworks. Applicative instances include fluid simulation and mesh quadrangulation tasks, in addition to the original shape correspondence problem. The goal of this paper is to propose an efficient and easy-to-code framework for computing data-driven functional map matrices.

The key idea behind FM is that instead of aligning points as in Iterative Closest Point (ICP) approaches [BM92], it is often simpler to align scalar functions defined on the input shapes. Thus, a typical pipeline for computing functional maps is composed of three steps. Given two shapes, one first collects a set of corresponding descriptors, such as the Wave Kernel Signature [ASC11]. Second, one performs dimensionality reduction by projecting the descriptors onto a spanning subspace of basis functions. Finally, one solves an optimization problem, seeking a matrix that best aligns the projected features, possibly while minimizing additional regularizing terms.

Numerous extensions to the original pipeline [OBCS*12] were proposed in the literature. These extensions can be classified into two research avenues. On the one hand, recent works focus on the formulation of novel regularization terms that can be incorporated into the functional map computation phase. For instance,

cycle-consistency is promoted in [HWG14], whereas [NO17] favor the preservation of the given descriptors. On the other hand, other works seek improved functional subspaces. For example, [KBB*13] design basis elements to account for sign or ordering ambiguities. In this context, our work contributes in that it *combines* the tasks of functional map computation and basis design into a single unified framework. Our formulation allows to harness the advancements in functional map regularization as well as to benefit from the increase in the search space of solutions when the bases are allowed to change during the optimization.

Choosing a good basis set is crucial in FM applications. In the original work [OBCS*12], the authors propose to use the eigenfunctions of the Laplace–Beltrami (LB) operator as the spanning subspace in the second step of the pipeline. More generally, over the last few years, LB bases became the prevailing choice for function representation in many geometry processing tasks such as computing descriptors [Rus07], distances [SRGB14], and generating shape segments [RBG*09], just to name a few. While this choice can be optimal under certain conditions [ABB*16], it may sometimes lead to subpar results. To this end, Kovnatsky and others [KBB*13, KBBV15, KGB16, LRBB17] optimize for joint diagonalizable (JD) basis elements to improve shape matching tasks.

In this paper, we address two limitations that appear in many existing work on functional maps. The first shortcoming is related to




Figure 1: We show above some of the modes associated with the Laplace–Beltrami (LB) operator (left) and the Proper Orthogonal Decomposition (POD) as computed for the Wave Kernel Signature [ASC11] and segment information [KO18] (right). While the LB modes encode the intrinsic geometry of the shape, they struggle with representing high frequency signals. In comparison, the POD modes faithfully represent the descriptors’ subspace while implicitly capturing the geometry as encoded in the given features. See also Fig. 2 for a numerical result the support improved feature matching of POD basis elements.

the common choice of the LB function space. While LB encodes the geometry of the surface, it is completely independent of the selected features, potentially introducing large representation errors. Indeed, high frequency signals such as locally supported functions will exhibit poor spectral representations [NMR*18]. Thus, and in contrast to previous work, our approach is based on designing basis elements that are tailored to a given collection of descriptors. In practice, we observe that employing LB-based representations often leads to the elimination of many degrees of freedom that could be re-introduced into the problem. Instead of using LB, we utilize the Proper Orthogonal Decomposition (POD) modes for dimensionality reduction purposes. POD subspaces share many of the advantageous properties of LB—they are orthogonal and have a natural ordering. However, POD modes are able to capture high frequency data and thus improve descriptors’ transfer between shapes.

The second limitation we alleviate deals with the split between the tasks of basis design and functional map computation. Indeed, most existing work focus only on one of these tasks: facilitating a fixed basis or alternatively using a closed-form solution for the functional map. Instead, we merge these objectives into one larger problem. In practice, this modification increases the search space, potentially allowing to find better solutions for a given problem. In addition, we can independently regularize and constrain the bases and the functional map to the specific requirements of the application at hand, leading to a flexible yet effective framework.

This work offers an effective minimization framework for computing functional basis sets on a pair of shapes and a corresponding functional map. The resulting optimization is unfortunately highly non-linear and non-convex. Nevertheless, we construct a novel and efficient Alternating Direction of Multipliers Method (ADMM) scheme. Specifically, we introduce auxiliary variables that allow to incorporate various regularizers promoting e.g., cycle-consistency or metric preservation. Moreover, our scheme converges empirically and we additionally show that our method is similar to a *provably convergent* procedure. We evaluate our approach on shape analysis tasks including shape matching, joint quadrangulation and function transfer. Our comparison indicates

that our method achieves comparable or moderately better results in shape correspondences on challenging scenarios where the shapes do not share the connectivity, or are only approximately isometric.

2. Related Work

Functional maps [OBBS*12] gain a lot of attention in geometry processing and related fields. Some of the main applications of FM include shape exploration [ROA*13], fluid simulation [AWO*14] and function transfer [NMR*18]. We refer the interested reader to a recent course discussing the functional map framework and a few related applications [OCB*16].

One of the main scenarios in which functional maps are employed is for computing shape correspondences between a given pair or collection of shapes. In this context, many works extend the original approach [OBBS*12] to include various regularization terms. For instance, Nogneng and Ovsjanikov [NO17] show that minimizing commutativity with descriptor operators leads to better functional maps. In [HWG14], the authors promote consistency with respect to the inverse mapping, and recently, [RPWO18] formulate orientation-preserving terms into the functional maps pipeline. In addition to improving the accuracy of functional map matrices, the approaches for extracting point-to-point maps are also under development. [RMC15] cast this problem as a probability density function estimation, whereas [EBC17] propose to minimize the error from projecting delta functions onto the basis and its orthogonal complement. Finally, several deep learning techniques were recently developed, see e.g., [MRMO20].

In all of the above works, while the functional map could be computed in any scalar basis, the eigenfunctions of the Laplace–Beltrami operator are typically used. This choice is natural given the wealth of theoretical results related to the LB spectrum, but, on the other hand, it is completely independent of the input descriptors. A related line of work compute Mexican hat wavelets which are able to encode high-frequency information [HQ12, KMP*21]. Recently, [SBL18] proposed to design a basis via a conformal deformation while the other basis set is fixed. Probably closest to our approach is the line of work of [KBB*13, KBBV15, KGB16,

LRBB17] where the authors look for spectral coefficients such that the resulting basis elements are as close as possible to the LB eigenfunctions while preserving the given constraints. In contrast to their perspective, we advocate the use of a linear domain in which the descriptors are better represented, in addition to the incorporation of different regularizers. We provide a detailed comparison between our method and theirs in 7.3.

3. Motivation and Background

To motivate our approach, we will need the following notation. Let $\mathcal{M}_1 = (\mathcal{V}_1, \mathcal{F}_1)$ and $\mathcal{M}_2 = (\mathcal{V}_2, \mathcal{F}_2)$ be a pair of manifold triangle meshes, where $\mathcal{V}_1, \mathcal{V}_2$ are their vertex sets and $\mathcal{F}_1, \mathcal{F}_2$ are their triangle sets. We represent scalar functions using real values on vertices, i.e., $f_1 : \mathcal{V}_1 \rightarrow \mathbb{R}$ is a scalar function on \mathcal{M}_1 , and similarly, $f_2 : \mathcal{V}_2 \rightarrow \mathbb{R}$ is a function on \mathcal{M}_2 . Thus, f_1 and f_2 are real-valued vectors of sizes $|\mathcal{V}_1| = m_1$ and $|\mathcal{V}_2| = m_2$, respectively. We define the inner product on \mathcal{M}_1 to be

$$\langle f_1, g_1 \rangle_{\mathcal{M}_1} := f_1^T G_1 g_1,$$

where $G_1 \in \mathbb{R}^{m_1 \times m_1}$ is the diagonal (lumped) mass matrix of the nodes of \mathcal{M}_1 (see e.g., [BKP*10, Chap. 3]), and similarly, we have $\langle f_2, g_2 \rangle_{\mathcal{M}_2} = f_2^T G_2 g_2$. The input to our problem is a collection of functional constraints $\{f_{1j}\}_{j=1}^n$ and $\{f_{2j}\}_{j=1}^n$ such that f_{1j} and f_{2j} encode the same information but on different meshes, for every j . Finally, we arrange the given constraints in matrices,

$$\tilde{F}_1 = [f_{11} \ f_{12} \ \dots \ f_{1n}] \in \mathbb{R}^{m_1 \times n}, \quad \tilde{F}_2 = [f_{21} \ f_{22} \ \dots \ f_{2n}] \in \mathbb{R}^{m_2 \times n}.$$

