

Supervised Descriptor Learning for Non-Rigid Shape Matching

Étienne Corman^{1,2}, Maks Ovsjanikov¹, and Antonin Chambolle²

¹ LIX, École Polytechnique

² CMAP, École Polytechnique, CNRS

Abstract. We present a novel method for computing correspondences between pairs of non-rigid shapes. Unlike the majority of existing techniques that assume a deformation model, such as intrinsic isometries, *a priori* and use a pre-defined set of point or part descriptors, we consider the problem of *learning* a correspondence model given a collection of reference pairs with known mappings between them. Our formulation is purely intrinsic and does not rely on a consistent parametrization or spatial positions of vertices on the shapes. Instead, we consider the problem of finding the optimal set of descriptors that can be jointly used to reproduce the given reference maps. We show how this problem can be formalized and solved for efficiently by using the recently proposed functional maps framework. Moreover, we demonstrate how to extract the functional subspaces that can be mapped reliably across shapes. This gives us a way to not only obtain better functional correspondences, but also to associate a confidence value to the different parts of the mappings. We demonstrate the efficiency and usefulness of the proposed approach on a variety of challenging shape matching tasks.

Keywords: Shape matching, Correspondences, Feature learning

1 Introduction

Finding high quality correspondences is a key component in many tasks including statistical shape analysis [12], deformation transfer [30] and interpolation (morphing) [16] among others. While a number of efficient techniques have been proposed to address the problem of rigid alignment [32], the problem of general non-rigid shape matching remains difficult.

Most existing methods for finding correspondences between non-rigid shapes rely on an *a priori* deformation model, which specifies the space of “reasonable” maps between shapes. Perhaps the most popular and widely used such model is that of approximate intrinsic isometries [5, 21], where the mapping is assumed to preserve geodesic distances between all pairs of points on the shapes. A more general possibility is to consider conformal deformations, which are only assumed to preserve angles [19, 17] or to parameterize the space of possible maps using a fixed deformation model [36]. Although these techniques can produce good results when the deformation satisfies the *a priori* model, they can fail badly

as soon as even moderate deviations from the model are introduced. This is especially critical since many natural deformations, such as articulated motion of humans or animals are known to induce potentially significant geodesic distortion [27]. Incorporating the possibility for such distortion into a deformation model is challenging especially using a purely axiomatic (theoretical) approach.

Rather than trying to devise a theoretical deformation model capable of adapting to known deformations, several communities have tackled this challenge by using a data-driven approach, where the space of “reasonable” maps or deformations is learned from a set of examples, e.g. [8]. Since obtaining example deformations is often significantly easier than devising a unified theoretical deformation model, such an approach allows the resulting techniques to remain flexible yet efficient in the particular settings where they are applied.

Most data-driven approaches for devising a deformation model, however, rely heavily on a consistent parametrization of the deformation domain (e.g. on a fixed grid in Euclidean space), and perform statistical analysis on the positions of vertices of the shapes [9, 4, 10, 11]. When computing correspondences between pairs of surfaces in 3D, such parametrization is often unavailable and moreover, shapes can undergo severe deformations which are difficult to capture using purely extrinsic approaches.

In this paper, we propose a purely intrinsic method for exploiting prior correspondence information between pairs of shapes to find better correspondences between a reference shape and a new previously unseen instance. Rather than doing the learning over, e.g., the positions of the vertices on the shapes, we propose to find the optimal set of *point descriptors* that can be jointly used to reproduce the given reference maps. While such an optimization is, in general, very complicated, since even to evaluate how well the descriptors can reproduce a given map would require a full solution of the shape matching problem, we show how this problem can be formalized and solved for efficiently by using the recently proposed functional maps framework [22]. Moreover, we demonstrate how to extract the functional subspaces that can be mapped reliably across shapes. This gives us a way to not only obtain better functional correspondences, but also to associate a confidence value to the different parts of the mappings. Our approach is also quite general since it can be used as a preprocessing step of other methods using functional maps [24, 14, 3] in order to improve the quality of the results and help to handle difficult deformation. Note that in this paper we focus on the shape matching problem which is the most developed application of the functional maps.

1.1 Related Work

Non-rigid shape matching is a very-well developed area and its complete overview is beyond the scope of this paper (see, e.g., [6, 35] for recent surveys of this field). We therefore concentrate on the work directly related to ours, namely near-isometric shape matching with special emphasis on approaches that utilise prior knowledge for establishing correspondences between pairs of shapes.

