

Adjoint Map Representation for Shape Analysis and Matching

Ruqi Huang and Maks Ovsjanikov

LIX, Ecole Polytechnique

Abstract

In this paper, we propose to consider the adjoint operators of functional maps, and demonstrate their utility in several tasks in geometry processing. Unlike a functional map, which represents a correspondence simply using the pull-back of function values, the adjoint operator reflects both the map and its distortion with respect to given inner products. We argue that this property of adjoint operators and especially their relation to the map inverse under the choice of different inner products, can be useful in applications including bi-directional shape matching, shape exploration, and pointwise map recovery among others. In particular, in this paper, we show that the adjoint operators can be used within the cycle-consistency framework to encode and reveal the presence or lack of consistency between distortions in a collection, in a way that is complementary to the previously used purely map-based consistency measures. We also show how the adjoint can be used for matching pairs of shapes, by accounting for maps in both directions, can help in recovering point-to-point maps from their functional counterparts, and describe how it can shed light on the role of functional basis selection.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: —Shape Analysis

1. Introduction

Encoding and analyzing maps between shapes lies at the core of many geometry processing applications, such as deformation transfer [SP04], shape interpolation [KMP07, VTSSH15], and visualization [PRMH10], to name a few. Perhaps the most primitive and intuitive representation of a map is given by the correspondences between points on the two shapes. However, pointwise maps both do not allow to express more complex (e.g., one-to-many) relations and can often lead to difficult optimization problems in practice.

Stemming from the framework of functional maps [OBCS*12], there has been a recent trend of employing functional approaches in map-based shape analysis, which overcome some of the difficulties arising from using pointwise correspondences. The key idea of the functional maps is to consider the correspondences between *functions* defined on shapes, which generalizes the notion of point-to-point maps and provides a flexible and simple framework for encoding, analyzing and inferring maps between shapes.

Although the original formulation of functional maps was primarily proposed for solving shape matching problems (map inference), it can also be used to study how geometry is deformed from one shape to the other. The connection is made explicit in a follow-up work [ROA*13], where the authors argue that given a map between a pair of shapes, the metric distortion induced by the map can be encoded by a pair of so-called shape difference operators. Moreover, they propose a closed-form expression for the shape difference operators in terms of functional maps.

Despite their appealing properties, both functional maps and shape difference operators have their limitations. For instance, functional maps focus on transporting function values across shapes and thus ignore the distortion induced by the maps. Meanwhile, although shape difference operators provide a compact way of encoding the distortion, they do not transport functions across the shapes of interest. In fact, one usually needs to choose a base shape and then construct the shape difference operators with respect to it. This makes it inconvenient to apply shape difference operators in analyzing shape collections, especially in cases when there are significant deformations among the shapes since the choice of the base shape might influence the final results.

In this paper, we propose to consider the *adjoint operators* (which we also call the *adjoint representation*) of functional maps, which, as we demonstrate, combine the properties of the functional maps and the shape difference operators. Similarly to a functional map, the adjoint transports functions across the spaces on different shapes, but at the same time distorts the function values in a way that reflects the deformation induced by a map.

Remark that the adjoint operator corresponds to a well-known classical construction that is closely related to the functional maps, and indeed can in certain cases be obtained simply via the transpose of the functional map matrix in the discrete setting. Nevertheless, we argue that it has certain properties that have not been fully exploited or analyzed so far, especially in relation to analyzing maps in shape collections. Thus, we start by describing the basic proper-

ties of the adjoint operators and their inherent connections to functional maps and shape difference operators. We then demonstrate three immediate applications, facilitated by the adjoint representation.

In the first application, we show how the adjoint operators can be used to regularize functional map computations for shape matching. Namely, we propose a simple scheme using the adjoint operators to promote invertibility of the map, while giving control over the conformality and area-preservation of the sought correspondence. The resulting method maintains the linear nature of the optimization problem and achieves improvement in map quality.

We then consider map-based exploration and analysis of shape collections, and show how the adjoint representation can be used to overcome the dependence of the approach proposed in [ROA*13] on using a prescribed base shape, by describing how this representation can be used within the cycle-consistency framework of [WHG13]. In particular, we show how the adjoint representation can be used to evaluate the consistency of shape distortions in a collection, and to reveal the jointly most distorted regions.

Finally, we consider recovering a point-to-point map from a given low-rank functional map, and demonstrate that the link between the inverse of a map and its adjoint suggests a change of the inner product that can be used to define new functional bases, which lead to smaller error in pointwise recovery. The resulting scheme is simple to implement and improves upon the original nearest neighbor search scheme in [OBCS*12] and even the more advanced approach proposed in [RMC15] in the case of non-isometric shape pairs, without using iterative optimization.

2. Related Work

Encoding and analyzing maps is one of best-studied areas of geometry processing, with its roots both in the classical problems of shape correspondence (matching), deformation design and analysis, and even parameterization, in the specific setting of mapping to some canonical domain. Therefore, since the complete overview of approaches for manipulating mappings is out-of-scope of the current article, below we concentrate primarily on various representations for mappings or general relations between non-rigid shapes.

By far the most common approach for representing and analyzing relations between geometric shapes is given by the classical notion of point-to-point correspondence, and most early methods for both rigid and non-rigid shape matching are based on finding the optimal pointwise map under some quality criterion, such as conformality or near-isometry (e.g., [BBK06, LF09, KLF11] among many others). Despite their simplicity, pointwise correspondences make it non-trivial to express more complex relations such as one-to-many maps, can lead to noisy distortion measures and perhaps most importantly often result in difficult non-convex non-linear optimization problems.

More recently another very successful set of approaches has been proposed based on soft or approximate (also called “fuzzy”) correspondences rather than point-to-point maps [SNB*12, OBCS*12]. This includes both maps between probability densities on the shapes [Mém11, SNB*12, SPKS16] and region-based maps [CK15,

GSTOG16], which can be used in a hierarchical way to obtain accurate correspondences. Many of these methods are closely related to the formalism of optimal transport, which has well-developed theoretical and computational foundations, with some recent techniques achieving remarkable scalability in solving certain problems, e.g., [SDGP*15]. These techniques are often more robust in the presence of geometric and discretization variability, although they still often become costly for large-scale matching problems.

Our work is most closely related to the functional maps framework introduced in [OBCS*12], and extended in several follow-up works, including [PBB*13, KBB*13, RMC15, KBBV15, RCB*16, LRB*16, KGB16, LRBB17, NO17] among others, with a recent overview provided in [OCB*16]. Rather than trying to establish correspondences between points or probability densities on the shapes, these methods consider general linear transformations between corresponding functional spaces. A key property exploited in these techniques is the fact that certain functional spaces enjoy a vector-space (Hilbert) structure. This allows to represent a map via a small matrix using a multi-scale functional basis such as the Laplace-Beltrami eigenfunctions, and, at the same time, to express many objectives, such as descriptor preservation, as linear constraints on the functional map, which can then be recovered by solving a least squares system.

In addition to being flexible and leading to relatively simple optimization problems, functional maps have also been used for finding correspondences jointly within a collection of shapes [WHG13, HWG14] and for analyzing distortion or differences between pairs [OBCCG13] or sets of shapes [ROA*13]. The latter work introduced the notion of shape differences, which intuitively measure the metric distortion induced by a functional map with respect to some fixed base shape.

Despite the success of these approaches in certain settings, they have some significant limitations. For example, the consistent map framework [WHG13, HWG14] is only built using the notion of loop closure of the maps themselves and thus does not allow to reveal or analyze distortion or differences in a collection jointly. On the other hand, shape difference operators [ROA*13] and the associated map visualization techniques assume a prescribed base shape and a particular direction for a map, which can lead to a bias on the reported distortion in a shape collection. Finally, estimating a point-to-point correspondence from a given functional map can be a difficult problem in itself [RMC15], especially for maps between significantly non-isometric shapes.

In this paper we argue that some of these difficulties can be overcome by considering the *adjoint* representation of a map. The adjoint is a classical construction, which is well-known to exist, in particular, for linear maps between Hilbert spaces [RS81] and has been exploited widely in areas as diverse as quantum mechanics [Mil08], control theory [GP00] and fluid simulation [MTPS04]. In the context of shape analysis, the adjoint appears ubiquitously in manipulating vector fields [dGDT16] and even in the definition of the classical Laplacian operator. Indeed, when the map is represented in a basis that is orthonormal with respect to the given inner products, then the adjoint is nothing but the transpose. Nevertheless, as we demonstrate below, by considering the adjoint, and in particular its relation to the map inverse, we can shed light on some

problems in shape analysis, including bi-directional shape matching, exploration and analysis via cycle-consistency of map distortions, and pointwise map recovery. Moreover, by exploring the effect of the inner product on the adjoint, we can recover functional bases that are better suited for some transformations.