In its most simple form, the task of computing functional maps consists of finding a matrix C that aligns the descriptors, i.e.,

$$C B_1^T G_1 \tilde{F}_1 \approx B_2^T G_2 \tilde{F}_2,$$

where B_j is a basis of scalar functions on \mathcal{M}_j for $j = 1, 2$. Typically, C is a moderately sized $k \times k$ matrix with $k < 300$. If we assume that the G_j and \tilde{F}_j matrices are fixed, it is natural to ask whether optimizing for C and for the B_j matrices will yield improved feature matching. We show in Fig. 16 an example of a functional map with its bases obtained in this way (left), leading to a high quality map between non-isometric shapes (right). Solving for the bases and map significantly increases the parameters from k^2 to $k^2 \times k \cdot m_1 \times k \cdot m_2$, resulting in a challenging-to-solve problem as m_j are very large. To deal with this issue, we can consider a subspace of solutions of size $k^2 \times k \cdot r_1 \times k \cdot r_2$, where r_j represent the spectral dimensions of some fixed bases. That is, instead of finding spatial bases, we look for spectral coefficient matrices onto predefined linear subspaces.

In practice, most existing work utilize the subspace spanned by the LB eigenfunctions. In this work, we propose to change this common choice and take subspaces that better fit the given descriptors. There are several approaches in the machine learning community that could be investigated to achieve this objective. In this work, we advocate the utilization of the Proper Orthogonal Decomposition (POD) modes [BHL93], which can be classified as a linear manifold learning method. Given a set of descriptors, the POD can be easily computed using the Singular Value Decomposition, see Sec. 6. One of the main reasons for preferring POD modes over other bases is due to the Karhunen–Loève theorem, stating that




Figure 2: We plot the error distributions of feature matching when LB and POD bases are used. The above results show that designing POD modes is beneficial both in the spectral and spatial domains.

these modes best approximate the input data under many choices of norms [Xiu10].

Therefore, incorporating POD modes instead of the LB spectrum may be considered as a data-driven approach for representing and manipulating signals. We show in Fig. 1 a few modes related to the LB operator (left) and resulting from POD computation (right). One significant difference between these bases is that POD modes allow for higher frequencies when compared to a similar truncation of LB. For example, the LB b_5 is significantly smoother than its related POD b_5 . Moreover, encoding descriptors in a POD subspace induces less information loss in comparison to LB representations in the context of designing bases and functional maps, as we show below.

To quantify the difference between the LB and POD subspaces, we measure the average matching error distributions per mode and per vertex, namely

$$e_1 = \frac{1}{n} \sum_{j=1}^n \left(C B_1^T G_1 f_{1j} - B_2^T G_2 f_{2j} \right)^2,$$

$$e_2 = \frac{1}{n} \sum_{j=1}^n \left(B_2 C B_1^T G_1 f_{1j} - f_{2j} \right)^2,$$

where the squares are taken pointwise, i.e., $e_1 \in \mathbb{R}^k$ and $e_2 \in \mathbb{R}^{m_2}$. In Fig. 2 we compare these errors when the bases are fixed as well as designed. Our results indicate that the fixed POD subspaces are extremely accurate for matching in the spectral domain, but yield the most error spatially. Moreover, LB bases produce poor results when designed, e.g., using joint diagonalization methods [KBB*13] for both error measures. Finally, designed POD modes give the most accurate estimation in the spatial domain and is second best in the spectral regime. We remark that for the POD case these errors naturally depend on the particular descriptors in use. In our applications, we employ a mixture of features such as the Wave Kernel Signature [ASC11] or hand-crafted landmarks.

4. Functional Map and Basis Search (FMBS)

Our main goal is to find basis matrices $B_1 \in \mathbb{R}^{m_1 \times k}$ and $B_2 \in \mathbb{R}^{m_2 \times k}$ and a functional map $C \in \mathbb{R}^{k \times k}$ such that these objects best align

the constraints \tilde{F}_1 and \tilde{F}_2 . To reduce clutter, we scale each \tilde{F}_j by its corresponding G_j and denote $F_j = G_j \tilde{F}_j, j = 1, 2$. Formally, we consider the problem

$$\begin{aligned} & \underset{B_1, B_2, C}{\text{minimize}} && \frac{1}{2} \left\| C B_1^T F_1 - B_2^T F_2 \right\|_F^2 \\ & \text{subject to} && B_1^T G_1 B_1 = I, B_2^T G_2 B_2 = I \end{aligned} \quad (1)$$

where $\|\cdot\|_F$ is the Frobenius norm, and the terms $B_j^T F_j$ can be viewed as projecting the constraints onto the basis matrices. The equality conditions given by $B_j^T G_j B_j = I$ constrain the bases to be orthogonal with respect to the mass matrix. Unfortunately, the minimization problem (1) is highly non-linear and non-convex, and thus practical solvers are challenging to construct. To alleviate these difficulties, we propose in the next section a splitting scheme that is based on the Alternating Direction Method of Multipliers (ADMM) [GM75a, GM75b].

5. An ADMM Approach to FMBS

The basic idea of ADMM depends on splitting the original complex optimization task into simpler subproblems that can be solved efficiently. Under certain conditions on the objective function and constraints, it can be shown that ADMM converges. Therefore, ADMM is often the optimization framework of choice, arguably due to its computational complexity and theoretical guarantees. To allow splitting in our problem (1) above, we introduce the auxiliary variables B'_1 and B'_2 and arrive at the following optimization

$$\begin{aligned} & \underset{B_1, B_2, C}{\text{minimize}} && \mathcal{E}_{\text{fid}}(B_1, B_2, C) \\ & \text{subject to} && B_1^T G_1 B'_1 = I, B_2^T G_2 B'_2 = I, B_1 = B'_1, B_2 = B'_2 \end{aligned} \quad (2)$$

where $\mathcal{E}_{\text{fid}}(B_1, B_2, C) = \frac{1}{2} \left\| C B_1^T F_1 - B_2^T F_2 \right\|_F^2$ is the *data fidelity* term. We stress that while ADMM can be viewed as a standard optimization technique, the choice of auxiliary variables is highly dependent on the problem and there is no general rule for how to “correctly” set these variables. In particular, our unique choice leads to an empirically converging scheme for a large range of parameters, and it further allows for a natural incorporation of novel regularizers (11). Finally, we mention that the auxiliary variables linearize the difficult orthogonality constraints which may lead to non-orthogonal bases in practice. However, this issue can be solved in a post-processing step.

