EUROGRAPHICS 2017 / L. Barthe and B. Benes
(Guest Editors)

Volume 36 (2017), Number 2

Informative Descriptor Preservation via Commutativity for Shape Matching

Dorian Nogneng Maks Ovsjanikov

LIX, Ecole Polytechnique

Abstract

We consider the problem of non-rigid shape matching, and specifically the functional maps framework that was recently proposed to find correspondences between shapes. A key step in this framework is to formulate descriptor preservation constraints that help to encode the information (e.g., geometric or appearance) that must be preserved by the unknown map. In this paper, we show that considering descriptors as linear operators *acting on functions through multiplication, rather than as simple scalar-valued signals, allows to extract significantly more information from a given descriptor and ultimately results in a more accurate functional map estimation. Namely, we show that descriptor preservation constraints can be formulated via* commutativity *with respect to the unknown map, which can be conveniently encoded by considering relations between matrices in the discrete setting. As a result, when the vector space spanned by the descriptors has a dimension smaller than that of the reduced basis, our optimization may still provide a fully-constrained system leading to accurate point-to-point correspondences, while previous methods might not. We demonstrate on a wide variety of experiments that our approach leads to significant improvement for functional map estimation by helping to reduce the number of necessary descriptor constraints by an order of magnitude, even given an increase in the size of the reduced basis.*

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—[Geometric algorithms, languages, and systems]

1. Introduction

In this paper we study the problem of non-rigid shape matching, which consists in trying to find a good correspondence between two shapes that might undergo a non-rigid transformation, such as articulated motion of humans. This problem has many applications such as deformation transfer [SP04], shape interpolation [KMP07] and even statistical shape modeling [HSS*09] among myriad others. A wide variety of methods has been used to tackle this problem over the years [VKZHCO11], primarily by restricting the search space either using feature-point correspondences [BBK06], or using a reduced model, such as conformal or isometric shape deformations.

In this paper, we concentrate on the functional map framework introduced in [OBCS*12], which has been widely adopted since its introduction due to its efficiency for representing and computing correspondences, which in the most basic case reduces to solving a linear system of equations. A key step in this framework, first introduced in the original article [OBCS*12] and then used in most follow-up works, including [PBB*13, COC14, LRB*16] among others, is to formulate function preservation constraints, which typically encode information (e.g., geometric or appearance) that must be preserved by the unknown map. These constraints are typically enforced simply by requesting that the function values must be globally preserved by the functional map. This means, however, that the constraints formulated using this approach often lead to underconstrained, badly defined optimization problems, especially when the number of linearly-independent descriptor functions is smaller than the number of basis functions, used to represent the map itself. This is especially problematic in the presence of non-rigid, possibly noisy deformations, for which obtaining a large set of informative, linearly independent descriptor functions can be very challenging.

Our main contribution is to notice that the standard approach for enforcing function preservation does not extract all of the information from a given descriptor. For example, the *level-sets* of the given function (i.e., the indicator functions of regions of constant value) are not necessarily preserved when using the basic function preservation constraint. This has two consequences: on the one hand, as mentioned above, this requires many descriptor functions to obtain a good approximation of a functional map, and on the other, perhaps more importantly, a solved-for functional map will not necessarily correspond to a point-to-point map, as it might not respect the "structural" properties of function preservation, as described in Section 5.1 in more detail. Indeed, one of our main motivations is to introduce constraints that would help guide the optimization process towards functional maps that are closer to point-to-point maps, without introducing additional computational complexity.

We show that much more information can be encoded into function (or descriptor) preservation constraints, while maintaining the overall *linear* system nature of the functional map framework, making it attractive from the computational standpoint. In particular, we show that when function preservation is encoded via *commutativity* with an underlying map, rather than simply via function value preservation, the resulting maps are both more accurate, and moreover can be obtained by using only a handful of descriptors (sometimes as few as 2-3), compared to hundreds required by the standard approach. We demonstrate on a wide range of experiments that our method helps to obtain better correspondences, largely removes the dependency of the descriptor number on the size of the reduced basis, and helps to obtain functional maps that are closer to point-to-point maps in a theoretically well-justified way.

2. Related work

Non-rigid shape matching is among the best-studied problems in digital Geometry Processing, and the full overview of related techniques is beyond the scope of our paper. Therefore, below we discuss the works that are most closely related to ours, especially those based on the functional map framework, and refer the interested reader to the recent surveys on shape matching for a more in-depth overview of the field [VKZHCO11].

Most early methods designed to find correspondences between shapes undergoing non-rigid transformation have concentrated on establishing mappings that minimize some distortion energy, such as conformality (locally angle preservation) [LF09, KLF11, APL15], or approximate intrinsic isometries (preserving geodesic distances) (e.g., [BBK06, TBW*09, OMMG10] among many others). Both the theoretical formalism and the computational methods associated with these approaches are mature and can often result in high-quality mappings whenever the deformations follow the prescribed models. However, such methods often lack flexibility, making it hard to introduce additional information, in the form of expected geometric or appearance properties (descriptors) that should be preserved by the map, and are badly-suited in the presence of more general non-rigid deformations.

Another, more recent, set of techniques has been proposed to obtain soft, or approximate correspondences rather than point-to-point maps [SNB*12,OBCS*12]. This includes both maps between probability densities [Mém11, SNB*12, SPKS16] and region-level maps [CK15, GSTOG16], which can be used in a multi-scale way to obtain accurate (sometimes even pointwise) correspondences. These techniques are often more robust in the presence of geometric and structural variability, and in many cases allow to inject domain-specific knowledge, such as expected descriptor preservation into the computational pipeline.