To summarize, our main contributions include:

- We introduce the adjoint operators of a functional map with respect to different inner products and describe their connection to both the map inverse and the previously proposed shape difference operators.
- We show how the adjoint operators can be used in applications such as bi-directional shape matching, shape exploration and pointwise map recovery.
- We demonstrate how using the adjoint operators within the cycle consistency framework of [WHG13,HWG14] can reveal the presence or lack of consistency of shape deformations in a collection, complementary to purely map-based consistency.

3. Adjoint Functional Maps

We first give a brief overview of the functional maps and the shape difference operators, which are inherently related to the adjoint operators. We also introduce the notations that we will use throughout the rest of the paper.

3.1. Notation and Background

Throughout this paper, we consider shapes to be two-dimensional connected, compact Riemannian manifolds embedded in \mathbb{R}^3 , and represented as triangle meshes in the discrete setting. Especially, we assume in the following that $T : M \rightarrow N$ is such that the induced functional map $C_{N,M}$ is a continuous operator. One sufficient condition for the continuity is that the Jacobian of T is non-degenerate. We refer the readers to [TY99] for a more detailed discussion.

Functional Maps Given a pair of shapes M, N and a map $T : M \rightarrow N$, a functional map $C_{N,M}$ is simply a pullback of T that maps each function in $L^2(N)$ to one in $L^2(M)$, where $L^2(\cdot)$ is the set of square integrable real-valued functions on a shape.

It has been demonstrated in [OBCS*12] that by choosing proper functional bases for both spaces $L^2(M)$ and $L^2(N)$ respectively, the functional map can be represented by a (possibly infinite-dimensional) matrix. In the rest of this paper, we denote by $C_{ij} : L^2(S_i) \rightarrow L^2(S_j)$ the functional map in the continuous case and its associated matrix representation in the discrete case for a lighter notion. Before we introduce other useful functional operators that are based on the functional maps, we formally define two types of functional inner products.

Definition 3.1 Given a shape S , we define the area-based inner product on $L^2(S)$ as

$$h_S^A(f, g) = \int_S f(x)g(x)dV(x), \quad (1)$$

and the conformal inner product on $H_0^1(S)$, the Sobolev space modulo constants, as

$$h_S^C(f, g) = \int_S \nabla f(x) \cdot \nabla g(x)dV(x). \quad (2)$$

Shape Difference Operators The authors of [ROA*13] propose shape difference operators, which are linear operators that capture and encode certain types of deformations between shapes with respect to a given functional map. Given a pair of shapes M, N and a functional map $C_{M,N} : L^2(M) \rightarrow L^2(N)$, in general, for a pair of functions f, g on M , we have

$$h_M(f, g) \neq h_N(C_{M,N}f, C_{M,N}g),$$

where h_M (resp. h_N) is a certain inner product on the function space on M (resp. N). The shape difference operators are then introduced to compensate for such a discrepancy. Namely, the area-based shape difference operator, $D_{M,N}^A : L^2(M) \rightarrow L^2(M)$, is a linear operator such that,

$$h_M^A(f, D_{M,N}^A(g)) = h_N^A(C_{M,N}(f), C_{M,N}(g)), \forall f, g \in L^2(M)$$

Corresponding to the other inner product defined in Eq. 2, the conformal shape difference operator, $D_{M,N}^C : H_0^1(M) \rightarrow H_0^1(M)$, is a linear operator such that,

$$h_M^C(f, D_{M,N}^C(g)) = h_N^C(C_{M,N}(f), C_{M,N}(g)), \forall f, g \in H_0^1(M)$$

3.2. Formulation

In this paper, our key observation is that the adjoint operator of a functional map provides a convenient way to combine information both from the functional maps and from the shape difference operators, which in particular will be shown useful for both shape matching and shape analysis.

The following theorem (which we prove in the appendix for completeness, to draw a link to the shape difference operators) guarantees the existence and uniqueness of the adjoint operator, $X_{M,N}$, induced by $T : M \rightarrow N$.

Theorem 3.1 Given two shapes M and N , endowed with inner products h_M, h_N respectively, and a map $T : M \rightarrow N$, there exists a unique linear operator $X_{M,N} : L^2(M) \rightarrow L^2(N)$ satisfying:

$$h_N(X_{M,N}(f_M), g_N) = h_M(f_M, g_N \circ T), \forall f_M, g_N. \quad (3)$$

for any pair of functions f_M, g_N on M, N respectively.

It is worth noting that the adjoint functional operator maps functions in the same direction as the given map T , which is opposite to that of functional maps.

As claimed in Theorem 3.1, like functional maps defined via the pull-back, the adjoint representation $X_{M,N}$ is well-defined without assuming T being injective nor surjective. Notice that since $g_N \circ T = C_{N,M}(g_N)$ simply by definition of functional maps, $X_{M,N}$ is the adjoint to $C_{N,M}$ with respect to inner products h_M, h_N . In this paper, we consider the two types of inner products defined in Definition 3.1, which are used ubiquitously in geometry processing, and denote by X^A (resp. X^C) the area-based (resp. conformal) adjoint operator.

3.3. Connection to Shape Difference Operators

As can be readily seen from the definitions above, the adjoint operators are closely related to both functional maps and shape difference operators. Below we formalize this intuition by showing

that under an additional assumption that the map T is bijective, the adjoint of functional map C , is simply a composition of a shape difference operator and an inverse of the map.

Proposition 3.1 Let $D_{M,N}^A$ (resp. $D_{M,N}^C$) be the area-based (resp. conformal) shape difference operator, $X_{M,N}^A$ (resp. $X_{M,N}^C$) be the area-based (resp. conformal) adjoint to $C_{N,M}$, and $C_{M,N} : L^2(M) \rightarrow L^2(N)$ be the functional map induced by T^{-1} . then we have

$$X_{M,N}^A = D_{N,M}^A C_{M,N}, \quad (4)$$

$$X_{M,N}^C = D_{N,M}^C C_{M,N}. \quad (5)$$

On the other hand, it was shown in [ROA*13] that a bijection T is locally area-preserving (resp. conformal) if and only if $D_{N,M}^A$ (resp. $D_{N,M}^C$) is an identity operator. Therefore, Proposition 3.1 leads to an alternative characterization of the above equivalence in terms of the adjoint functional maps. Under the condition of Proposition 3.1, we claim that T is an area-preserving (resp. conformal) map if and only if $C_{M,N}$ is the area-based (resp. conformal) adjoint to $C_{N,M}$. In other words, *the adjoint of the map will equal its inverse*, when both are represented as functional operators, if and only if the map is area-preserving or conformal, depending on the choice of inner products. In the rest of the paper we will consider ways to explore and use this property in a variety of shape processing applications.

Another connection between the area-based adjoint operator and the functional map is stated in the following proposition (which will be used in the pointwise map recovery scheme in Section 7).

Proposition 3.2 Let $X_{M,N}^A$ be the area-based adjoint operator of $C_{N,M}$, and let $\kappa = X_{N,M}^A(\mathbf{1}_N)$, where $\mathbf{1}_N$ is the constant function on N . Then we have

$$X_{M,N}^A(\kappa \cdot f) = C_{M,N}(f), \forall f \in L^2(M).$$

4. Discretization

In practice, we assume that all shapes are represented as manifold triangle meshes without boundary, and that functions are represented by values assigned to vertices of the mesh.

Given a pair of shapes M, N let C be a functional map between functional spaces on N and M , $C : L^2(N) \rightarrow L^2(M)$. For example, C can represent the pull-back with respect to some point-to-point map $T : M \rightarrow N$, or can be an arbitrary linear map across the two functional spaces. In both cases, C is represented as a matrix with dimensions $k_M \times k_N$ where k_M, k_N correspond to the dimensionality of the chosen functional bases.

Setting aside the exact choice of basis for now, in the discrete setting the inner product between two functions f, g , represented via their coefficients \mathbf{f}, \mathbf{g} in some functional basis, is given as: $\langle f, g \rangle_H = \mathbf{f}^T H \mathbf{g}$, where H is some symmetric positive-definite matrix. Therefore, the discrete version of Eq. 3 above becomes:

$$\langle X \mathbf{f}, \mathbf{g} \rangle_N = \langle \mathbf{f}, C \mathbf{g} \rangle_M, \Leftrightarrow \mathbf{f}^T X^T H_N \mathbf{g} = \mathbf{f}^T H_M C \mathbf{g}. \quad (6)$$

Here, H_M and H_N are inner product matrices for functions on M




Figure 1: Map accuracy on 100 pairs of shapes in the FAUST dataset obtained using the regular functional maps pipeline, consistent maps approach of [ERGB16] and using adjoint regularization in the case of a small number of descriptors.

and N respectively. Since this must hold for all pairs of functions f, g , we get $X^T H_N = H_M C$, or:

$$X = H_N^{-1} C^T H_M. \quad (7)$$

Remark that X has size $k_N \times k_M$ and maps functions from shape M to N , i.e., in the opposite direction to that of C .