To minimize (2) we facilitate an iterative scheme $k = 0, 1, \dots$ where at each step, the unknowns are updated in an alternating style. Namely, all the variables are kept fixed except for the one which is being updated. In our case, the update order for the primal variables is $(B_1, B_2, B'_1, B'_2, C)$, followed by the update of the dual variables (P_1, P_2, Q'_1, Q'_2) . We note that each of the subproblems is at most quadratic in the unknown, and thus can be solved efficiently. In what follows, we discuss in detail each of the update tasks including their formulation and solution. To shorten the mathematical formulations below, we omit the step k with the understanding that the variables are updated in a sequential fashion as shown in Alg. 1. In addition, we denote by $\mathcal{L}_j(B_j, B'_j, P_j, Q'_j)$ the *scaled Lagrangian* terms, i.e.,

$$\mathcal{L}_j(B_j, B'_j, P_j, Q'_j) = \frac{\rho}{2} \left\| B_j^T G_j B'_j - I + P_j \right\|_F^2 + \frac{\rho}{2} \left\| B_j - B'_j + Q'_j \right\|_{\mathcal{M}_j}^2$$

where $j = 1, 2$, $\|A\|_{\mathcal{M}_j} = \text{trace}(A^T G_j A)$ is an operator norm on the surface \mathcal{M}_j , and $\rho \in \mathbb{R}^+$ is a penalty parameter provided by the user and it may be updated during the optimization.

5.1. Updating the bases, B_1 and B_2

The variable B_1 is being updated first, using the estimations of the other variables from the previous step. Specifically, we have

$$B_1^{k+1} = \underset{B_1}{\text{arg min}} \mathcal{E}_{\text{fid}}(B_1, B_2, C) + \mathcal{L}_1(B_1, B'_1, P_1, Q'_1). \quad (3)$$

Computing the first order optimality conditions of (3) lead to a *Sylvester Equation* of the form

$$\begin{aligned} & F_1 F_1^T B_1 C^T C + \left(\rho G_1 B'_1 B_1^T G_1 + \rho G_1 \right) B_1 = \\ & F_1 F_2^T B_2 C + \rho G_1 B'_1 (I - P_1)^T + \rho G_1 (B'_1 - Q'_1), \end{aligned} \quad (4)$$

which can be efficiently solved with numerical algorithms such as [GNVL79] implemented via e.g., `dlyap` in MATLAB. We emphasize that the dimensionality of Eq. (4) introduces a practical challenge, as it involves dense matrices of size $m_1 \times m_1$. These concerns, along with other implementation aspects, are considered in Section 6.

The update for B_2 is carried after the update of B_1 , but before the other variables. Therefore, we use the estimate of B_1 at step $k+1$, whereas the rest of the variables are taken from the k th step. The minimization takes the following form

$$B_2^{k+1} = \underset{B_2}{\text{arg min}} \mathcal{E}_{\text{fid}}(B_1, B_2, C) + \mathcal{L}_2(B_2, B'_2, P_2, Q'_2). \quad (5)$$

Problem (5) is quadratic in B_2 , and its solution can be computed through the following linear system

$$\begin{aligned} & \left(F_2 F_2^T + \rho G_2 + \rho G_2 B'_2 B_2^T G_2 \right) B_2 = \\ & F_2 F_1^T B_1 C^T + \rho G_2 B'_2 (I - P_2)^T + \rho G_2 (B'_2 - Q'_2). \end{aligned} \quad (6)$$

5.2. Updating the auxiliary variables, B'_1 and B'_2

The minimization problems associated with the unknowns B'_1 and B'_2 are similar. These optimization problems take the form

$$B_j^{k+1} = \underset{B'_j}{\text{arg min}} \frac{\rho}{2} \left\| B_j^T G_j B'_j - I + P_j \right\|_{\mathcal{M}_j}^2 + \frac{\rho}{2} \left\| B_j - B'_j + Q'_j \right\|_{\mathcal{M}_j}^2, \quad (7)$$

for $j = 1, 2$. The solution is given via the linear system

$$\left(\rho G_j + \rho G_j B_j B_j^T G_j \right) B'_j = \rho G_j B_j (I - P_j) + \rho G_j (B_j + Q'_j). \quad (8)$$

5.3. Updating the functional map, C

Given the basis matrices B_1 and B_2 , finding the best functional map that aligns the constraints in a least squares sense has a closed-form solution. Namely, we want to minimize the term $\mathcal{E}_{\text{fid}}(B_1, B_2, C)$ with respect to C , and the solution is given by

$$C^{k+1} = \left(B_2^T F_2 \right) \left(B_1^T F_1 \right)^+, \quad (9)$$

where A^+ is the pseudo-inverse of the matrix A .




Figure 3: Our *fmbs* algorithm shows good empirical behavior for a large range of parameters. We compute the normalized energy and primal/dual residuals per normalized minimization step, across all our FAUST tests. The graphs above show the averaged and standard deviation of the energy (left), primal residual (middle) and dual residual (right), where all exhibit decay.

5.4. Updating the dual variables, P_j and Q'_j

The last step of our scheme is trivial and for $j = 1, 2$, it is given by

$$\begin{aligned} P_j &= P_j + B_j^T G_j B'_j - I, \\ Q'_j &= Q'_j + B_j - B'_j. \end{aligned} \quad (10)$$

We summarize the above steps in pseudocode in Alg. 1. We note that generating \mathcal{O}_1^{-1} is computationally prohibitive as \mathcal{O} is a large and dense matrix. However, we significantly reduce the computation costs by representing B_j in a spectral subspace, as we discuss in Sec. 6.

Algorithm 1 Functional Map and Basis Search (*fmbs*)

- 1: Input matrices $F_1 \in \mathbb{R}^{m_1 \times n}$, $F_2 \in \mathbb{R}^{m_2 \times n}$ and a scalar $\rho \in \mathbb{R}^+$
 - 2: Initialize $B_1 = B'_1 = \text{SVD}(\tilde{F}_1, k)$, $B_2 = B'_2 = \text{SVD}(\tilde{F}_2, k)$
 $P_1 = 0, P_2 = 0, Q'_1 = 0, Q'_2 = 0$
 - 3: **for** $k = 0, 1, 2, \dots$ **do**
 - 4: Solve $B_1 = \text{dlyap}(-\mathcal{O}_1^{-1} A_1, B_1, \mathcal{O}_1^{-1} C_1)$, s.t. \triangleright Eq. (4)
 $A_1 = F_1 F_1^T$
 $B_1 = C^T C$
 $C_1 = F_1 F_2^T B_2 C + \rho G_1 B'_1 (I - P_1)^T + \rho G_1 (B'_1 - Q'_1)$
 $\mathcal{O}_1 = \rho G_1 B'_1 B_1^T G_1 + \rho G_1$
 - 5: Solve $A_2 B_2 = B_2$, s.t. \triangleright Eq.(6)
 $A_2 = F_2 F_2^T + \rho G_2 + \rho G_2 B'_2 B_2^T G_2$
 $B_2 = F_2 F_1^T B_1 C^T + \rho G_2 B'_2 (I - P_2)^T + \rho G_2 (B'_2 - Q'_2)$
 - 6: Update B'_1 by solving Eq. (8) with $j = 1$
 - 7: Update B'_2 by solving Eq. (8) with $j = 2$
 - 8: Set $C = (B_2^T F_2) (B_1^T F_1)^+$ \triangleright Eq. (9)
 - 9: Update the dual variables using Eq. (10)
 - 10: Update ρ following Section 3.4 in [BPC*11]
 - 11: **end for**
-

5.5. Provably convergent FMBS

Most convergence results related to ADMM handle problems with convex objective functions and linear constraints. Recently, [WYZ18] and [GGC18] extended the convergence analysis of ADMM to a significantly larger class of problems including non-convex objective terms and non-linear constraints. In particular, in the latter work, the authors investigate the case where *biaffine* constraints are given, namely, constraints involving two variables which become linear when one variable is kept fixed. For instance, our orthogonality conditions $B_j^T G_j B'_j = I$ are exactly of this form. Moreover, [GGC18] relax the convexity requirements on the objective function and allow to include differentiable terms instead.