Most closely related to our work are the methods based on the functional map framework, initially introduced in [OBCS*12], and later extended significantly in follow-up works (e.g., [PBB*13, RCB*16, LRB*16] to name a few). These methods are based on the notion that it is often easier to obtain correspondences between functions, rather than points, by first using a reduced functional basis and second by formulating many *linear* constraints that allow to recover the functional map by solving a least squares system. An approach that tackles the problem of extracting a good point-to-point correspondence from a functional map can be found in [RMC15]. This framework has a particular advantage of being flexible and allowing to easily incorporate constraints including preservation of geometric quantities (descriptors), while at the same time being able to incorporate deformation models (e.g., isometries) via commutativity with various operators.

Despite this flexibility, one notable difficulty of using the functional map representation is that typically a large number of constraints is necessary to obtain a good solution. This includes using many descriptor preservation constraints [OBCS*12], even in the case of partial maps [RCB*16] (where for example, the authors use 352-dimensional descriptors). Unfortunately, obtaining a large set of high-quality robust and informative descriptor functions can be challenging [COC14], and moreover noisy descriptor functions can severely affect the resulting quality of the functional map. This is especially problematic since in the original formulation [OBCS*12], which has been also used in follow-up works, the number of descriptor preservation constraints is tightly linked to the size of the reduced basis, meaning that in order to obtain better correspondences more constraints are necessary, even in the absence of noise. Thus, several previous methods have tried to use regularization to improve the conditioning of the functional map computation, e.g., via sparsity [KBBV15].

In this paper we argue, that the previously proposed approach for function preservation constraints in the functional map framework does not extract all of the available information from a given function. By drawing a link between theoretical guarantees under which functional maps correspond to point-to-point maps, we show that it is possible to formulate the descriptor preservation constraints in a way that is both more informative, and results in higher quality functional maps even as the number of basis functions increase. Remarkably, we show that this is possible without sacrificing the overall linear least squares computational advantage of this framework. Our approach is general and can be used within any other method based on the functional map representation (e.g., [OBCS*12, PBB*13, RCB*16, LRB*16]), by simply changing the way that constraints are formulated and solved for.

To summarize, our main contributions include:

- A novel approach to formulating function (e.g., descriptor) preservation constraints within the functional maps framework.
- Theoretical analysis demonstrating that our method results in desired point-to-point maps, in the presence of perfect descriptors.
- Both theoretical guarantees and experimental evidence that our constraints allow to extract strictly more information from descriptor functions compared to previous approaches.

We evaluate our method on a wide variety of data, and show that using our simple modification can result in significant improvement in the quality of functional maps and reduce the number of necessary descriptor constraints by an order of magnitude.

3. Overview

The rest of the paper is organized as follows: in Section 4 we introduce the notation and give an overview of the basics of the functional maps framework and the associated pipeline for computing pairwise shape correspondences. Section 5 describes our proposed modification to this pipeline and discusses the main properties of the constraints that we introduce. We start by giving a general motivation and theoretical justification for our constraints in Section 5.1 and then describe how they can be introduced into the functional map estimation pipeline in Section 5.2. We list some of the properties of these constraints in Section 5.3, in particular proving that our approach *is strictly more informative* than the standard method for function preservation, and that it provably allows us to extract more information from the same given descriptors. Section 6 is dedicated to the experiments, which demonstrate that our constraints result in more accurate functional maps, and allow to obtain high quality maps with significantly fewer descriptors. Finally we conclude with Section 7 by mentioning some interesting challenges and directions for future work.

4. Overview of the Functional Maps Framework

In this section we describe the general setting of non-rigid shape matching, introduce the main notation that will be used in the rest of the paper, and give an overview of the functional map framework introduced in [OBCS*12], including the main computational steps required for estimating functional maps in practice. We concentrate in particular on the way that various constraints, and especially the function preservation constraints are encoded in this pipeline, as this forms the basis of our contribution described in the following Section 5.

4.1. Setup

The main goal in the problem of shape matching is to try to find a correspondence or a mapping between a pair of shapes M and N that represent similar physical objects, for which one would expect a natural correspondence to exist. For example, M and N can represent the same cat or human shape in two different poses. Throughout our discussion we will assume that in the continuous setting, these shapes can be modeled as smooth surfaces (two-dimensional manifolds) embedded in $\mathbb{R}^3$, and that in the discrete setting they are stored as triangle meshes, having n_M and n_N vertices respectively.

The simplest and most common approach is to represent a solution to the shape matching problem as a correspondence $T : N \to M$ that maps each vertex in N to a vertex in M according to some quality criterion. Such a correspondence can also be written as a matrix Π of size $n_N \times n_M$, that has exactly one 1 on each line, and zeros everywhere else. When the number of points $n_N = n_M$ and the map T is a bijection then Π is a standard permutation matrix. If we allow convex combinations of vertices (or equivalently probability distributions) as solutions, as in [SPKS16], then we can relax the binary $0,1$ constraint to allow the entries of Π to lie in the interval $[0,1]$ with the additional constraint that all lines of Π sum to 1.

Another, more general relaxation that was considered in [OBCS*12] is via linear mappings between real-valued functions defined on the shapes. Thus, given a function f defined on shape M, $f : M \to \mathbb{R}$, we can use T to transfer f onto N through composition to define $g = f \circ T$. Here, $g : N \to \mathbb{R}$ and $g(y) = f(T(y))$, for any point y on N. For any fixed T, the mapping between functions $f \mapsto g$ is linear, and thus can be represented as a matrix in the discrete setting. If functions are represented as discrete vectors, then using the notation above, we can simply write: $g = \Pi f$, if the functions are expressed as vectors with respect to the standard basis. Note that *any* real-valued matrix Π corresponds to a valid linear functional map, even if it does not represent a correspondence between points or probability distributions. This means that this representation is complete and no additional constraints have to be imposed on the estimated matrix.