As done in [ROA*13], we consider the area-based (L^2) and conformal (H^1) inner products. In the discrete setting, the former corresponds to the area-matrix A (which can either be diagonal in the case of lumped area weights, or more accurate FEM approximation [MDSB03]), and the latter is the cotangent weight matrix W [PP93]. Note that given a basis encoded as columns of a matrix B , the corresponding area-based and conformal inner product matrices, expressed in that basis become: $B^T A B$ and $B^T W B$.

This leads to the following expressions for area-based and conformal adjoint operators:

$$X^A = (B_N^T A_N B_N)^{-1} C^T (B_M^T A_M B_M) \quad (8)$$

$$X^C = (B_N^T W_N B_N)^+ C^T (B_M^T W_M B_M) \quad (9)$$

The last expression involves the pseudo-inverse ($^+$) of a matrix, to account for the fact that W_M contains a kernel, which consists of the constant functions.

4.1. Influence of the choice of basis

The expression above is valid for any choice of functional bases. However, in certain specific cases, it can be simplified further. Namely, suppose that the f is expressed in some basis such that if the basis functions are stored as columns of the matrix B_N we have: $B_N^T H_N B_N = Id$. In that case, we have for any pair of functions, $\langle f, g \rangle_{H_N} = \mathbf{f}^T B_N^T H_N B_N \mathbf{g} = \mathbf{f}^T \mathbf{g}$. For example if the map C is written in the Laplace-Beltrami eigenbasis Φ chosen on each shape s.t. $\Phi^T A \Phi = Id$ and $W \Phi = A \Phi \Delta$ on each shape, where A is the area matrix, and Δ is the diagonal matrix of eigenvalues then:

$$X^A = C^T, \text{ and } X^C = \Delta_N^+ C^T \Delta_M. \quad (10)$$




Figure 2: Comparison of the quality of the map obtained using the regular functional maps approach with the one obtained with adjoint regularization based on the same descriptors.

5. Adjoint Regularization for Map Computation

To illustrate the relation between the map inverse and its adjoint, we consider a simple modification to the functional map estimation pipeline, proposed in the original article [OBCS*12]. Namely, we compute a functional map C between shapes M and N by solving the optimization problem (see also Section 2.5 in [OCB*16]):

$$C_{\text{opt}} = \arg \min_{C_{M,N}} \|C_{M,N}F - G\|^2 + \alpha \|\Delta_N C_{M,N} - C_{M,N} \Delta_M\|^2.$$

Here, F and G are matrices storing in each column the coefficients of corresponding descriptor functions in the basis of Laplace-Beltrami eigenfunctions, and Δ_M, Δ_N are the diagonal matrices of eigenvalues of the LB operators. As has been remarked in several works before, e.g., [ERGB16], this formulation is not symmetric with respect to the shapes M and N . Therefore, we consider the problem of solving for the two maps $C_{M,N}$ and $C_{N,M}$ jointly, first proposed in [ERGB16]. In our approach, we couple the two problems via the two adjoint operators introduced above. This leads to the following optimization problem:

$$\begin{aligned} \min_{C_{M,N}, C_{N,M}} & E_1(C_{M,N}) + E_2(C_{N,M}) + E_3(C_{M,N}, C_{N,M}), \\ E_1(C_{M,N}) &= \|C_{M,N}F - G\|^2 + \alpha \|\Delta_N C_{M,N} - C_{M,N} \Delta_M\|^2, \\ E_2(C_{N,M}) &= \|C_{N,M}G - F\|^2 + \alpha \|\Delta_M C_{N,M} - C_{N,M} \Delta_N\|^2, \\ E_3(C_{M,N}, C_{N,M}) &= \beta \|C_{M,N} - C_{N,M}^T\|^2 + \gamma \|\Delta_N C_{M,N} - C_{N,M}^T \Delta_M\|^2, \end{aligned}$$

where β, γ are scalar regularization parameters. Note that the coupling energy E_3 involves two terms, which approximate the map inverse using its area-based and conformal adjoint operators respectively. I.e., using the notation above, we can rewrite: $E_3(C_1, C_2) = \beta \|C_2 - X^A(C_1)\|^2 + \gamma \|C_2 - X^C(C_1)\|^2$. These two terms give an explicit control over promoting area-preservation and conformality as well as overall invertibility of the map.

We evaluated this modification on 100 random pairs of shapes in the FAUST dataset [BRLB14], where we used a small number of descriptor functions (20 Wave Kernel Signature [ASC11] descriptors for a map represented with 60×60 eigenfunctions). After estimating the maps $C_{M,N}, C_{N,M}$ jointly, we then evaluated $C_{M,N}$ using the same protocol as done in [KLF11, OBCS*12] and compared to the baseline standard approach. Note that even though we used the inverse map $C_{N,M}$ together with the coupling to estimate $C_{M,N}$, we did not use $C_{N,M}$ for pointwise map computation




Figure 3: The fraction of coverage of the target shape with the maps computed using the regular pipeline, Coupled Maps of [ERGB16] and adjoint regularization after ICP [BM92] post-processing. Before ICP, the coverage is 5%, 38.3% and 49.8% respectively.

or evaluation in any way. Thus, the adjoint coupling is only used for regularization. Since we considered a particularly challenging case with a relatively small number of descriptors, the basic approach of [OBCS*12] fails to produce good results. On other hand, as shown in Figure 1, our simple regularization can produce a significant improvement in the map quality without using any additional information, and while still maintaining the linear (least squares) nature of the optimization problem. We also compared our regularization with the Coupled Functional Maps approach of [ERGB16] using exactly the same input descriptors. Unlike the coupled approach which enforces orthonormality as a hard constraint, here we consider both the area-based and conformal adjoint operators together as soft constraints, since in our experiments both β and γ play a role and activating both terms improves results over using only one. We also notice that the framework of [ERGB16] enforces invertibility directly using a penalty $\|C_{M,N}C_{N,M} - Id\|$, which is unlikely to hold in the reduced basis and leads to a non-convex problem. Contrastingly, our regularization results in a simple convex problem (quadratic in $C_{M,N}, C_{N,M}$) so that we can find a global optimum.

It is also worth noting that the regularization terms weighted by α in $E_1(C_{M,N})$ and $E_2(C_{N,M})$ are used so that the resulting functional maps represent approximately isometric maps, which in practice promotes diagonal structure in the functional map matrices. In the same spirit, the authors of [ERGB16] propose a term which penalizes the output functional maps with large off-diagonal entries. To control for this difference, we compare the results of our method with the ones obtained by [ERGB16] after removing the respective regularization terms favoring the diagonal-structure output. As we show in Figure 10, the performance of [ERGB16] drops significantly, while ours remains almost the same.

To illustrate the effect of using our adjoint regularization, we also visualize the map that we obtained using the standard functional maps approach with ICP compared to the one obtained with our scheme in Figure 2. Note that the map obtained with our approach is significantly more continuous and consistent even though the computation uses exactly the same descriptor functions. We also plot the percentage of area covered by the pointwise map obtained using different approaches on the same pair of shapes in Figure 3. Here we plot a binary function on the target shape, which equals to

1 if and only if a particular point is used as a target point of some point from the source shape. Intuitively this function demonstrates how invertible the map is. Note again that although our adjoint regularization is linear in nature and the conversion from a functional map to a point-to-point map is done using a simple nearest-neighbor scheme we achieve a significant increase in map invertibility. Using more advanced pointwise map recovery schemes such as [RMC15] this can certainly be improved further.

We also performed the same test on 76 pairs of shapes in the TOSCA [BBK08] and 71 pairs of shapes in the SCAPE [ASK*05] datasets, while keeping all parameters the same throughout all experiments: $(\alpha, \beta, \gamma) = (0.01, 5000, 1.13)$. In terms of runtime, both our method and the method of [ERGB16] take between 5 and 15 seconds to process (including the ICP refinement) a pair of shapes on the FAUST, SCAPE and TOSCA datasets, depending on the shape complexity, on a machine with 3.3GHz i5 processor and 32GB memory. The corresponding results are shown in Figures 11 and 12 in the appendix. Note that although the amount of improvement varies, our regularization consistently helps to improve the map quality, especially in poorly mapped areas and reduces the number of significantly incorrect matches.