In practice, Alg. 1 behaves well and it exhibits energy decrease for many choices of parameters as we show in Fig. 3 and in Sec. 7, however, it does not satisfy the conditions given in [GGC18]. To show convergence, we consider in App. A a different minimization (14) for which we can show the following result.

Proposition 1 Under some mild conditions, problem (14) satisfies all the requirements in [GGC18] and thus its ADMM converges to a constrained stationary point. That is, the sequence of variable updates $\{\mathcal{X}^k, \mathcal{Z}^k\}_{k=0}^\infty$ is bounded and every limit point $(\mathcal{X}^*, \mathcal{Z}^*)$ is a constrained stationary point.

5.6. Regularized FMBS

One of the key aspects of our minimization (2) is that it introduces many degrees of freedom via the unknowns B_1, B_2 and C . While in general it is a positive feature of our approach, the associated optimization requires a significant amount of descriptors n . To relax this dependency, we propose to incorporate regularization terms into our problem. In particular, we add a consistency regularizer that takes into account the inverse functional map D . Moreover, we add an isometry promoting term which is given by commutativity with the LB operator [OBBS*12] and Dirichlet energies that favor smooth basis elements. We note that other regularizers such as descriptor commutativity [NO17] or orientation preservation [RPWO18] may be also considered. Formally, we propose the following objective function

$$\begin{aligned} \mathcal{E} &= \mathcal{E}_{\text{fid}} + \mu_{\text{cfid}} \mathcal{E}_{\text{cfid}} + \mu_{\text{iso}} \mathcal{E}_{\text{iso}} + \mu_{\text{dir}} \mathcal{E}_{\text{dir}}, \\ \mathcal{E}_{\text{cfid}} &= \frac{1}{2} |B_1^T F_1 - D B_2^T F_2|_F^2, \\ \mathcal{E}_{\text{iso}} &= \frac{1}{2} |C B_1^T W_1 B'_1 - B_2^T W_2 B'_2 C|_F^2, \\ \mathcal{E}_{\text{dir}} &= \frac{1}{2} \text{Tr} (B_1^T W_1 B'_1) + \frac{1}{2} \text{Tr} (B_2^T W_2 B'_2), \end{aligned} \quad (11)$$

where $\mu_{\text{cfid}}, \mu_{\text{iso}}, \mu_{\text{dir}} \in \mathbb{R}^+$ are penalty scalars, Tr yields the trace of a matrix, and W_j is the cotangent weights matrix [PP93] of shape \mathcal{M}_j for $j = 1, 2$. One of the key aspects of our framework resulting from our ADMM formulation (2) is that it allows to combine challenging regularizers (11) in a straightforward way. Thus, the formulation with \mathcal{E} and its associated ADMM is somewhat technical, and we defer the derivation to the appendix.

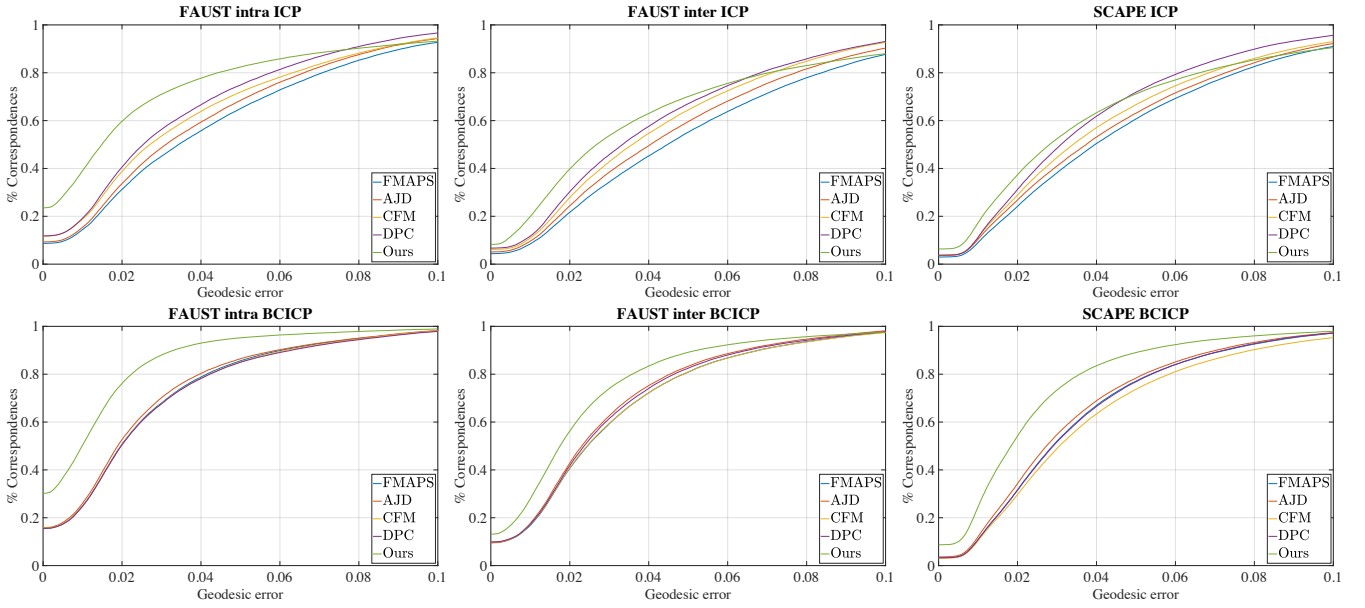


Figure 4: We compute the geodesic error between the ground-truth mapped vertices to those mapped with the above methods, and we accumulate the amount of points that reached a certain error [KLF11] and visualize it in the above graphs. Noticeably, our method yields the best averaged error over other methods for both ICP and BCICP map extraction. We note that basis design methods such as ours and AJD [KBB*13] gain the most from advanced mapping methods such as BCICP [RPWO18].




Figure 5: We visualize the quality of the computed correspondences using texture transfer. We note that the alternative methods struggle with elongated or small areas as the head and legs, whereas our method achieves improved results as can be compared to the ground-truth (GT).

6. Implementation Details

In what follows we describe a few technical aspects related to our method including dimensionality reduction of problem (2), variable initialization, stopping condition and the development platform.

Dimensionality reduction. Solving (2) directly is computationally prohibitive when the shapes consist of many vertices. To overcome this difficulty, we propose to reduce the spatial dimension and use a spectral domain instead, allowing for fast computation times while retaining a significant amount of degrees of freedom. Specifically, we take the left singular vectors obtained by computing the Singu-

lar Value Decomposition (SVD) of the given constraints. Namely,

$$\tilde{U}_j \tilde{S}_j \tilde{V}_j^T = \text{SVD}(\tilde{F}_j), \quad j = 1, 2,$$

where we denote $U_j = \tilde{U}_j(:, 1:r)$ the r most significant modes. In our experiments we choose r such that U_j covers at least 90% of the spectrum. We note that other spectral bases could be considered, e.g., the LB basis itself [KBB*13]. However, each choice leads to a different optimization with its own assumptions and challenges. In Sec. 7, we compare our approach to other methods.

To incorporate U_j into our optimization, we denote the changes in boldface and perform the following modifications,

$$\mathbf{B}_j = U_j^T \mathbf{B}_j, \quad \mathbf{F}_j = U_j^T \mathbf{F}_j, \quad \mathbf{G}_j = U_j^T \mathbf{G}_j U_j, \quad j = 1, 2, \quad (12)$$

yielding matrices of sizes $r_j \times k$, $r_j \times n$ and $r_j \times r_j$, respectively. Substituting the above components with their high dimensional counterparts is the only change needed to obtain a spectral version of Alg. 1. Finally, given \mathbf{B}_j , we reconstruct B_j via $B_j = U_j \mathbf{B}_j$. We note that while our approach strongly depends on the input features for deriving the low-dimensional subspaces, in our experiments we observed that it works quite well with a variety of descriptors such as WKS and segmentation information.