The key aspect of the functional map representation proposed in [OBCS*12] is to use a reduced basis to encode the functional map, instead of working in the full spaces $\mathbb{R}^{n_M}$ and $\mathbb{R}^{n_N}$. Thus, suppose we are given some set of basis functions on shapes M and N, encoded as matrices Φ_M, Φ_N respectively, having sizes $n_M \times k_M$ and $n_N \times k_N$ for some $k_M \ll n_M$ and $k_N \ll n_N$, where each column corresponds to a basis function on the corresponding shape. Then, the *functional map* matrix can be written as $C = {\Phi_N}^{+} \Pi \Phi_M$, where $+$ denotes the Moore-Penrose pseudoinverse. For example, if the basis functions are orthonormal with respect to the standard inner product then $C = \Phi_N^T \Pi \Phi_M$, whereas if the basis functions are orthonormal with respect to a weighted inner product, so that $\Phi_N^T A_N \Phi_N = I_{n_N}$ where A_N is a matrix of weights, then $C = {\Phi_N}^T A_N \Pi \Phi_M$.

The expression above assumes that the initial correspondence matrix Π is known. In practice, however, the shape matching problem consists precisely in trying to recover this correspondence for a given pair of shapes. For this the authors of [OBCS*12] have proposed a pipeline, which was further extended in several follow-up works [PBB*13, COC14, RCB*16], and which consists of the following general steps (1) Computing a fixed set of the basis functions, Φ_M, Φ_N, (2) Estimating a set of pairs of descriptor (also called "probe") functions $f^{(p)}, g^{(p)}$, where $p \in \{1, ..., P\}$ such that the unknown functional map C should satisfy $C f^{(p)} \approx g^{(p)}$ for each p, (3) Solving the linear system of equations to estimate the unknown functional map C, and (4) Recovering a point-to-point correspondence from a given functional map matrix C. Note that the calculations in step (2) are done in the given functional basis, which significantly reduces the dimension of the considered space (the number of unknowns) and therefore makes the problem scalable, since the matrix C is of size $k_N \times k_M$.

In practice, the most commonly-used basis in functional map computations is given by the eigenfunctions, corresponding to the smallest eigenvalues of the Laplace-Beltrami operator, which has a multi-scale effect, since the eigenfunctions are ordered from low-frequency (smoothest) to high-frequency according to the eigenvalues. By far the most common discretization of this operator is the classical cotangent-weight scheme [PP93, MDSB03], which allows to represent it as a matrix $L = A^{-1} W$, where A is a diagonal matrix of area weights and W is a sparse matrix of cotangent weights. In this case the eigenfunctions can be found by solving the generalized eigenvalue problem $W\phi = \lambda A \phi$, and the matrix of eigenfunctions Φ satisfies the relation $\Phi^T A \Phi = Id$, and $\Phi^+ = \Phi^T A$.

Functional Map Estimation The key aspects in estimating functional maps therefore consists in formulating pairs of function preservation constraints $f^{(p)}, g^{(p)}$ and solving the linear system of equations to recover the unknown matrix C. The pairs of probe functions $f^{(p)}, g^{(p)}$ can represent descriptors such as Gaussian or mean curvature, or multi-scale descriptors such as the Heat or Wave Kernel signatures [SOG09, ASC11] for some range of parameter choices (i.e., each $f^{(p)}, g^{(p)}$ corresponds to a parameter, such as time in the HKS). Alternatively function preservation constraints can also represent knowledge of parts or feature points that are known to match, in which case the functions can be either indicators of given parts, or derived quantities, such as distance function to a feature. Once all of the function preservation constraints are computed, they can be stacked into matrices F, G whose corresponding columns represent the pairs of functions expressed in the given (Laplace-Beltrami) basis. Then, the optimal functional map is found by solving the following system in the least squares sense:

$$C_{\text{opt}} = \arg\min_{C} \|CF - G\|^2 + \alpha \|\Delta_N C - C\Delta_M\|^2. \quad (1)$$

Here, Δ_N, Δ_M are diagonal matrices of eigenvalues of the Laplace-Beltrami operator and α is a small scalar weight. In other words, the optimal functional map C can be computed so that it preserves the given functions and commutes with the Laplace-Beltrami operator itself. This latter constraint is associated with the standard assumption that the sought map should be approximately intrinsically isometric.

Limitations Although simple and efficient, the pipeline described above has several limitations: first the number of linearly independent function preservation constraints must be sufficiently high to ensure that the least squares system leads to a good approximation of the functional map. Without additional assumptions, such as sparsity, in most cases this implies that the number of descriptors must be approximately equal to the number of basis functions (which typically ranges between 80-100). Unfortunately, obtaining a large number of descriptor functions that are robust, informative and linearly independent can often be difficult. Moreover, as described below, this basic method for enforcing descriptor preservation does not extract the full information from the given functions, leading to sub-optimal results. Finally, and perhaps most importantly, this approach does not have constraints or regularizers that would lead to the solution to point-to-point maps, which can affect the overall accuracy of the correspondence estimation pipeline.