6. Consistency of Adjoint Functional Maps

In this section, we consider the problem of map analysis and visualization. Given a pair (or a collection) of shapes and the corresponding functional maps, our aim is to detect and highlight the regions on the shapes that undergo significant deformation.

6.1. A Brief Overview of Previous Approaches

A previous approach towards this problem has been taken in [OBCCG13], which can be seen as a special application of the shape difference operators [ROA*13]. For the sake of generality, we below review the previous frameworks in terms of the shape difference operators.

Given a pair of shapes M, N and the functional map $C : L^2(M) \rightarrow L^2(N)$, the associated shape difference operator $D_{M,N}^A$ is a linear operator mapping functions from $L^2(M)$ to itself. In [OBCCG13], the authors propose a functional measuring the area distortion induced by C between M and N , written as:

$$E_M(f) = \frac{\int_M f D_{M,N}^A(f) d\nu_M}{\int_M f^2 d\nu_M} \quad (11)$$

As discussed in [OBCCG13], $E_M(f)$ is large (resp. small) whenever there is an area expansion (resp. contraction) from the support of f on M to the support of $C(f)$ on N . The optimizers of $E_M(f)$ are then expected to detect and highlight the regions on M that undergo significant area deformations. Note that the same functional can be written with respect to the conformal shape difference $D_{M,N}^C$.

By adding a regularity condition $\int_M f^2 d\nu_M = 1$, in the discrete setting, maximizing (resp. minimizing) $E_M(f)$ boils down to computing the largest (resp. smallest) eigenvalue. Furthermore, the eigenfunctions of $D_{M,N}^A$ are computed as a collection of multi-scale *highlighted functions*, each of which corresponds to an eigenvalue of $D_{M,N}^A$. Typically, users need to check the highlighted functions

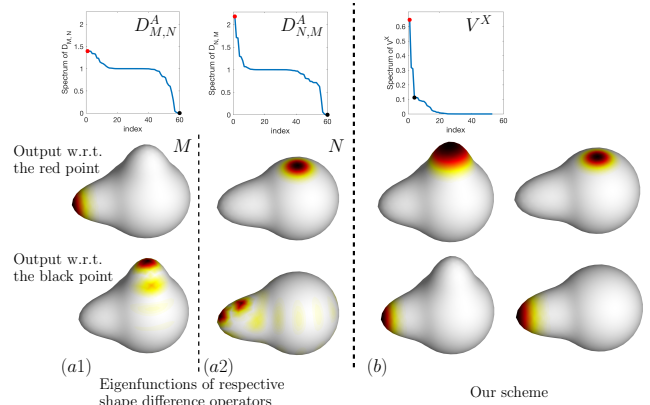


Figure 4: Comparison of a pair of deformed spheres with the shape difference operators and with our scheme. In column (a1) and (a2), we demonstrate the spectrum and the eigenfunctions of $D_{M,N}^A$ and $D_{N,M}^A$ corresponding to the largest (red point) and smallest eigenvalues (black point), respectively. In column (b), we show the results obtained by our scheme. Unlike the previous frameworks, our method produces a pair of functions at each point in the spectrum, and reveals the fact that the deformation at the top is more significant than the one on the left.

with respect to the largest and the smallest several eigenvalues to get a general picture of how the area deformations from M to N take place.

The major limitation of this framework is that its output relies on the base shape on which we construct the shape difference operator, i.e., if we choose N as the base shape, the resulting functions would be supported on N and the notions of contraction or expansion are measured with respect to its inner products.

To illustrate this, consider a pair of deformed spheres as shown in Figure 4. Intuitively, by comparing them, we expect to capture the area expansion on the left and the area contraction at the top, from M to N . If we single out one of them as the base shape, say, N , then we need to check the highlighted functions corresponding to both ends of the spectrum. However, as shown in the bottom row, the function in the second panel is noisy due to the low-rank approximation of the functional map. In practice, as the functional maps are computed with truncated basis, they are not consistent even if they are induced by consistent pointwise maps. On the other hand, one can as well consider both the eigenfunctions corresponding to the largest eigenvalues of $D_{M,N}^A$ and $D_{N,M}^A$ respectively, which are relatively more stable (see the middle row of the left part of Figure 4). Nevertheless, this brings in another problem — in this case the functions are living in distinct function spaces and measured by different functionals, we can not compare the deformations captured by the obtained functions. In other words, it is not justified to compare directly the spectrum plots on the top of column (a1) and (a2).

6.2. Consistency of Adjoint Representations

Motivated by the problems posed in Section 6.1, in this part we propose a new scheme, which exploits the informativeness of the adjoint operator and provides a unified way to jointly detect and

highlight the deformations in a collection of shapes. Instead of constructing some functional measuring distortion based on a single map, we incorporate the adjoint representation in the framework on cycle-consistency of functional maps [WHG13, HWG14].

Now assume that we are given a collection of shapes $\{S_i\}_{i=1}^n$, and a collection of cycle consistent bijections $T_{ij} : S_i \rightarrow S_j$, which satisfy $T_{i_k i_1} \circ T_{i_{k-1} i_k} \circ \dots \circ T_{i_1 i_2} = T_{i_1, i_1}, \forall \{i_1, i_2, \dots, i_k\} \subset \{1, 2, \dots, n\}$. Let C_{ij}, X_{ij} be the functional map and the adjoint representation from $L^2(S_i)$ to $L^2(S_j)$, induced by the consistent maps $\{T_{ij}\}_{1 \leq i, j \leq n}$. First of all, we justify in the continuous case the exact cycle consistency of X_{ij} .

Proposition 6.1 *Given a collection of shapes $\{S_i\}_{i=1}^n$, and consistent bijections $T_{ij} : S_i \rightarrow S_j, \forall 1 \leq i, j \leq n$, let X_{ij} (either the area-based or conformal one) be the adjoint representation induced by T_{ij} , then $\{X_{ij}\}_{1 \leq i, j \leq n}$ satisfy the cycle-consistent condition below.*

$$X_{ii} = Id_i, \quad (12)$$

$$X_{ji} X_{ij} = Id_i, \quad (13)$$

$$X_{ki} X_{jk} X_{ij} = Id_i. \quad (14)$$

Proposition 6.1 implies that the adjoint representations can be adapted in the consistency framework proposed in [WHG13], where the cycle consistency is posed as a soft constraint and optimized for using an iterative procedure.

As we mentioned in Section 4, in practice, instead of $L^2(S_i)$, we consider a reduced m -dimensional function space on S_i , meaning that any function in the reduced space is represented by a m -dimensional vector and each C_{ij} is an m by m matrix. Now given n functions $\{f_i\}_{i=1}^n$, one on each of the shapes, their consistency with respect to the adjoint operators can be measured by:

$$E_{cons}(f) = \sum_{i,j} \|X_{ij} f_i - f_j\|^2, \text{ where } f = (f_1^T, f_2^T, \dots, f_n^T)^T \quad (15)$$

Following the same argument as in [WHG13], we can show that $E_{cons}(f) = f^T V^X f$, where V^X is a square matrix of dimension mn , and its (i, j) -th block is given as follows:

$$V_{ij}^X = \begin{cases} \sum_j (I + X_{ij}^T X_{ij}) & \text{if } i = j, \\ -(X_{ji} + X_{ij}^T) & \text{otherwise.} \end{cases} \quad (16)$$

Therefore, after adding a regularity condition $f^T f = 1$ (to avoid trivial solutions), finding a set of consistent functions with respect to a collection of given adjoint operators maps amounts to computing the eigenvectors of V^X corresponding to the smallest several eigenvalues.

Although the procedure described above provides a simple way to construct the most and least consistent functions with respect to the adjoint operators X_{ij} , a more interesting observation we leverage here is that functions consistent with respect to the functional maps are not necessarily consistent with respect to the adjoint representations at the same time. In fact, consider n functions f_1, f_2, \dots, f_n such that they are consistent with respect to both $\{C_{ij}\}$ and $\{X_{ij}\}$, i.e., $C_{ij} f_i = f_j$ and $X_{ij} f_i = f_j, \forall i, j$. Then, according to Eq. 4 and Eq. 5, it implies that $D_{ji} f_j = f_j$, which in turns suggests that f_j is supported in a region on S_j undergoing little deformation.

Algorithm 1: Highlighting Jointly Area-based Deformed Regions

input : n shapes $\{S_i\}_{i=1}^n, C_{ij} : L^2(S_i) \rightarrow L^2(S_j)$, and a parameter ϵ .

output: A collection of functions on each of the shapes highlighting the jointly deformed regions.