Variable initialization and stopping rule. In our tests we noticed that our method is robust to the choice of initial values. Nevertheless, we describe the particular values we used in our experiments. To initialize the primal variables B_j and B'_j , we take the first k singular vectors of the respective descriptors, \tilde{F}_j . This computation is denoted by $\text{SVD}(\tilde{F}_j, k)$ in Alg. 1. Using these bases, we can solve Eq. (9) to obtain an initial C . The dual variables P_j and Q'_j are set to zero matrices of an appropriate size. The stopping condition we used is based on the primal and dual residuals and is detailed in [BPC*11, Sec. 3.3], where the maximum number of steps is 10,000.

Development platform and parameters. We implemented our method in MATLAB, using its built-in optimization tools such as `dlyap` and `mldivide`. Our approach was tested on an Intel Core i7 2.6GHz processor with 16GB RAM. We show in Fig. 6 a runtime comparison to AJD [KBB*13] and OPC [RPWO18] on meshes of sizes $1k - 500k$ vertices. The parameters of our method include the penalty scalars μ_{cfid} , μ_{iso} and μ_{dir} for the different energy terms (11). We list our choices in Tab. 1, which also shows how much of the spectrum we employ, given by the r parameter. Finally, the size of the functional map and the associated bases was $k = 20$ unless noted otherwise.

7. Evaluation and Results

To evaluate our method, we consider several applications in which functional maps are useful such as extraction of point-to-point maps [OBSC*12], function transfer [NMR*18], and consistent quadrangulation [ACBCO17]. We test our approach on a variety of datasets including SCAPE [ASP*05], TOSCA [BBK08] FAUST [BRLB14], and SHREC07 [GBP07]. In our comparison, we consider several baseline methods such as functional maps (FMAPS) [OBSC*12], approximate joint diagonalization (AJD) [KBB*13], coupled functional maps (CFM) [ERGB16], descriptor preservation via commutativity (DPC) [NO17] and orientation preserving correspondences (OPC) [RPWO18]. In our com-

Dataset	r	μ_{cfid}	μ_{iso}	μ_{dir}
FAUST intra	0.9	$1e-4$	$1e-6$	$1e-2$
FAUST inter	0.9	$1e-4$	$1e-5$	$1e-4$
SCAPE	0.9	$1e-3$	$1e-6$	$1e-4$
Remeshed FAUST intra	0.99	$1e-2$	$1e-5$	$1e-6$
Remeshed FAUST inter	0.99	$1e-2$	$1e-5$	$1e-6$
Remeshed SCAPE	0.99	$1e-1$	$1e-5$	$1e-6$

Table 1: The parameter values used in our tests for each dataset.

parison, we only use the functional map matrices as computed using the above techniques, and we discard any other improvements related to a specific application. In all cases, we used the authors' recommended parameters or we searched for the best ones.

7.1. Extracting point-to-point maps

One of the main applications of functional maps is the computation of point-to-point correspondences between pairs of shapes. In our comparison, we consider two different scenarios. The first includes the original FAUST and SCAPE shapes using 20 landmarks and 100 Wave Kernel Signature (WKS) [ASC11] features. In the second case, we remesh the shapes and use consistent segmentation data [KO18] with WKS descriptors. We emphasize that the latter scenario is extremely challenging as it is completely automatic, it involves approximate features, and the meshes have different connectivities. The pairs we use appeared previously in [KLF11, CK15]. For map extraction we employ the ICP method proposed in [OBSC*12] and the recent BCICP approach [RPWO18], although other methods [RMC15, EBC17] could be used. Our evaluation metrics include the computation of cumulative geodesic errors [KLF11] and visualization of transferred scalar functions or textures.

In Fig. 4 we show the average cumulative geodesic errors of the first scenario. We note that our approach achieves a significant improvement over all the other competing methods. In particular, when ICP extraction is facilitated, our method yields very good results on FAUST intra which involves pairs of different poses of the same people. Interestingly, our method benefits the most from recent advances in map extraction techniques [RPWO18] as can be seen in the second row. Specifically, using BCICP increases the gap between our results vs. others on FAUST inter (different people, different pose) and SCAPE. This hints that our functional map and associated bases introduce more degrees of freedom which could be exploited in elaborated methods such as [RPWO18]. This behavior can be additionally seen in AJD [KBB*13] BCICP results which surpass most methods even though their ICP measures were lower




Figure 6: We compare the total pre-processing and computation times of the above methods on a pair of shapes for a large range of vertex counts, m . Our method is significantly faster than AJD and OPC for high vertex counts, where for low number of vertices OPC is more efficient than our approach.




Figure 7: To compare the advantages of regularization, we solve our problem with various combinations of energy terms. Using the obtained maps, we transport the coordinate functions from the target to the source and show the results above. We achieve a significant improvement when we regularize as can be seen around the chest and head (middle-right and right) vs. the non-regularized result (middle-left).

than others in general. We do not compare to OPC [RPWO18] in this setup as we use non-symmetric landmarks and thus there is no advantage in using their orientation preserving regularization over, e.g., DPC [NO17].

We demonstrate the error measures of Scenario 2 in Fig. 15. We stress that this setup is particularly challenging as the shapes do not share the connectivity and we use automatically computed features. Nevertheless, our method exhibits the best results on FAUST both for the isometric and non-isometric cases when ICP map extraction is applied. Moreover, when we utilize BCICP on FAUST our method and OPC yield the best scores compared to the alternative methods. Finally, the remeshed SCAPE was an extremely difficult test case, leading to mappings of poor quality in general for most methods (notice the y-axis gets to 0.6 instead of 1). For this dataset, CFM and DPC produced good measures for ICP, and our method and OPC were the highest with BCICP refinement.

The point-to-point correspondence allows to map information from the target to the source. In Fig. 5, we compare the mappings generated in Scenario 1 on a single pair of FAUST intra using texture transfer. The meshes in this dataset are in 1 – 1 correspondence and thus we can use the ground-truth (GT) map for comparison. Overall, the performance of the tested methods was generally good. However, small parts of the body such as hands and legs were less accurate for FMAPS, AJD and DPC. Moreover, other methods exhibit large errors in the head, whereas ours correctly finds the symmetry line (see the zoom below).

7.2. Effect of regularization

To evaluate the benefits of utilizing regularizing terms, we visualize the map quality via coordinate function transfer in Fig. 7. Indeed, there is a clear improvement when \mathcal{E}_{iso} is introduced (middle right) vs. using \mathcal{E}_{fid} alone (middle left) as can be seen on the chest and head. Adding \mathcal{E}_{dir} (right) is not beneficial in this case as it is visually indistinguishable from the \mathcal{E}_{iso} (middle right) case.




Figure 8: We plot the cumulative geodesic error for maps computed using various regularization settings. Our results indicate that the regularized problems yield better correspondences. See the text for additional details.




Figure 9: We show above the first three basis functions designed using AJD (top row) and our approach (bottom row). Different from AJD, our energy does not favor matching basis elements which allows for a more flexible design process. See also Fig. 16.

In addition to this visualization, we also run our algorithm on FAUST and SCAPE in scenario 1, using different regularization configurations. We show in Fig. 8 the cumulative geodesic error of these tests. For each dataset the solid line represents using only \mathcal{E}_{fid} , the dashed is the result when we incorporate \mathcal{E}_{iso} , the dotted line is produced by adding \mathcal{E}_{dir} , and we get the dash-dot line by minimizing the full \mathcal{E} . On average, regularization consistently improves the results. Specifically, we identify a 2.6% and 1.5% improvement on FAUST and SCAPE, respectively. See also the zoomed plots in Fig. 8. Further, the consistency term (dash-dot line) helps both with respect to the accuracy of the results and the empirical convergence of the problem. We note that the Dirichlet penalization (dotted line) improves the results of FAUST, whereas for SCAPE its contribution is less apparent.