5. Novel Approach for Functional Correspondences

5.1. Motivation

One of the primary motivations behind our approach to function preservation within the functional maps framework is a classical result that states that any non-trivial linear functional map C corresponds to a point-to-point map if and only if it preserves pointwise products of functions $C(f \cdot h) = C(f) \cdot C(h)$ for any pair of smooth functions $f, h : M \to \mathbb{R}$ (See for example Corollary 2.1.14 of [SM93] for a proof). Here $f \cdot h$ represents a function whose value at every point x equals to the product $f(x)h(x)$. Intuitively, this is because a functional map that preserves products of functions must satisfy $C(f^2) = C(f)^2$. If f is an indicator function of a region then $f^2 = f$ and this latter condition implies that $C(f) = C(f^2) = C(f)^2$ which means that $C(f)$ must itself be an indicator function of a region. Thus, the preservation of products of functions is directly related to guiding general functional maps to correspond to point-to-point maps in both the continuous and the discrete setting. Here non-trivial means that $C(\mathbf{1}_M) = \mathbf{1}_N$, where $\mathbf{1}_M$ is the constant function equal to one everywhere on M. This constraints is indeed trivial, because $\mathbf{1} \circ T = \mathbf{1}$ for any T. We also note that without this trivial constraint, the preservation of products of functions is still a very strong condition on a functional map and guarantees a partial correspondence coming from a generalized composition operator (See Example 2.1.10 on p. 21 of [SM93] for a discussion).

Perhaps the simplest way to introduce this result and intuition into the pipeline described above is by taking multiple pairs of descriptor functions $f^{(p_1)}, g^{(p_1)}$ and $f^{(p_2)}, g^{(p_2)}$ for which we expect $Cf^{(p_1)} \approx g^{(p_1)}$, $Cf^{(p_2)} \approx g^{(p_2)}$, and producing new function preservation constraints $f^{(p_3)}, g^{(p_3)}$ via $f^{(p_3)} = f^{(p_1)} \cdot f^{(p_2)}$ and $g^{(p_3)} = g^{(p_1)} \cdot g^{(p_2)}$. There are however, several issues with such an approach: first, it is not clear how many additional constraints are necessary and what pairs of descriptor functions should be taken. Secondly, any noise in the descriptors will be amplified when pairs of such functions are taken. Thus, we take a slightly different approach as described below.

To motivate our construction further, consider a pair of descriptors $f^{(p)}, g^{(p)}$ that are "fully discriminative," in the sense that for every point $x \in M$ there exists a unique point y on N such that $g^{(p)}(x) = f^{(p)}(y)$. Given such a pair of descriptors, we would expect to recover the underlying point-to-point map using a single function (descriptor) preservation constraint. However, if we simply enforce $Cf^{(p)} = g^{(p)}$, then even in the full basis we will not be able to recover the underlying map, since the function preservation constraint only leads to k_N linear equations instead of the required $k_M k_N$ equations. This is because the simple function preservation constraint does not preserve the individual level-sets of the function values, which should be expected from a map. A simple method might be to decompose a single pair of descriptor functions $f^{(p)}, g^{(p)}$, into multiple function preservation constraints by introducing new functions by considering level-sets (or Gaussians around certain values, as in [OMPG13], Section 4.1), but this again can result in more noise and additional parameters.

5.2. Our constraints

In this paper, we propose a different approach to function preservation constraints. Namely, we start with the observation that in the full basis, given a pair of corresponding probe functions $f^{(p)}, g^{(p)}$, we would expect the mapping matrix Π to be such that $\Pi_{i,j} \cdot (f_j^{(p)} - g_i^{(p)}) = 0$ for all i, j, which is equivalent to saying that indicator functions of regions of constant values of $f^{(p)}$ and $g^{(p)}$ are preserved. This corresponds to the intuition that the level-sets of functions must be preserved along with the values of the functions themselves. This constraint can be rewritten via commutativity as $\Pi Diag(f^{(p)}) = Diag(g^{(p)})\Pi$, where $Diag(v)$ is the matrix that contains the values of the vector v along the diagonal and is zero elsewhere. This form also makes apparent the relation between preservation of level sets and function products. Indeed, if $h : M \to \mathbb{R}$ is any function on M, then its pointwise product with $f^{(p)}$ is obtained, in the discrete formulation, via the matrix vector product $Diag(f^{(p)})h$. Therefore, $\Pi Diag(f^{(p)}) = Diag(g^{(p)})\Pi$ implies $\Pi Diag(f^{(p)})h = Diag(g^{(p)})\Pi h$, which implies that the associated linear mapping C between functions must satisfy $C(f^{(p)} \cdot h) = g^{(p)} \cdot C(h)$ for any $h : M \to \mathbb{R}$, which corresponds exactly to the product rule.

Our constraints in the reduced basis As discussed above, in the functional map framework, the key map estimation step is done in the reduced basis. Thus, we introduce our constraints by following the idea of commutativity with an operator based on the descriptor, as discussed in the previous paragraph. However, in the reduced basis the commuting matrices will not remain diagonal. For a given pair of descriptor functions $f^{(p)}, g^{(p)}$, we therefore create matrices $X^{(p)} = \Phi_M^+ Diag(f^{(p)})\Phi_M$ and $Y^{(p)} = \Phi_N^+ Diag(g^{(p)})\Phi_N$. Finally, we add the corresponding constraints to the system into Eq. 1 by requiring the unknown map C to satisfy in the least squares sense:

$$\sum_s \|CX^{(p)} - Y^{(p)}C\|^2, \tag{2}$$

where the summation is across the available descriptors, and we use the Frobenius matrix norm.