- (1) Set the area-based adjoint operators $X_{ij} = C_{ji}^T$;
- (2) Construct the (i, j) -th block of W^C as follows:

$$W_{ij}^C = \begin{cases} \sum_j (I + C_{ij}^T C_{ij}) & \text{if } i = j, \\ -(C_{ji} + C_{ij}^T) & \text{otherwise.} \end{cases} \quad (17)$$

and compute the eigenvalues of W^C no larger than ϵ and collect the corresponding eigenvectors as a matrix $F \in \mathbb{R}^{mn \times k}$,

- (3) Construct V^X with respect to X_{ij} , compute the eigenvectors of $F^T V^X F$ and sort them in descending order as b_1, b_2, \dots, b_k ;
 - (4) Lastly, for j from 1 to k , let $F b_j = (f_1^T, f_2^T, \dots, f_n^T)^T$ and visualize the function f_i on S_i .
-

Thus, in order to obtain functions that are supported on significantly deformed regions, we propose to find functions that are (1) consistent with respect to $\{C_{ij}\}$ and (2) the least consistent with respect to $\{X_{ij}\}$. Such functions can be obtained by first learning a set of functions that are consistent with respect to $\{C_{ij}\}$ using the framework [WHG13], and then searching in the subspace spanned by these functions for the ones that are the least consistent with respect to $\{X_{ij}\}$. The consistency measurement is given as a quadratic form with respect to V^X . We describe how to implement the idea in Algorithm 1. It is worth noting that our scheme provides a unified measurement of the inconsistency of the produced highlighted functions, which are equivalently the eigenvalues of $F^T W^X F$, where F contains the consistent functional bases with respect to the functional maps (see step 2 in Algorithm 1), and at each eigenvalue, instead of producing a single function on a specified shape, our scheme produces a set of functions, one for each shape in the collection.

We compare our result with the one based on shape difference operators directly: as shown in column (b) of Figure 4, our result captures both of the deformed regions jointly. Perhaps more remarkably, our result reveals the fact that the deformation at the top is more significant than the one on the left (as the red point higher than the black point in the spectrum), though the two deformations would be considered in the opposite directions from the point of view of shape difference operators.

7. Pointwise Recovery with the Adjoint Representation

In this section, we consider another fundamental problem in shape analysis with the functional approaches. Namely, recovering a pointwise map from a given functional map $C : L^2(M) \rightarrow L^2(N)$.

The original approach proposed for this task in [OBCS*12] is to construct heat kernel functions for infinitesimal time t , and for

each heat kernel function $k_t(y, \cdot)$ on N an $x \in M$, such that $Ck_t(x, \cdot)$ is the nearest to $k_t(y, \cdot)$ in L^2 -distance. When the functional map is expressed in the Laplace-Beltrami basis, this can be done by simply finding for each row of Φ_N the nearest neighbor among the rows of $(C\Phi_M^T)^T = \Phi_M C^T$. A slight variation of the same idea would be to find for each row of $A_N \Phi_N$ the nearest neighbor in the rows of $A_M \Phi_M C$, which would correspond to approximating Dirac δ -functions. The authors of [OBCS*12] also propose an ICP-based technique in the embedded functional space as a refinement on a given functional map by iterating between finding the pointwise map and projecting the functional map C to the nearest orthonormal matrix, forcing $C^T C = Id$.

In general, among the difficulties of recovering a pointwise map from a given functional map, two are significant: one is the information loss when representing the functional map in truncated functional basis; the other is the potential non-isometric distortions, which can make it difficult to align functions across the shapes. Inspired by the properties of the adjoint, we propose a modification of this scheme, which takes into account the deformation captured by the difference between the adjoint representation and the functional map (as described in Proposition 3.1), which alleviates the latter obstacle and can thus be applied in the more difficult non-isometric scenarios. Our key observation is that, we can modify the functional inner product on one of the shapes to make the two functional spaces more comparable. Furthermore, this modification leads to a simple change of basis under which the orthonormality $C^T C = Id$ is meaningful even for non-area preserving maps.

Below, we first give a description for recovering the pointwise map associated with the functional map $C_{M,N}$ under the assumption that an approximation of the inverse map $C_{N,M}$ is given and then show how this assumption can be lifted.

Our method proceeds in the following three steps:

1. We change the inner product on the function space on N , so that new inner product $\tilde{H}_N = C_{N,M}^T H_M C_{N,M}$;
2. We modify the function basis on shape N to make it orthonormal with respect to \tilde{H}_N .
3. We then apply the same technique as described in the original article [OBCS*12].

We observe that after applying the first two steps described above the functional map $C_{N,M}$ expressed in the new basis will equal to the adjoint of $C_{M,N}$ and therefore the latter will satisfy the orthonormality property $C_{M,N}^T C_{M,N} = Id$.

When the inner product H_M is the simple L^2 inner product and the maps are expressed in the orthonormal LB basis, then we obtain $\tilde{H}_N = C_{N,M}^T C_{N,M}$ (since $H_M = Id$ in the standard basis Φ_M). To construct a new basis we solve the generalized eigenvalue problem $\Delta_N \phi = \tilde{H}_N \phi$ where Δ_N is the diagonal matrix of eigenvalues of the Laplace-Beltrami operator on N . We then store the vectors ϕ as columns of a matrix e and set the new basis to $\tilde{\Phi}_N = \Phi_N e$.

In the following proposition, we demonstrate the difference between the nearest-neighbor search (NN-search) before and after applying the change of basis.

Proposition 7.1 *Let M, N be two discrete surfaces each consisting*

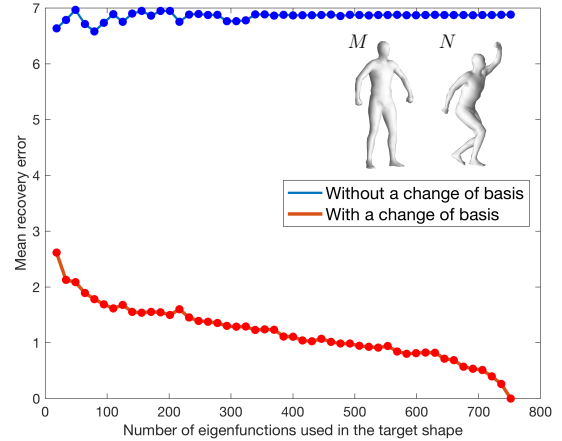


Figure 5: Using a fixed (3-dimensional) eigenbasis on shape M , we compared the nearest neighbor search with the original eigenbasis on N and the modified one in our scheme. As shown in the plot, as the dimension of the eigenbasis on N increases, the recovery error of our scheme vanishes while the old scheme fails to improve.

of n vertices, and a bijective map $T : N \rightarrow M$. Given a functional map $C_{M,N}$ induced by T with eigenbases $\Phi_M \in \mathbb{R}^{n \times k_M}$ and $\Phi_N \in \mathbb{R}^{n \times k_N}$. Assume that $3 \leq k_M \ll n$ is fixed. Then the recovery error converges to 0 as k_N increases to n after applying the change of basis, while it does not converge in the NN-search scheme.

To illustrate this difference, we compared a pair of shapes M, N , using only 3 eigenfunctions on the source shape M , and increased the number of basis functions in the target shape N . Figure 5 shows the mean recovery errors (in Euclidean distance) with respect to the increasing number of eigenfunctions on N . It is obvious that the error regarding our new scheme decreases and reaches zero error when the full basis on N are taken. Meanwhile, the old scheme without changing basis does not improve as the size of Φ_N grows.

In the discussion above we assumed that the inverse map $C_{N,M}$ is known (indeed, in the previous example we used $C_{N,M}$ in the full basis). One possible way to $C_{N,M}$ is to directly compute the inverse of $C_{M,N}$. However, in the presence of significant deformations between M and N , this does not lead to reasonable results.

Instead, according to Proposition 3.1, we propose to optimize for a semi positive-definite matrix \tilde{H} such that $\|X_{M,N}^A - \tilde{H} C_{M,N}\|_{Fro}$ is minimized. Therefore, the problem reduces to computing $X_{M,N}^A$ given $C_{M,N}$. On the other hand, thanks to Proposition 3.2, we can find a function κ such that $X_{M,N}^A(\kappa \cdot f) = C_{M,N}$, and as proven in Proposition 3.2, $\kappa = X_{N,M}^A(\mathbf{1}_N) = C_{M,N}^T(\mathbf{1}_N)$. To conclude, we first use the given $C_{M,N}$ to estimate $X_{M,N}^A$, then solve a convex optimization problem $\tilde{H} = \min_{H \in \text{SPD}} \|X_{M,N}^A - H C_{M,N}\|_{Fro}$, and finally compute e such that $e^T \tilde{H} e = Id$. In the end, direct computation shows that the modified basis on N is $\tilde{\Phi}_N = \Phi_N e$ and the functional map in the new basis is $\tilde{C}_{M,N} = e^{-1} C_{M,N}$.