7.3. Comparison with AJD [KBB*13]

Perhaps closest to our approach is the method that finds approximate joint diagonalized bases of Kovnatsky et al. [KBB*13]. In this work, the authors explore an optimization problem which is

conceptually similar to ours. However, there are several key differences between our technique and theirs as we detail below.

In terms of the energy functional, our technique is fundamentally different from theirs. Their approach favors basis elements which diagonalize the LB operator, leading to smooth functions. However, the disadvantage in this point of view is that one implicitly assumes that smooth basis functions span the descriptors subspace. Unfortunately, many practical descriptors that are currently used in functional map pipelines do not fit into this assumption. Indeed, any high frequency signal such as segment information will undergo a low pass filter which may lead to data loss in practice, as we show in Fig. 2. In contrast, our method does not favor smooth basis elements and may output high frequency functions, see e.g., Fig. 1. Finally, even when we include Dirichlet terms in our minimization, they are weighted weakly.

Another significant difference is in the data fidelity term. The formulation in [KBB*13] and others [LRBB17] fixes the associated functional map C to attain a *particular structure*. Namely, they include a term that takes the following form

$$\tilde{\mathcal{E}}_{\text{fid}} = \frac{1}{2} |B_1^T F_1 - B_2^T F_2|_F^2, \quad (13)$$

which can be interpreted as setting C to be $C \approx B_2^T B_1$. There are two disadvantages to formulation (13) which our approach overcomes. First, regularizing the functional map C is not straightforward as in our formulation (11), and may lead to *quartic* expressions in the unknowns B_j . Indeed, our formulation allows to independently constrain the bases or the functional map and its inverse, manifesting greater flexibility alongside the natural utilization of state-of-the-art regularizers. Second, our method allows for *general* functional map matrices and thus it increases the search space of solutions when compared with AJD frameworks.

To summarize, our approach generalizes AJD methods in that it combines work on joint diagonalization and functional map optimization in a unified framework. There are three key differences in our technique. First, we consider a much larger search space of solutions as we utilize the Proper Orthogonal Decomposition (POD) modes which are better suited to the given features, and we further allow for general functional map matrices. Second, since we jointly optimize for the functional map and the bases, we can naturally incorporate regularization terms. Finally, on the algorithmic side, AJD approaches facilitate a constrained minimization tool which is inefficient in practice as can be seen in Fig. 6 and its convergence is not guaranteed. In contrast, we analyze our approach and show that it is similar to a provably convergent problem.

In addition to this qualitative comparison, we show in Figs. 9 and 16 the differences between the designed basis elements. Indeed, AJD (top row) produces highly consistent basis functions compared to ours (bottom row). However, we believe that this behavior limits the design process significantly, which may lead to less accurate matching results as can be seen in Figs. 4 and 15. Specifically, we select a pair of shapes from TOSCA and visualize the correspondence differences via texture transfer in Fig. 5. Overall, AJD produces reasonable results as compared to the ground-truth (GT). However, various parts of the shape such as head, legs and tail, dis-




Figure 10: We compute functional maps and bases using AJD and our method, and we compare the results to the ground-truth (GT) via texture transfer.

play large errors. In contrast, our technique was able to accurately match most areas of the shapes including the challenging parts.

To conclude our qualitative comparison, we modify AJD to use POD modes in their design process instead of the LB eigenfunctions and we plot the cumulative geodesic error that was obtained for FAUST in Fig. 11. Indeed, switching to POD modes (blue dashed line) yields a large improvement compared to LB-based AJD (blue line). However, our method (red line) is still significantly more accurate, which can be attributed in part to the state-of-the-art regularization terms we include in our optimization.

7.4. Consistent quadrangulation and function transfer

The increasing interest in the functional map approach over the last few years lead to the development of techniques which can utilize a given functional map directly, without the need to convert it to a point-to-point map. For instance, [ACBCO17] proposed an optimization framework for designing consistent cross fields on a pair of shapes for the purpose of generating approximately consistent quadrangular remeshings of the input shapes. Our computed functional map and bases can be directly used within their method




Figure 11: Switching to POD-based design with AJD (blue dash line) yields an improvement over the LB subspaces (blue line). Still, our framework generates correspondences that are more accurate.




Figure 12: Our technique fits methods whose input is a functional map with its associated bases such as [ACBCO17]. We demonstrate the consistent quadrangulations obtained by using their method with input generated by our approach.

to produce quad meshes. In Fig. 12, we show an example of the remeshing results of two pairs of shapes having different connectivities (left and right) and genus (right). Still, we obtain highly consistent results as can be seen in the matching singularity points (red spheres). We provide an additional instance of this pipeline in Fig. 13 comparing the quadrangulation achieved with fixed LB bases (left) vs. our technique with designed POD modes (right). Indeed, we observe a much better alignment of isolines and singularity points with our approach compared to Fixed LB.

The last application we consider involves the transfer of scalar valued information between shapes. Recently, [NMR*18] showed that by extending the usual functional basis to include basis products, an improved function transfer can be performed which we call *Product transfer*. In Fig. 14, we utilize this pipeline using our functional map and bases to transfer an extremely challenging data given by a localized Gaussian function. Indeed the transfer is improved using the extended basis as the noise is less severe and the maximum is more localized.

8. Limitations

One limitation of our framework is related to the dependencies between the given constraints and our choice of dimensionality reducing subspaces U_j . Indeed, one can always add the standard LB spectrum to these subspaces. However, we observe that in general, the results may change depending on the particular subspace in use and its size. For instance, while increasing r allows for greater flexibility for representing scalar functions, it also requires more regularization, otherwise unwanted solutions may potentially become local minimizers. Another shortcoming of our approach is that it tends to produce maps that are less smooth compared to those generated with LB bases. This behavior is somewhat expected, as our bases are designed to potentially transfer high frequency information which in turn leads to less uniform correspondences. We leave further investigation of these aspects to future work.




Figure 13: We compare the quadrangulations produced by using fixed LB bases (left) and designed POD modes (right). Overall, there is a significant improvement in the designed case in terms of isolines alignment, singularity matches and adherence to curvature (see e.g., the index finger, left).

9. Conclusions and Future Work

In this paper, we proposed a method for designing basis elements on a pair of triangle meshes along with an associated functional map. Unlike most existing work which utilize the spectrum of the Laplace–Beltrami operator, our technique adopts the Proper Orthogonal Decomposition (POD) modes to reduce the dimensionality of the problem. This choice introduces many degrees of freedom and it significantly extends the space of potential solutions. To effectively solve the problem, we incorporate state-of-the-art regularization terms which promote consistency, isometry and smoothness. Our optimization scheme is based on the Alternating Direction Method of Multipliers (ADMM) and it consists of easy-to-solve linear or Sylvester-type equations. We show that in practice our method behaves well in terms of convergence, and we additionally prove that a similar problem to ours is guaranteed to converge. We evaluate our machinery in the context of shape matching, function transfer and consistent quadrangulation, and we demonstrate that our results yield a significant improvement over state-of-the-art approaches for computing functional maps.

In the future, we would like to characterize the dependencies between the subspaces spanned by the bases to the given constraints and the relation to the functional map. Moreover, we believe that




Figure 14: Mapping Gaussian function between shapes with different connectivities is a challenging task, whose results may be exploited in context of shape matching to construct an accurate correspondence or to improve a given one.