5.3. Properties

As mentioned above, the descriptor preservation constraints alone do not extract all of the information that is present in a given descriptor. In particular even if the descriptors are perfect and identify each vertex uniquely, the classical constraints $CF = G$ may still not be sufficient to identify each vertex. As a toy example, if $n_M = 2$, $n_N = 2$ and $f = (1,2)$ and $g = (1,2)$, then there is a unique point-to-point map that preserves these functions. However, using only the constraint $\Pi \begin{pmatrix} 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$ could also lead to the solution $\Pi_1 = \begin{pmatrix} 0 & 0.5 \\ 2 & 0 \end{pmatrix}$. Thus, we can see that although Π_1 preserves the descriptor functions, it fails to preserve the values pointwise: it fails to map vertices that have a value to vertices that have a similar value. Π_1 does not preserve the commutativity constraint, and indeed Π_1 is not a permutation matrix.

Below we show that the above phenomenon cannot happen if Π is enforced to be a doubly stochastic matrix: i.e., having entries that all lie in the interval $[0,1]$ and whose rows and columns sum to 1:

Theorem 1 Let $f \in \mathbb{R}^n$ and $g \in \mathbb{R}^n$ be such that the multiset of values contained in f and in g are the same. Let Π be an $n \times n$ matrix such that $\forall i,j, 0 \leq \Pi_{i,j} \leq 1$ and $\sum_k \Pi_{i,k} = 1$, $\sum_k \Pi_{k,j} = 1$. Then $\Pi f = g$ implies $\Pi_{i,j} = 0$ whenever $f_j \neq g_i$.

Proof: We proceed by induction on the values of f. Let $L = \max(f) = \max(g)$. By assumption, the sets $I_f = \{k | f_k = L\} \subset \{1,\ldots,n\}$ and $I_g = \{k | g_k = L\} \subset \{1,\ldots,n\}$ must have the same cardinality. Moreover, each g_k for $k \in I_g$ can only be obtained from combinations of f_k for $k \in I_f$. Thus, $\Pi_{k,k'} = 0$ if $k \in I_g$ and $k' \notin I_f$. These constraints also imply $\Pi_{k,k'} = 0$ if $k \notin I_g$ and $k' \in I_f$, because each column of Π should sum to 1. $\square$

Unfortunately, enforcing a matrix to be a stochastic matrix involves inequality constraints that do not translate well in the reduced basis: e.g., the projection of a stochastic matrix in the reduced basis may not remain stochastic. Thus, it is not easy to restrict to such matrices in the reduced basis. Instead, rather than enforcing inequality constraints we propose to introduce the commutativity with respect to the operators derived from the descriptor functions as described above. In our toy example, the commutativity constraint would be

$$\Pi \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} \Pi$$

We can see that enforcing such a constraint eliminates the wrong solution Π_1.

Recall that any functional map that corresponds to a point-to-point map should satisfy $\Pi \mathbf{1} = \mathbf{1}$, which would thus be a natural constraints to use in our optimization. Interestingly our new commutativity constraint along with the additional regularization, requiring the map to preserve the constant function $\Pi \mathbf{1} = \mathbf{1}$, implies the previously used constraints $\Pi f = g$ even in the reduced basis, as proved in the following theorem:

Theorem 2 If $f \in \mathbb{R}^{n_M}$, $g \in \mathbb{R}^{n_N}$ and $\Pi \in M_{n_N,n_M}(\mathbb{R})$, then $\Pi Diag(f) = Diag(g)\Pi$ and $\Pi \mathbf{1} = \mathbf{1}$ implies that $\Pi f = g$. Similarly, if C is in the reduced basis, where the first basis function is a constant function, and $Ce_1 = e_1$, then the commutativity constraint $C\Phi_M^+ Diag(f)\Phi_M = \Phi_N^+ Diag(g)\Phi_N C$ implies that $C\Phi_M^+ f = \Phi_N^+ g$.

Proof: We consider the first case, in the full basis: $\Pi f = \Pi Diag(f)\mathbf{1} = Diag(g)\Pi \mathbf{1} = Diag(g)\mathbf{1}$. For the second case, by assumption, we have: $C\Phi_M^+ f = C\Phi_M^+ Diag(f)\Phi_M e_1 = \Phi_N^+ Diag(g)\Phi_N Ce_1 = \Phi_N^+ Diag(g)\Phi_N e_1 = \Phi_N^+ g$. $\square$

However, in practice, it might still be useful to enforce both the commutativity constraint and the $CF = G$ constraint as the latter might give more control on the importance of preserving the function globally, by e.g., adding a scalar weight. Note that this theorem is only meant to be a theoretical guarantee that our formulation includes at least as much information as the previous one.

6. Experiments

As described above, the new commutativity constraints allow us to extract more information from the same set of descriptors, and furthermore allow us to guide the functional map estimation process closer to point-to-point maps, while still maintaining the linear (least squares) complexity of the optimization. Below we demonstrate the utility of these constraints on a wide range of shapes and deformations and show that our approach allows to significantly reduce the number descriptor functions necessary to estimate an accurate functional map, and even improve results with the increase in the size of the basis for a fixed number of descriptors, which is not true for the previously used constraints.

6.1. Using few descriptors

In our first set of experiments, we plot the average correspondence error for several methods on three standard benchmarks: FAUST [BRLB14], SCAPE [ASK*05], and TOSCA [BBK08].