8. Additional Experimental Results

As described in the previous sections, the adjoint operators allow us to jointly detect and visualize the deformations among a collection of shapes, and to develop a simple but efficient algorithm for map recovery. In this section, we demonstrate respectively how these two methods work in practice. In both parts, we assume that we are given a pair or a collection of shapes, and the functional maps between them, which are computed from some given point-to-point map. Typically, the input functional maps are represented by matrices of dimension 60×60 .

8.1. Joint Map Analysis and Visualization

We first verify the relevance of our approach by comparing the results generated by the previous method based on computing eigen-decompositions of shape difference operators and by ours in the pairwise setting. Each of the methods produces a collection of eigenfunctions f_i with respect to eigenvalues λ_i , we demonstrate in Figure 6 a weighted function $f = \sum_i \lambda_i f_i^2$ for compact visualization. In Figure 6(a), we plot the weighted eigenfunctions with respect to the largest 5 eigenfunctions of $D_{M,N}^A$ (resp. $D_{N,M}^A$) on M (resp. N). While in Figure 6(b), we plot f combining the eigenfunctions associated to the largest 10 eigenvalues of V^X . As we can see, in Figure 6(a), the highlighted areas on M and N are inconsistent, while in Figure 6(b), our method produce more consistent highlighted functions across M and N . Moreover, our result detects all of the areas highlighted separately by eigenfunctions of $D_{M,N}^A$ and of $D_{N,M}^A$. (see Figure 13 in the appendix for a detailed version of the same experiment).




Figure 6: (a) The weighted sum of eigenfunctions associated to the largest 5 eigenvalues of $D_{M,N}^A$ on M , and that of $D_{N,M}^A$ on N ; (b) The weighted sum of the 10 eigenfunctions produced by our method. We demonstrate both the front and the back of each shape for a complete view of the highlighted areas.

As shown in Algorithm 1, the only parameter of our method is ϵ , which controls the consistency of the resulting functions across shapes in the collection. In our implementation, we manually choose ϵ to be the largest value that is smaller than the k^{th} eigenvalue of W^C at which a significant deviation from zero occurs. Choosing a larger ϵ typically increases the inconsistency of the highlighted functions with respect to the functional maps, while allowing a larger functional subspace. This means that more distorted areas can be captured at the expense of potential loss of consistency.

In Figure 7, we tested our algorithm on two collections of shapes:

humans in 5 different poses (each consisting of 12500 vertices) and horses in 4 varying gaits (each consisting of 8431 vertices). In each case, from top to bottom, we plotted the eigenfunctions with respect to the largest 5 eigenvalues of V^X . As can be seen, the plotted functions highlight the joint, knee, chest and shoulders of the human poses and the joint and hip of the galloping horses, which are evidently deformed in the respective collection.

8.2. Pointwise Recovery

In this part, we compared our method for pointwise map recovery introduced in Section 7 with the nearest neighbor search proposed in [OBCS*12] and a more recent framework of [RMC15].

In general, assume we are given a pair of shapes M, N , the ground truth map $T : N \rightarrow M$, and the functional map $C_{M,N}$ computed with T . As mentioned at the beginning of this section, we fixed the dimension of $C_{M,N}$ as 60×60 , meaning that on each shape we truncated the first 60 eigenfunctions of the respective Laplace-Beltrami operator. We first evaluated the different schemes in a synthetic data set consisting of 64 deformations of the unit sphere deformed by adding two protrusions of varying sizes. The sizes of protrusions sample an evenly spaced 2-dimensional grid of values. Each deformed sphere consists of 1922 vertices. We randomly sampled 25 pairs of these deformed spheres and then evaluated the three schemes on them. In particular, after computing $\tilde{C}_{M,N}, \tilde{\Phi}_N$ with the above procedure, we only used the first 20 columns of $\tilde{C}_{M,N}$ (and accordingly the first 20 eigenfunctions in Φ_M) for the nearest neighbor search in the modified basis. The average map recovery accuracy is plotted in Figure 8. In this simple case, our scheme achieved a remarkably high recovery accuracy, in fact, in 24 out of 25 pairs we obtained exact map recovery.

Then we conducted tests on four pairs of more sophisticated data: elephant vs horse (each consists of 1199 vertices), conformal bunnies (each consists of 14290 vertices), cat vs lion (7207 vs 5000 vertices), and deformed faces (each consists of 6918 vertices). Since the geometric structures of these shapes are obviously more complicated than the ones of the synthetic deformed spheres, we use more eigenbasis functions on M – set to 40 in our scheme. All these three pairs are non-isometric shapes, therefore it is in general more challenging to recover the pointwise map with the given low-rank approximation of the functional maps. In Figure 9, we plotted the recovery accuracy curves in the left column, and to its right, we demonstrated the pairs in comparison. In particular, as we recovered a map from N to M , we computed for each vertex on N how far its matching point is from the ground truth on M , and we plotted the error distribution on N for the results coming from the iterative method of [RMC15] and from ours. In general, our result shows an improvement over the naive nearest neighbor search and is comparable with that from the more advanced approach. It is worth noting that our method is more computationally efficient than the one of [RMC15], since ours only consists of one step for changing the basis and one step for nearest neighbor search while the latter uses an iterative procedure to refine the output map. Our method takes around 3 seconds for each of the examples in Figure 9, while the procedure of [RMC15], using 10 iterations, takes up to 63 seconds (e.g., on the cat vs. lion example). However, the

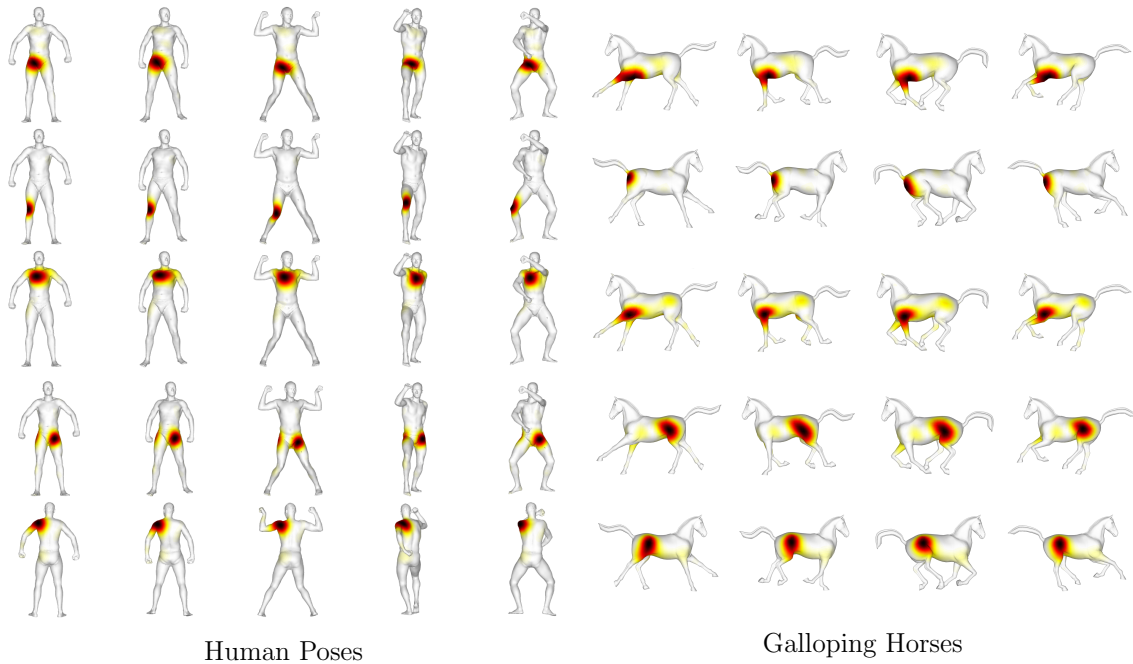


Figure 7: We tested our method for jointly highlighting the deformed regions in two collections of shapes and plotted the eigenfunctions corresponding to the largest 5 eigenvalues of V^X respectively. In each case, the eigenfunctions capture the deformed regions across the respective collection without relying on a single base shape. Thus, the regions detected on each shape are consistent.




Figure 8: We evaluated our scheme on 25 random pairs of deformed spheres. Note that the recovery accuracy of our scheme outperforms both the naive nearest neighbor search and the framework of [RMC15] in this simple synthetic dataset.

computational ease can come at a price, for instance, in the second row of Figure 9, the errors aggregate at the tip of the bunny's ear.

9. Conclusion and Future Work

In this paper, we argue that adjoint operators of functional maps provide a tool that can be used within a variety of map analysis and processing scenarios. In particular, we establish a connection between the adjoint operators, functional maps, their inverses and shape difference operators. We demonstrated the utility of the adjoint operators for (1) adding regularity to bi-directional shape

matching procedure; (2) encoding and revealing the distortions within shape collections (3) pointwise map recovery via a basis-modification scheme.