The vast majority of techniques for non-rigid shape matching implicitly make use of a deformation model for finding correspondences between geometric shapes. Perhaps the most common model in the context of intrinsic (i.e., not relying on vertex positions and not assuming approximate alignment) approaches is approximate isometries, introduced by Bronstein et al. [5] and Mémoli [21]. This model has been used by a large number of methods, (e.g., [13, 33, 23, 28, 22] among many others) that all assume that the sought correspondences must approximately preserve pairwise geodesic distances. Another set of approaches is based on a more relaxed model, conformal mappings, used by, e.g., [19, 17] where only angles are assumed to be preserved. Other techniques, such as the one used by Zhang et al. [36] explicitly deform a shape using a fixed deformation model to find correspondences between non-rigid shapes.

All of these approaches use a model given *a priori* to find correspondences, which can be problematic if the real deformations do not agree with the given model. Interestingly, it has recently been observed [27] that even articulated motion of humans can induce noticeable isometric distortion, which could explain some of the difficulties encountered by previously proposed techniques.

In contrast, other works have proposed to *learn* an appropriate deformation model from a set of examples, and then use this model for shape matching. Perhaps the best-known example of this approach are Active Shape and Active Appearance Models [8, 9] and their variants (see, e.g., [11]) used widely in Computer Vision. In a similar vein, techniques in Statistical Shape Analysis [10] use the distribution of positions of pre-specified landmark points in 3D to learn a statistical deformation model over which inference can be made. Related techniques are commonly used in medical imaging and Morphometrics [4] and in Geometry Processing communities, e.g. [1, 12] among many others. However, all of these methods assume the existence of a common domain over which learning can be made, and which most often is done using vertex coordinates of either landmark points or all points on a fixed reference shape. In the context of intrinsic shape matching, where shapes lack labeled landmark points and can undergo severe deformations, vertex coordinates are often not relevant, limiting the applicability of such techniques.

Rather than relying on vertex positions, recent methods have considered using derived properties such as point or triangle *descriptors* for learning. Thus, Kalogerakis et al. [15] and Van Kaick et al. [34] have proposed using a set example shape to train classifiers for part segmentation and labeling, which can then be used to establish part-level correspondences. Similarly, Chen et al. [7] explore the predictive power of various descriptors for detecting distinctive landmark (schelling) points identified by users. These methods, while similar to ours in learning on the level of descriptors do not, however, specifically address the shape matching problem.

Perhaps most closely related to ours are recent works by Litman et al. [20] and Rodolà et al. [26], where the authors use a set of examples to learn the most informative descriptors that are used directly in the context of shape matching. Our approach is fundamentally different, however, since rather than trying

to identify descriptors that can distinguish different points, we propose to find the optimal *descriptor set* that can be used to *jointly produce the entire map across shapes*. We thus avoid the problem of obtaining *consistent* correspondences present in these approaches (and obtained during post-processing), since consistency is incorporated directly in the learning stage. Crucially, we use the recently proposed functional map representation [22] that allows us to formulate the learning problem purely intrinsically, while permitting to directly control and optimize for the influence of descriptors on the quality of the final map.

Goals Given a collection of (training) shapes with known correspondences our goal is to identify the most informative descriptor set that can be used to solve the non-rigid shape matching problem on new (test) instances. Besides we want to learn where are the most stable correspondences.

2 Consistent Functional Maps

Our method is based on the functional map representation introduced in [22]. In this section, we give a brief overview of the representation and the method used in [22] to construct a functional map for a given pair of shapes.

While our method is general, throughout the paper we assume that all shapes are represented as triangle meshes, and all functions are expressed as vectors in the basis of the eigenfunctions of the Laplace-Beltrami operator. This basis needs to be computed beforehand on each shape. The objective is to output a uniquely defined functional map.

2.1 Functional Map Representation

The functional map representation is based on the observation that given two surfaces S_0 and S_i , a point-to-point map $T_i : S_i \rightarrow S_0$ induces a map between function spaces $C_i : L^2(S_0) \rightarrow L^2(S_i)$, where $L^2(S)$ is the set of square integrable functions defined on the surface S . The functional map C_i is defined by composition with T_i as $C_i f = f \circ T_i$. The operator C_i is a linear transformation and given a basis it can be represented as a matrix in the discrete setting. This matrix can be easily computed if the map T is known.

The basic method described in [22] approximates the functional map C_i using a set of linear constraints. The first type of constraints is given by a set of pairs of functions, which we refer to below as “probe functions”, that are expected to be preserved by the deformation. The second is a regularization term coming from the deformation model. This leads to the least square problem:

$$X_i = \arg \min_C \|CG_0 - G_i\|_F^2 + \alpha \|C \odot W\|_F^2, \quad (1)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. The use and meaning of each term will be detailed in the following paragraphs.