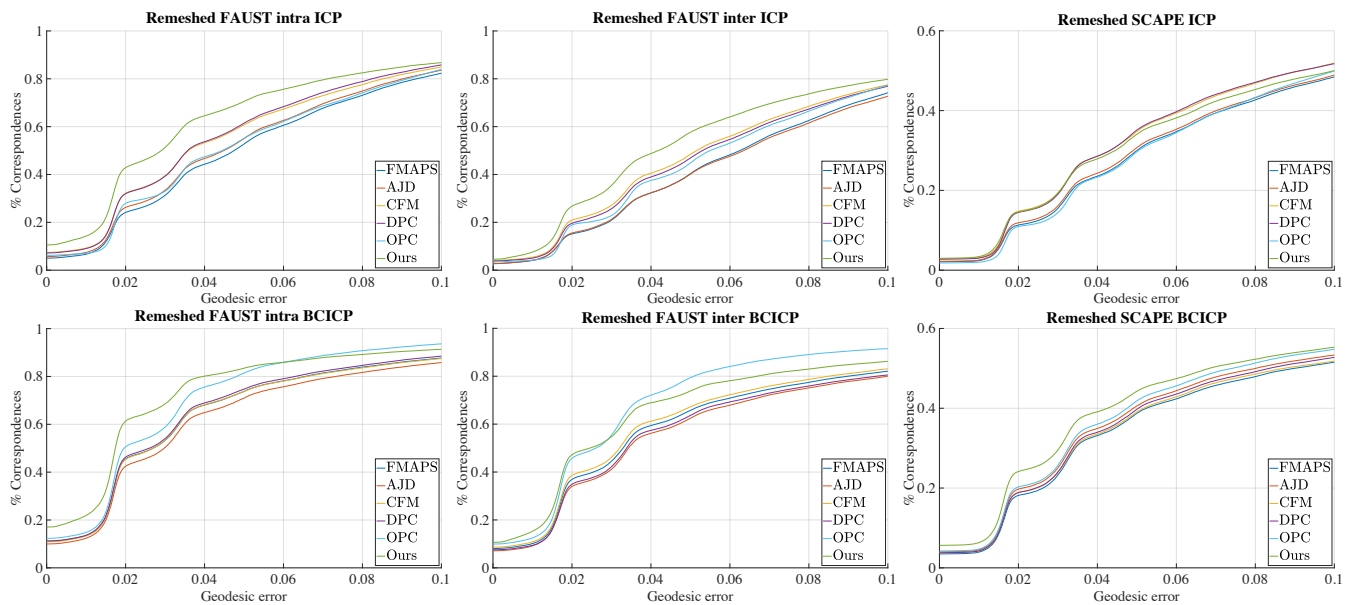


Figure 15: We evaluate our approach in the challenging scenario of shapes with different connectivities and automatically generated descriptors. Our method exhibits very good error measures as compared to state-of-the-art approaches for computing functional maps.

many applications may benefit from the proposed pipeline on a single shape. Namely, generate a self functional map with a set of basis elements defined on the shape. Examples include symmetry detection, fluid simulation and data interpolation, among many other possibilities. On the other hand, extending our framework to handle multiple shapes is an interesting direction as well. Finally, we believe that many of the questions that we consider in our work could benefit from the recent advancements in machine learning with deep neural networks. We plan to investigate how the task of designing a basis and a functional map can be solved using deep learning approaches.

References

- [ABB*16] AFLALO Y., BREZIS H., BRUCKSTEIN A., KIMMEL R., SOCHEN N.: Best bases for signal spaces. *Comptes Rendus Mathématique* 354, 12 (2016), 1155–1167. 1
- [ACBCO17] AZENCOT O., CORMAN E., BEN-CHEN M., OVSJANIKOV M.: Consistent functional cross field design for mesh quadrangulation. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 92. 7, 9, 10
- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on* (2011), IEEE, pp. 1626–1633. 1, 2, 3, 7
- [ASP*05] ANGUELOV D., SRINIVASAN P., PANG H.-C., KOLLER D., THRUN S., DAVIS J.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in neural information processing systems* (2005), pp. 33–40. 7
- [AWO*14] AZENCOT O., WEISSMANN S., OVSJANIKOV M., WARDETSKY M., BEN-CHEN M.: Functional fluids on surfaces. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 237–246. 2
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 7
- [BHL93] BERKOOZ G., HOLMES P., LUMLEY J. L.: The proper orthogonal decomposition in the analysis of turbulent flows. *Annual review of fluid mechanics* 25, 1 (1993), 539–575. 3
- [BKP*10] BOTSCH M., KOBELT L., PAULY M., ALLIEZ P., LÉVY B.: *Polygon mesh processing*. AK Peters/CRC Press, 2010. 3
- [BM92] BESL P. J., MCKAY N. D.: Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures* (1992), vol. 1611, International Society for Optics and Photonics, pp. 586–607. 1
- [BPC*11] BOYD S., PARIKH N., CHU E., PELEATO B., ECKSTEIN J., ET AL.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122. 5, 7
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2014), pp. 3794–3801. 7
- [CK15] CHEN Q., KOLTUN V.: Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 2039–2047. 7
- [EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 165–174. 2, 7
- [ERGB16] EYNARD D., RODOLA E., GLASHOFF K., BRONSTEIN M. M.: Coupled functional maps. In *3D Vision (3DV), 2016 Fourth International Conference on* (2016), IEEE, pp. 399–407. 7
- [GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: Shape retrieval contest 2007: Watertight models track. *SHREC competition* 8, 7 (2007). 7
- [GGC18] GAO W., GOLDFARB D., CURTIS F. E.: Admm for multiaffine constrained optimization. *arXiv preprint arXiv:1802.09592* (2018). 5, 13

- [GM75a] GABAY D., MERCIER B.: *A dual algorithm for the solution of non linear variational problems via finite element approximation*. Institut de recherche d'informatique et d'automatique, 1975. 4
- [GM75b] GLOWINSKI R., MARROCO A.: Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de dirichlet non linéaires. *Revue française d'automatique, informatique, recherche opérationnelle. Analyse numérique* 9, R2 (1975), 41–76. 4
- [GNVL79] GOLUB G., NASH S., VAN LOAN C.: A hessenberg-schur method for the problem $ax + xb = c$. *IEEE Transactions on Automatic Control* 24, 6 (1979), 909–913. 4
- [HQ12] HOU T., QIN H.: Continuous and discrete mexican hat wavelet transforms on manifolds. *Graphical Models* 74, 4 (2012), 221–232. 2
- [HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 36. 1, 2
- [KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 439–448. 1, 2, 3, 6, 7, 8, 9
- [KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015), pp. 905–914. 1, 2
- [KGB16] KOVNATSKY A., GLASHOFF K., BRONSTEIN M. M.: Madmm: a generic algorithm for non-smooth optimization on manifolds. In *European Conference on Computer Vision* (2016), Springer, pp. 680–696. 1, 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)* (2011), vol. 30, ACM, p. 79. 6, 7
- [KMP*21] KIRGO M., MELZI S., PATANÈ G., RODOLÀ E., OVSIJANIKOV M.: Wavelet-based heat kernel derivatives: Towards informative localized shape analysis. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 165–179. 2
- [KO18] KLEIMAN Y., OVSIJANIKOV M.: Robust structure-based shape correspondence. In *Computer Graphics Forum* (2018), Wiley Online Library. 2, 7
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Fully spectral partial shape matching. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 247–258. 1, 2, 9
- [MRMO20] MARIN R., RAKOTOSAONA M.-J., MELZI S., OVSIJANIKOV M.: Correspondence learning via linearly-invariant embedding. *Advances in Neural Information Processing Systems* 33 (2020). 2
- [NMR*18] NOGNENG D., MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M., OVSIJANIKOV M.: Improved functional mappings via product preservation. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 179–190. 2, 7, 10
- [NO17] NOGNENG D., OVSIJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 259–267. 1, 2, 5, 7, 8
- [OBBS*12] OVSIJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 30. 1, 2, 5, 7
- [OCB*16] OVSIJANIKOV M., CORMAN E., BRONSTEIN M., RODOLÀ E., BEN-CHEN M., GUIBAS L., CHAZAL F., BRONSTEIN A.: Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses* (2016), ACM, p. 9. 2
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental mathematics* 2, 1 (1993), 15–36. 5
- [RBG*09] REUTER M., BIASOTTI S., GIORGI D., PATANÈ G., SPAGNUOLO M.: Discrete laplace-beltrami operators for shape analysis and segmentation. *Computers & Graphics* 33, 3 (2009), 381–390. 1
- [RMC15] RODOLÀ E., MOELLER M., CREMERS D.: Point-wise map recovery and refinement from functional correspondence. *arXiv preprint arXiv:1506.05603* (2015). 2, 7
- [ROA*13] RUSTAMOV R. M., OVSIJANIKOV M., AZENCOT O., BEN-CHEN M., CHAZAL F., GUIBAS L.: Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 72. 2
- [RPWO18] REN J., POULENARD A., WONKA P., OVSIJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *arXiv preprint arXiv:1806.04455* (2018). 2, 5, 6, 7, 8
- [Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the fifth Eurographics symposium on Geometry processing* (2007), Eurographics Association, pp. 225–233. 1
- [SBL18] SCHONSHECK S. C., BRONSTEIN M. M., LAI R.: Nonisometric surface registration via conformal laplace-beltrami basis pursuit. *arXiv preprint arXiv:1809.07399* (2018). 2
- [SRGB14] SOLOMON J., RUSTAMOV R., GUIBAS L., BUTSCHER A.: Earth mover's distances on discrete surfaces. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 67. 1
- [WYZ18] WANG Y., YIN W., ZENG J.: Global convergence of ADMM in nonconvex nonsmooth optimization. *Journal of Scientific Computing*, 2018. 5
- [Xiu10] XIU D.: *Numerical methods for stochastic computations: a spectral method approach*. Princeton university press, 2010. 3