We believe that the use of functional map adjoints can open the door to many other tasks, such as accurate map estimation in shape collections by exploiting consistency between distortions, and even finding correspondences between non-isometric shape pairs by directly estimating X rather than C , as they allow to incorporate and reveal shape deformation in a compact manner.

Acknowledgement The authors thank the anonymous reviewers for their valuable comments. This work is supported by Marie-Curie CIG-334283, a CNRS chaire d'excellence, chaire Jean-Marjoutet from Ecole Polytechnique, FUI project TANDEM 2, and a Google Focused Research Award.

References

- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV Workshops* (2011), pp. 1626–1633. 5
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: SCAPE: Shape Completion and Animation of People. In *ACM Transactions on Graphics (TOG)* (2005), vol. 24, ACM, pp. 408–416. 6
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized Multidimensional Scaling: A Framework for Isometry-Invariant Partial Surface Matching. *PNAS* 103, 5 (2006), 1168–1172. 2
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer Science & Business Media, 2008. 6

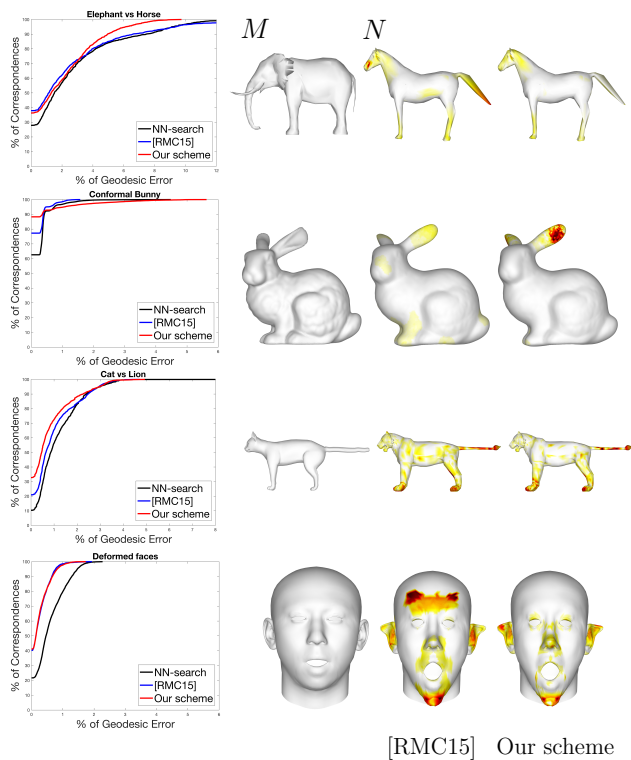


Figure 9: The left-most column shows the recovery accuracy curves. The shapes in comparison are plotted to its right, we computed the pointwise recovery errors on N as shown in the right-most two columns (the error increases as the color gets darker).

- [BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-d shapes. In *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1992), vol. 14, pp. 239–256. 5
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and Evaluation for 3d Mesh Registration. In *Proc. CVPR* (2014), pp. 3794–3801. 5
- [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. 2
- [dGDT16] DE GOES F., DESBRUN M., TONG Y.: Vector field processing on triangle meshes. In *ACM SIGGRAPH 2016 Courses* (2016). 2
- [ERGB16] EYNARD D., RODOLA E., GLASHOFF K., BRONSTEIN M. M.: Coupled functional maps. In *Fourth International Conference on 3D Vision (3DV), 2016* (2016), IEEE, pp. 399–407. 4, 5, 6, 12
- [GP00] GILES M. B., PIERCE N. A.: An introduction to the adjoint approach to design. *Flow, turbulence and combustion* 65, 3-4 (2000), 393–415. 2
- [GSTOG16] GANAPATHI-SUBRAMANIAN V., THIBERT B., OVSIANIKOV M., GUIBAS L.: Stable Region Correspondences Between Non-Isometric Shapes. In *Computer Graphics Forum* (2016), vol. 35, pp. 121–133. 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. 2, 3, 7
- [KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Computer Graphics Forum* 32, 2 PART4 (2013), 439–448. 2
- [KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In

Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2015), vol. 07-12-June-2015, pp. 905–914. 2

- [KGB16] KOVNATSKY A., GLASHOFF K., BRONSTEIN M. M.: MADMM: A generic algorithm for non-smooth optimization on manifolds. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (2016), vol. 9909 LNCS, pp. 680–696. 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended Intrinsic Maps. In *ACM Transactions on Graphics (TOG)* (2011), vol. 30, ACM, p. 79. 2, 5
- [KMP07] KILIAN M., MITRA N. J., POTTMANN H.: Geometric Modeling in Shape Space. In *ACM Transactions on Graphics (TOG)* (2007), vol. 26, ACM, p. 64. 1
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius Voting for Surface Correspondence. In *ACM Transactions on Graphics (TOG)* (2009), vol. 28, ACM, p. 72. 2
- [LRB*16] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D.: Non-Rigid Puzzles. *Computer Graphics Forum* 35, 5 (2016), 135–143. 2
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Fully spectral partial shape matching. *Computer Graphics Forum* 36, 2 (2017), 247–258. 2
- [MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In *Visualization and mathematics III*. Springer, 2003, pp. 35–57. 4
- [Mém11] MÉMOLI F.: Gromov–Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of computational mathematics* 11, 4 (2011), 417–487. 2
- [Mil08] MILLER D. A.: *Quantum mechanics for scientists and engineers*. Cambridge University Press, 2008. 2
- [MTPS04] MCNAMARA A., TREUILLE A., POPOVIĆ Z., STAM J.: Fluid control using the adjoint method. In *ACM Transactions On Graphics (TOG)* (2004), vol. 23, ACM, pp. 449–456. 2
- [NO17] NONGENG D., OVSIANIKOV M.: Informative descriptor preservation via commutativity for shape matching. In *Eurographics* (2017). 2
- [OBCCG13] OVSIANIKOV M., BEN-CHEN M., CHAZAL F., GUIBAS L.: Analysis and visualization of maps between shapes. *Computer Graphics Forum* 32, 6 (2013), 135–145. 2, 6
- [OBCS*12] OVSIANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31 (2012), 30:1—30:11. 1, 2, 3, 5, 7, 8, 9
- [OCB*16] OVSIANIKOV 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, 5
- [PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse Modeling of Intrinsic Correspondences. In *Computer Graphics Forum* (2013), vol. 32, pp. 459–468. 2
- [PP93] PINKALL U., POLTHIER K.: Computing Discrete Minimal Surfaces and their Conjugates. *Experimental mathematics* 2, 1 (1993), 15–36. 4
- [PRMH10] PRASSNI J. S., ROPINSKI T., MENSMAJN J., HINRICHS K.: Shape-based transfer functions for volume visualization. In *PacificVis* (2010), pp. 9–16. 1
- [RCB*16] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial Functional Correspondence. In *Computer Graphics Forum* (2016). 2
- [RMC15] RODOLÀ E., MÖLLER M., CREMERS D.: Point-wise Map Recovery and Refinement from Functional Correspondence. *VMV* (2015), 25–32. 2, 6, 9, 10

- [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* 32, 4 (2013), 1. 1, 2, 3, 4, 6, 13
- [RS81] REED M., SIMON B.: *I: Functional Analysis. Methods of Modern Mathematical Physics*. Elsevier Science, 1981. 2
- [SDGP*15] SOLOMON J., DE GOES F., PEYRÉ G., CUTURI M., BUTSCHER A., NGUYEN A., DU T., GUIBAS L.: Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 66. 2
- [SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft Maps Between Surfaces. In *Computer Graphics Forum* (2012), vol. 31, pp. 1617–1626. 2
- [SP04] SUMNER R. W., POPOVIĆ J.: Deformation Transfer for Triangle Meshes. *ACM Transactions on Graphics* 23, 3 (2004), 399–405. 1
- [SPKS16] SOLOMON J., PEYRÉ G., KIM V. G., SRA S.: Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 72. 2
- [TY99] TAKAGI H., YOKOUCHI K.: Multiplication and composition operators between two L^p spaces. In *Contemp. Math.* (1999), vol. 232, pp. 321–338. 3
- [VTSSH15] VON-TYCOWICZ C., SCHULZ C., SEIDEL H.-P., HILDEBRANDT K.: Real-Time Nonlinear Shape Interpolation. *ACM Transactions on Graphics (TOG)* 34, 3 (2015), 34. 1
- [WHG13] WANG F., HUANG Q., GUIBAS L. J.: Image co-segmentation via consistent functional maps. In *Proceedings of the IEEE International Conference on Computer Vision* (2013), pp. 849–856. 2, 3, 7, 13

Appendix A:




Figure 10: Map accuracy on 100 pairs of shapes in the FAUST dataset obtained using consistent maps approach of [ERGB16] and using adjoint regularization without constraints forcing diagonality of the resulting functional maps.