Our main goal in this experiment is to show that by formulating the descriptor preservation constraints via commutativity, rather

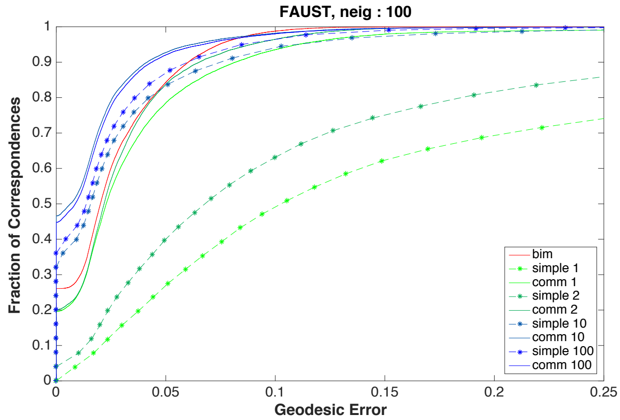

Figure 1: *Error plots showing the accuracy of our descriptor preservation via commutativity (solid lines) compared to simple value preservation (dotted) and the Blended Intrinsic maps (red) on shape pairs from the FAUST dataset. Our method allows to obtain superior performance using even a very small descriptor set.*

Figure 2: *Example maps obtained by formulating the descriptor preservation with the simple method (left) and using our commutativity approach (right), using exactly the same descriptor functions.*

than using the original approach based on preservation of values, results in more accurate functional map inference, without requiring any additional information.

We compare our results to the following methods:

- Blended Intrinsic Maps [KLF11]
- The original method proposed in [OBCS*12] and used in follow-up works [PBB*13,RCB*16], that formulates function preservation constraints, based on values.

We note that our approach can be incorporated into *any* pipeline for estimating functional maps, which uses function (e.g., descriptor) preservation constraints. As such, it can be easily combined with the other techniques that have been introduced for optimizing functional map computations, e.g., based on sparsity [PBB*13] or specific prior structure existing in partial correspondences [RCB*16]. Therefore, our goal is not to demonstrate that our particular choice of descriptors or parameters results in state-of-the-art correspondences on these benchmarks, but rather to show that our formulation of function preservation via commutativity allows to obtain more accurate results than the one based on function values alone.

For this, we used the original functional map estimation pipeline introduced in [OBCS*12], with the same code and parameters. Namely, we used a varying number $(1,2,10,$ and $100)$ of descriptor functions, and compared the results of estimating the functional map either by minimizing:

$$C_{\text{opt}} = \arg\min_{C} \|CF - G\|^2 + \alpha \|\Delta_2 C - C\Delta_1\|^2,$$

as described in Section 4 or using our constraints, which simply adds an extra term to the energy above, given by $\sum_i \|CX_i - Y_i C\|^2$, as described in Eq. 2 above. In both cases we computed the functional map C by solving a linear least squares system, using a vectorization of the functional map C (re-writing it as a vector c), and solving the system $Ac = b$, where A and b are obtained by rewriting the above energy in matrix-vector form. After estimating the functional map C we used the post-processing technique of [OBCS*12] based on high-dimensional ICP to both refine the functional map and convert it to a point-to-point correspondence.

In all of the experiments in this section we used $n_{\text{eig}} = 100$ eigenfunctions to represent the functional basis, by discretizing the Laplace-Beltrami operator using the standard cotangent-scheme [PP93, MDSB03], and a sparse eigensolver, to estimate the basis. We followed the exact pipeline suggested in [OBCS*12] for map estimation and simply sub-sampled the set of descriptors used in that work (Wave Kernel Signature [ASC11] and Wave Kernel map based on segment correspondences).

Figures 1, 3, 5 show the results obtained using our approach compared to the basic function-preservation method on the three benchmarks, using a varying number of descriptor preservation constraints. We evaluated each method on 100 pairs of shapes in the FAUST dataset, 76 shape pairs in TOSCA and 71 pairs in SCAPE, by taking each shape to be a source in exactly one pair. We follow the evaluation protocol introduced in [KLF11], by plotting on the x-axis a geodesic threshold, and on the y-axis, the fraction of the correspondences obtained by each method that are at a distance that is less than this threshold from the ground truth map. The geodesic distances (approximated using Dijkstra's algorithm) are divided by the scaling factor $\sqrt{Area}$, where *Area* is the total area of the shape. In this work, similarly to other non-symmetry aware intrinsic methods, we do not disambiguate left-right symmetries, and thus take the minimum between the distance between the matched vertex and the ground truth, and the distance between the matched vertex and the symmetric of the ground truth as our distance to target.

In each plot, the red curves represent the performance of [KLF11], the blue/green curves represent approaches based on the functional map framework (green: using only 1 descriptor, blue: using 100 descriptors). The curves with $*$ symbols use the original function-preservation formulation, whereas those with solid lines use our new approach.

We notice the following general trend: the method that uses our commutativity constraints usually leads to better results compared to the classic descriptor preservation constraints, and the improvement is particularly large when only a few descriptors are used. Using more descriptors leads to little improvement for this new method, which, together with the previous fact, highlights that the new formulation helps in extracting more information from the same descriptors.