Updating the bases, B_1 and B_2

The necessary first-order optimality conditions of the augmented Lagrangian of B_1 , yield the following equation.

$$\begin{aligned}
C_1 &= \mathcal{A}_1 B_1 B_1 + \mathcal{O}_1 B_1, \\
A_1 &= F_1 F_1^T + \mu_{\text{iso}} W_1 B_1' B_1'^T W_1^T, \\
B_1 &= C^T C, \\
C_1 &= F_1 F_2^T B_2 C + \mu_{\text{iso}} W_1 B_1' C^T B_2'^T W_2 B_2 C - \mu_{\text{dir}} W_1 B_1' \\
&\quad + \mu_{\text{cfid}} F_1 F_2^T B_2 D^T + \rho G_1 B_1' (I - P_1)^T + \rho G_1 (B_1' - Q_1'), \\
\mathcal{O}_1 &= \mu_{\text{cfid}} F_1 F_1^T + \rho G_1 B_1' B_1'^T G_1 + \rho G_1.
\end{aligned} \tag{20}$$

Similar to the non-regularized case, this is a Sylvester-type equation which can be solved using, e.g., MATLAB's `dlyap` with the call $B_1 = \text{dlyap}(-\mathcal{O}_1^{-1} A_1, B_1, \mathcal{O}_1^{-1} C_1)$. The update of B_2 consists of another Sylvester-type equation given by

$$\begin{aligned}
C_2 &= \mathcal{A}_2 B_2 B_2 + \mathcal{O}_2 B_2, \\
A_2 &= \mu_{\text{cfid}} F_2 F_2^T, \\
B_2 &= D^T D, \\
C_2 &= F_2 F_1^T B_1 C^T + \mu_{\text{iso}} W_2 B_2' C B_1'^T W_1^T B_1 C^T - \mu_{\text{dir}} W_2 B_2' \\
&\quad + \mu_{\text{cfid}} F_2 F_1^T B_1 D + \rho G_2 B_2' (I - P_2)^T + \rho G_2 (B_2' - Q_2'), \\
\mathcal{O}_2 &= F_2 F_2^T + \mu_{\text{iso}} W_2 B_2' C C^T B_2'^T W_2^T + \rho G_2 B_2' B_2'^T G_2 + \rho G_2.
\end{aligned} \tag{21}$$

Updating the auxiliary variables, B_1' and B_2'

The update for B_1' is very similar to the non-regularized case. Namely, we obtain the following linear system.

$$\begin{aligned}
B_3 &= \mathcal{A}_3 B_1' \\
A_3 &= \mu_{\text{iso}} W_1 B_1 C^T C B_1'^T W_1 + \rho G_1 B_1 B_1'^T G_1 + \rho G_1, \\
B_3 &= \mu_{\text{iso}} W_1 B_1 C^T B_2'^T W_2 B_2' C - \mu_{\text{dir}} W_1 B_1 \\
&\quad + \rho G_1 B_1 (I - P_1) + \rho G_1 (B_1 + Q_1').
\end{aligned} \tag{22}$$

Unlike the non-regularized version, the update of B_2' involves a Sylvester-type equation of the form.

$$\begin{aligned}
C_4 &= \mathcal{A}_4 B_2' B_4 + \mathcal{O}_4 B_2', \\
A_4 &= \mu_{\text{iso}} W_2 B_2 B_2'^T W_2, \\
B_4 &= C C^T, \\
C_4 &= \mu_{\text{iso}} W_2 B_2 C B_1'^T W_1 B_1' C^T - \mu_{\text{dir}} W_2 B_2 \\
&\quad + \rho G_2 B_2 (I - P_2) + \rho G_2 (B_2 + Q_2'), \\
\mathcal{O}_4 &= \rho G_2 B_2 B_2'^T G_2 + \rho G_2.
\end{aligned} \tag{23}$$

Updating the functional maps, C and D

Updating the functional map can be done by vectorizing the linear system we obtain by taking the gradient of the augmented Lagrangian with respect to C . This leads to the following vectorized

equation

$$\begin{aligned}
\text{vec}(B_5) &= A_5 \text{vec}(C), \\
A &= (B_1^T F_1 F_1^T B_1 + \mu_{\text{iso}} W_{B,1} W_{B,1}^T + \rho D D^T) \otimes I \\
&\quad + I \otimes (\mu_{\text{iso}} W_{B,2} W_{B,2}^T + \rho D^T D) \\
&\quad - \mu_{\text{iso}} W_{B,1}^T \otimes W_{B,2}^T - \mu_{\text{iso}} W_{B,1} \otimes W_{B,2}, \\
B_5 &= B_2^T F_2 F_1^T B_1 + \rho (I - R_1) D^T + \rho D^T (I - R_2),
\end{aligned} \tag{24}$$

where $W_{B,j} = B_j^T W_j B_j'$. For the matrix D , we obtain the following

$$\begin{aligned}
C_6 &= A_6 D + D B_6, \\
A_6 &= \rho C^T C, \\
B_6 &= \mu_{\text{cfid}} B_2^T F_2 F_2^T B_2 + \rho C C^T, \\
C_6 &= \mu_{\text{cfid}} B_1^T F_1 F_2^T B_2 + \rho C^T (I - R_1) + \rho (I - R_2) C^T.
\end{aligned} \tag{25}$$

Updating the dual variables, P_j and Q_j'

Finally, the last step of our scheme is equivalent to the non-regularized case and for $j = 1, 2$, it is given by

$$\begin{aligned}
P_j &= P_j + B_j^T G_j B_j' - I, \\
Q_j' &= Q_j' + B_j - B_j', \\
R_1 &= R_1 + C D - I, \\
R_2 &= R_2 + D C - I.
\end{aligned} \tag{26}$$