Proof of Theorem 3.1 For a fixed function $f_M \in L^2(M)$, we define a functional on $L^2(N)$: $S(g_N) = h_M(f_M, g_N \circ T)$. Note that $S(g_N)$ is a continuous linear functional on $L^2(N)$, therefore by the Riesz-Frechet Theorem, there exists a unique function w_N (for the fixed f_M) such that $S(g_N) = h_M(w_N, g_N)$. We then define $X_{M,N}(f_M) = w_N$, it is easy to verify that $X_{M,N}$ is linear in f_M , thus we prove the existence of a linear operator satisfying $h_N(X_{M,N}(f_M), g_N) = h_M(f_M, g_N \circ T)$. On the other hand, the uniqueness of $X_{M,N}$ is obvious, thus we finish the proof.




Figure 11: Map estimation accuracy on 71 pairs of shapes in the SCAPE dataset with the standard functional map pipeline vs. the coupled approach of [ERGB16] and our adjoint regularization.




Figure 12: Map estimation accuracy on 76 pairs of shapes in the TOSCA dataset with the standard functional map pipeline vs. the coupled approach of [ERGB16] and our adjoint regularization.

Proof of Proposition 3.1 We deal with the area-based case below, and the conformal case can be proven in the same way. It follows from the definition of the area-based shape difference operator that

$$\begin{aligned} h_N^A(D_{N,M}^A C_{M,N}(f_M), g_N) &= h_M^A(C_{N,M} C_{M,N}(f_M), C_{N,M}(g_N)) \\ &= h_M^A(f_M, C_{N,M}(g_N)). \end{aligned}$$

The second line is due to the assumption that $C_{M,N}$ is induced by T^{-1} , which means $C_{N,M} C_{M,N} = Id_M$. Then according to Theorem 3.1, $X_{M,N}$ is the unique linear operator satisfying $h_N^A(X_{M,N}(f_M), g_N) = h_M^A(f_M, C_{N,M}(g_N))$, thus it must be identical to $D_{N,M}^A C_{M,N}$.

Proof of Proposition 3.2 We define two functions on M, N respectively: $\rho_M(x) = \sqrt{\det(g_M(x))}$, $\forall x \in M$ and similarly $\rho_N(y) = \sqrt{\det(g_N(y))}$, $\forall y \in N$. The Riemannian measure ν_M on shape M is such that $d\nu_M = \rho_M d\lambda$, where λ is the Euclidean measure. We first




Figure 13: (a) The eigenfunctions of $D_{M,N}^A$ associated to the largest 5 eigenvalues; (b) The eigenfunctions of $D_{N,M}^A$ associated to the largest 5 eigenvalues; (c) The eigenfunctions associated to the largest 10 eigenvalues of V^X , produced by our scheme, as mentioned in Section 8.1, the detected areas in our scheme are consistent, so we only show the functions on M . Note that all the areas highlighted in the top row on two different shapes are captured in a specific order in the bottom row.

prove the following equation:

$$X_{M,N}^A \left(\frac{f}{\rho_M} \right) = \frac{f \circ T^{-1}}{\rho_N}. \quad (18)$$

Since $\langle \frac{f \circ T^{-1}}{\rho_N}, g \rangle_N = \int_N \frac{f \circ T^{-1}}{\rho_N} g dV_N = \int_N \frac{f \circ T^{-1}}{\rho_N} g \rho_N d\lambda = \int_N (f \circ T^{-1}) g d\lambda$. Following the same argument, we prove that $\langle \frac{f}{\rho_M}, g \circ T \rangle_M$ equals to $\int_M f (g \circ T) d\lambda$. Since T is a diffeomorphism, we know that $N = T(M)$. Thus by a change of variable and the definition of $X_{M,N}^A$, we prove Eq. 18.

Then we define a function κ on M , such that $(\kappa \rho_M) \circ T^{-1} = \rho_N$. Using Eq. 18, one can verify that $X_{M,N}^A(\kappa f) = C_{M,N}(f)$. On the other hand, since $X_{M,N}^A(\kappa) = X_{M,N}^A(\frac{\kappa \rho_M}{\rho_M}) = \frac{(\kappa \rho_M) \circ T^{-1}}{\rho_N} = 1_N$. According to Proposition 6.1 we prove later, we have κ is well-defined and $\kappa = X_{N,M}^A(1_N)$.

Proof of Proposition 6.1 First, it follows from Proposition 3.1 that $X_{ii} = D_{ii} C_{ii} = Id_i$. Secondly, according to the functionality of shape difference operators proven in [ROA*13], we have $D_{ij}^{-1} = C_{ji} D_{ji} C_{ij}$. Thus thanks to Proposition 3.1, $X_{ji} X_{ij} = D_{ij} C_{ji} D_{ji} C_{ij} = D_{ij} D_{ij}^{-1} = Id_i$ It has been shown in [WHG13] that the functional maps $\{C_{ij}\}$ satisfy the exact cycle-consistency, then we have

$$\begin{aligned} X_{jk} X_{ij} &= D_{kj} C_{jk} D_{ji} C_{ij} = D_{kj} C_{jk} D_{ji} (C_{ij} C_{ki}) C_{ik} \\ &= (D_{kj} C_{jk} D_{ji} C_{kj}) C_{ik} = D_{ki} C_{ik} = X_{ik} \end{aligned}$$

Thus $X_{ki} X_{jk} X_{ij} = Id_i$ (the last line is due to the functionality of the shape difference operators).

Proof of Proposition 7.1 Without loss of generality, we assume that the map between M and N be an identity map (thus in the following we do not differentiate f and $f \circ T$), so that $C_{M,N} = \Phi_N^T A_N \Phi_M$ and $C_{N,M} = Id$ expressed in the full basis, so that $\tilde{H}_N =$

A_M . Then, for any function f (including δ functions):

$$\begin{aligned} \|C_{M,N} \Phi_M^+ f - \Phi_N^+ f\| &= \|\Phi_N^T A_N (\Phi_M \Phi_M^T A_M) f - \Phi_N^T A_N f\| \\ &= \|D_{\Phi_N} \circ P_{\Phi_M} f - D_{\Phi_N} f\|, \end{aligned}$$

where P_{Φ_M} is an operator project a function onto the subspace spanned by Φ_M , and D_{Φ_N} is another projection operator by expresses the result in the basis of Φ_N . Now, in the new basis, since $\tilde{\Phi}_N^T A_M \tilde{\Phi}_N = Id$, we have $\tilde{C}_{M,N} = \tilde{\Phi}_N^T A_M \Phi_M$, and since $\tilde{C}_{M,N}$ is orthonormal, we get, after pre-multiplying by $\tilde{C}_{M,N}^T$:

$$\begin{aligned} \|\Phi_M^+ f - \tilde{C}_{M,N}^T \tilde{\Phi}_N^+ f\| &= \|\Phi_M^T A_M f - \Phi_M^T A_M (\tilde{\Phi}_N \tilde{\Phi}_N^T A_M) f\| \\ &= \|D_{\Phi_M} f - D_{\Phi_M} \circ P_{\tilde{\Phi}_N} f\|, \end{aligned}$$

This implies that in the new basis, finding nearest neighbors of $\tilde{\Phi}_N^T A_M$ among $\tilde{C}_{M,N} \Phi_M^T A_M$ is equivalent to finding nearest neighbors of $\tilde{C}_{M,N}^T \tilde{\Phi}_M^T A_M$ among $\Phi_M A_N$ and moreover, in this case, the order of projection is reversed: in the old scheme, one compares the projection onto the first the subspace spanned by Φ_M then onto Φ_N with the projection only onto Φ_N ; while in the new scheme, we are comparing the projection onto Φ_M with the composition between projection onto Φ_M and $\tilde{\Phi}_N$.

Now let $k_N = n$, the projection of f onto $\tilde{\Phi}_N$, $P_{\tilde{\Phi}_N}$, is an identity operator. Thus for any f , the projection error $\|D_{\Phi_M} f - D_{\Phi_M} \circ P_{\tilde{\Phi}_N} f\| = \|D_{\Phi_M} f - D_{\Phi_M} f\| = 0$. Contrastingly, in the old scheme, since $k_M \ll n$, $\|D_{\Phi_N} \circ P_{\Phi_M} f - D_{\Phi_N} f\| \neq 0$, unless f is spanned by the first k_M eigenfunctions Φ_M .