We also notice that increasing the number of descriptors used to 100 is not generally the best choice for getting better results, which

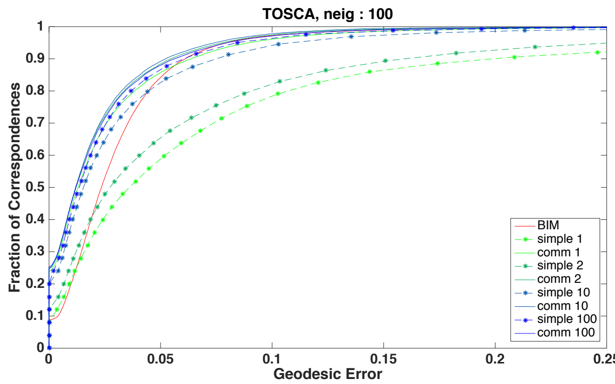

Figure 3: *Error plots showing the accuracy of descriptor preservation via our commutativity approach (solid lines) compared to simple value preservation (dotted) and the Blended Intrinsic maps (red) on the TOSCA dataset.*

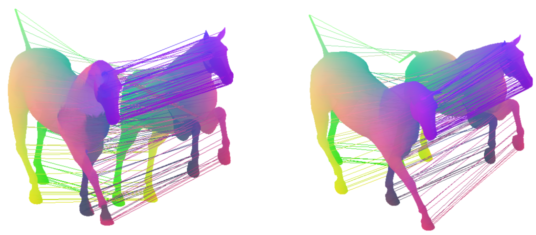

Figure 4: *Example maps obtained by formulating the descriptor preservation with the simple method (left) and using our commutativity approach (right), using exactly the same descriptor functions.*

shows that this method is well-suited for performing on few reliable descriptors.

Perhaps most remarkably, our new formulation allows to obtain results with *only two* descriptor functions (e.g., on the SCAPE dataset) that are better than the ones produced by the original method using the full set of 100 descriptors.

In Figures 2, 4 and 6 we also provide some example maps computed using descriptor preservation with commutativity vs. the simple value-based approach. Note that the resulting maps are typically less noisy and more globally consistent, despite using exactly the same information in the optimization, which also suggests that our formulation helps to obtain more accurate functional map. For visualization, we sampled $100, 300$ and 200 points on the source shape uniformly from the list of all vertices, for the pairs from the FAUST (Fig. 2), TOSCA (Fig. 4), and SCAPE datasets (Fig. 6) respectively.

6.2. Changing the dimension of the reduced space

In our second range of experiments, we show the dependence of the results on the number of functions used in the basis for functional maps. Here, rather than changing the number of descriptor functions, we fix the descriptor set and change the dimensionality of the basis and evaluate the quality of the approximation of the point-to-point correspondences using the functional map pipeline.

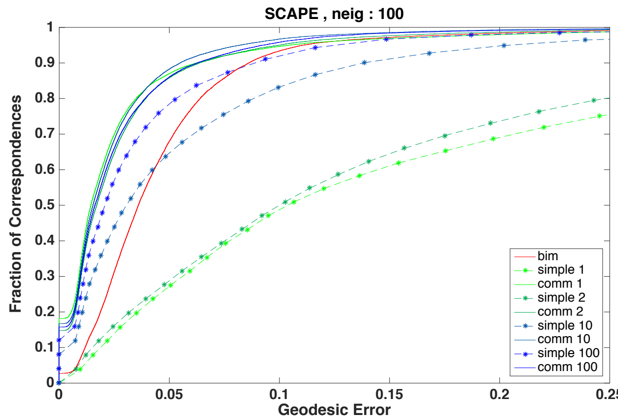

Figure 5: *Error plots showing the accuracy of our descriptor preservation via commutativity (solid lines) compared to simple value preservation (dotted) and the Blended Intrinsic maps (red) on the SCAPE dataset. Our method allows to obtain superior performance using even a very small descriptor set.*

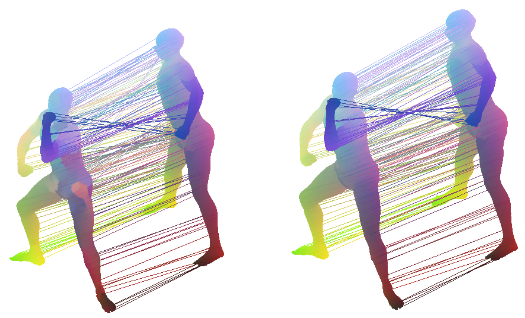

Figure 6: *Example maps obtained by formulating the descriptor preservation with the simple method (left) and using our commutativity approach (right), using exactly the same descriptor functions.*

In Figure 7 we show the average correspondence error between a subset of shapes in the FAUST dataset [BRLB14], using the same pipeline described in the previous section, for a fixed number (in this case two) of descriptor functions. In particular, we used the Wave Kernel Signature for a single energy value, along with a single descriptor function that is aimed at segment preservation using the Wave Kernel Map with a fixed energy value. This gives us two descriptor functions, which we incorporate into the functional map energy using either the standard descriptor preservation constraints, as done in [OBCS*12] or using our commutativity-based approach. We then convert the estimated functional map to a point-to-point correspondence and evaluate its accuracy using the distance to the ground truth. We plot the average pointwise map error, computed the same way as in the previous experiment, across the shape pairs for a varying number of basis functions.

As can be seen in Figure 7, compared to the basic method for descriptor preservation, our approach allows not only to improve quality of the correspondences significantly, without using any additional information, but also provides more resilience with respect

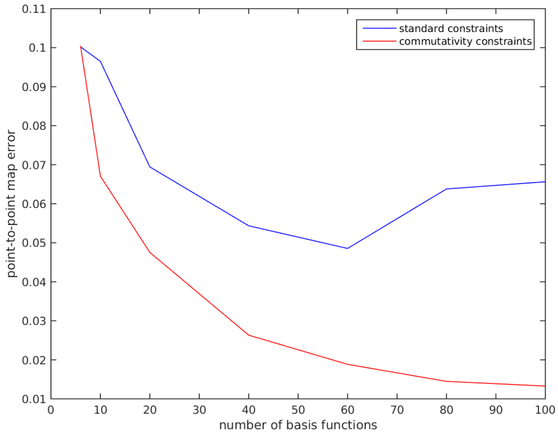

Figure 7: *Average error on pairs of shapes in the FAUST dataset, depending on the number of basis functions in the functional map representation, for a fixed number (two in this case) descriptors. Unlike the standard approach of [OBCS*12], which deteriorates when the size of the basis significantly exceeds the number of descriptors, our method continues to produce high-quality results even for a small number of descriptors and a large number of basis functions. The average error is computed as the average geodesic distance to the ground truth correspondence, symmetries allowed.*

to the choice of the number of basis functions, for a fixed descriptor set. This implies that our approach can potentially enable more accurate correspondence computation based on the functional map pipeline, without requiring any additional information, and supports the idea that using our formulation allows to extract additional information from descriptor functions, which in turns results in better pointwise maps.

7. Conclusion, Limitations & Future work

We proposed a new formulation for incorporating descriptor (or more general function) preservation constraints within the functional map framework, which enables finding better solutions to the non-rigid shape matching problem. Our formulation is especially useful when the number of descriptors is lower than the dimension of the reduced space of functions since it allows to extract more information from the same set of given descriptors. Our formulation is applicable in the same settings as the original functional maps framework, with the main limiting factor being the computational time necessary to assemble and solve our optimization problem, which nevertheless remains linear in the unknown map. We also note that we do not enforce the preservation of the constant function in practice, and thus our formulation should, in principle be also applicable even to partial maps. We leave the exploration of this as an interesting direction for future work.

Conceptually, we propose to consider descriptors or functions as *linear functional operators* acting on other functions through multiplication, unlike the standard approach which views them simply as scalar-valued signals. We believe that this idea is particularly exciting for future work, and are planning to investigate other ways in which informative descriptors and constraints can be defined directly as functional operators, opening the door to a much richer way of characterizing shapes and their geometry, which can be useful in shape matching problems.

8. Acknowledgments

This work was supported in part by Marie-Curie CIG-334283, a CNRS chaire d'excellence, chaire Jean Marjoulet from Ecole Polytechnique, FUI project TANDEM 2 and a Google Focused Research Award. We would also like to thank Marc Glisse, Etienne Corman and Frederic Chazal for useful discussions.

References

[APL15] AIGERMAN N., PORANNE R., LIPMAN Y.: Seamless Surface Mappings. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 72. 2

[ASC11] AUBRY M., SCHLICKEWEI 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. 4, 6

[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. 5

[BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized Multidimensional Scaling: A Framework for Isometry-Invariant Partial Surface Matching. *Proceedings of the National Academy of Sciences 103*, 5 (2006), 1168–1172. 1, 2

[BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer Science & Business Media, 2008. 5

[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. 5, 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. 2

[COC14] CORMAN É., OVSJANIKOV M., CHAMBOLLE A.: Supervised Descriptor Learning for Non-Rigid Shape Matching. In *European Conference on Computer Vision* (2014), Springer, pp. 283–298. 1, 2, 3

[GSTOG16] GANAPATHI-SUBRAMANIAN V., THIBERT B., OVSJANIKOV M., GUIBAS L.: Stable Region Correspondences Between Non-Isometric Shapes. In *Computer Graphics Forum* (2016), vol. 35, pp. 121–133. 2

[HSS*09] HASLER N., STOLL C., SUNKEL M., ROSENHAHN B., SEIDEL H.-P.: A Statistical Model of Human Pose and Body Shape. In *Computer Graphics Forum* (2009), vol. 28, pp. 337–346. 1

[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. 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, 6

[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. 1, 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. 3, 6

[Mém11] MÉMOLI F.: Gromov–Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of computational mathematics 11*, 4 (2011), 417–487. 2

[OBCS*12] OVSJANIKOV 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, 3, 6, 7, 8

[OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One Point Isometric Matching With the Heat Kernel. In *Computer Graphics Forum* (2010), vol. 29, pp. 1555–1564. 2

[OMPG13] OVSJANIKOV M., MÉRIGOT Q., PĂTRĂUCEAN V., GUIBAS L.: Shape matching via quotient spaces. In *Computer Graphics Forum* (2013), vol. 32, pp. 1–11. 4

[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. 1, 2, 3, 6

[PP93] PINKALL U., POLTHIER K.: Computing Discrete Minimal Surfaces and their Conjugates. *Experimental mathematics 2*, 1 (1993), 15–36. 3, 6

[RCB*16] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial Functional Correspondence. In *Computer Graphics Forum* (2016). 2, 3, 6

[RMC15] RODOLÀ E., MÖLLER M., CREMERS D.: Point-wise map recovery and refinement from functional correspondence. In *Vision, Modeling & Visualization, VMV* (2015), pp. 25–32. 2

[SM93] SINGH R. K., MANHAS J. S.: *Composition Operators on Function Spaces*, vol. 179. Elsevier, 1993. 4

[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

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. In *Computer graphics forum* (2009), vol. 28, pp. 1383–1392. 4

[SP04] SUMNER R. W., POPOVIĆ J.: Deformation Transfer for Triangle Meshes. *ACM Transactions on Graphics (TOG) 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, 3

[TBW*09] TEVS A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: Isometric Registration of Ambiguous and Partial Data. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on* (2009), IEEE, pp. 1185–1192. 2

[VKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A Survey on Shape Correspondence. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 1681–1707. 1, 2