Probe Functions The probe functions can be represented by two matrices G_0 and G_i , where each pair of corresponding columns represents a pair of functions g_0, g_i such that $C_i g_0 \approx g_i$ is expected to hold for the unknown C_i . In practice we normalize the corresponding column so that each column has the same L^2 norm. Thus, the functional map should verify $CG_0 \approx G_i$. In the context of isometric matching the probe functions are given by classical descriptors, such as the HKS [31], or WKS [2].

Regularization The existence of intrinsic symmetries can lead to an ambiguity between two equally good solutions. In general, the probe functions do not resolve the symmetry ambiguity, that is G_0 and G_i are composed only of symmetric functions. So, in addition to the probe function constraints, the authors of [22] have proposed a regularization using the assumption that the deformation is nearly isometric. This assumption is equivalent to the commutativity of C_i with the Laplace-Beltrami operator, namely $C_i \Delta_0 = \Delta_i C_i$. In the discrete setting the eigenfunctions of the Laplace-Beltrami operator are used as function basis. Thus, this equation can be written as $C_i \odot W = 0$ where “ \odot ” denotes the component-wise multiplication and the matrix W is defined by $W_{kl} = \lambda_k^i - \lambda_l^0$ with λ_k^i the k^{th} eigenvalue of the Laplace-Beltrami operator on the surface S_i .

Uniqueness of the solution. In practice the eigenvalues of the Laplace-Beltrami operator of two different shapes are always numerically different except for the zero eigenvalues. Thus, the only zero coefficient of W is $W_{1,1}$ which weights the coefficient $C_{1,1}$ of the functional map. Since the corresponding eigenfunctions are constant, $C_{1,1}$ maps the constant functions of $L^2(S_0)$ into the constant functions of $L^2(S_i)$. The coefficient $C_{1,1}$ should always be one. Therefore, since W is non zero everywhere, the solution of (1) is unique without any assumptions on G_0 .

2.2 Main Challenge

In the original article [22] the probe functions are assumed to be given, so how to choose them was not discussed. As mentioned in introduction, this choice can already be challenging. For example in Figure 1a the smoothed Gaussian curvature computed on two different meshes provides a decent functional correspondence. At the same time, in Figure 1b the logarithm of the Gaussian curvature, while intrinsic in theory, does not result in a useful correspondence.

One option to identify the best descriptors would be to simply find the most stable probe functions in the example (training) set, by learning spectral descriptors [20] for example. However, some descriptors (e.g., the constant function) can be stable without at all being informative. More importantly, however, as can be seen from Equation (1), the descriptors influence the resulting functional maps X_i *jointly*. As an example, if a correspondence is described by several probe functions the resulting functional map will tend to respect this constraint while other meaningful correspondences will be arbitrarily put aside due to their low redundancy. So picking the best descriptors independently will not necessarily result in high quality maps.




Fig. 1: Probe functions computed independently on two shapes. One carries meaningful information (a) and the other is misleading (b).

Thus, the key idea developed in this paper is to introduce weights for probe functions, over which learning can be made. As explained below, the probe function constraint will be replaced by: $\|CG_0D - G_iD\|_F^2$. The weights D will be optimized so that the weighted descriptors are jointly as informative as possible. This will allow us to improve the quality of the functional maps and to extract the most stable functional subspaces.

2.3 Algorithm Outline

We propose a two-step method described in following two sections and summarized in Figure 2. Given collection of shapes, we learn the most informative set of weight D by solving an optimization problem. We then extract a function basis whose components are ordered by quality of correspondence. When given a previously unseen shape, we use this information to compute a high-quality functional map using the optimal weights and to discard the badly mapped functions by reducing the functional space.

3 Selection of the Best Functional Correspondences

The idea developed here is to assign a weight to each pair of probe functions. These weights can then be tuned according to their consistency in the matching. Since *a priori* there is no reason to choose one probe function over another, we propose to learn the optimal weights given a training set of shapes.

As input we need a set of N triangulated meshes with known correspondences representing the same object undergoing a set of deformations. Our main assumption is that the optimal weights on the probe functions should be stable across the shapes in the collection. Thus, if we are given a new deformation of the same shape, the learned weight should also select the consistent probe functions. The output of our algorithm will be a set of weights for the probe functions, which, as we will show below, can then be used to find correspondences between new, unseen shape instances.

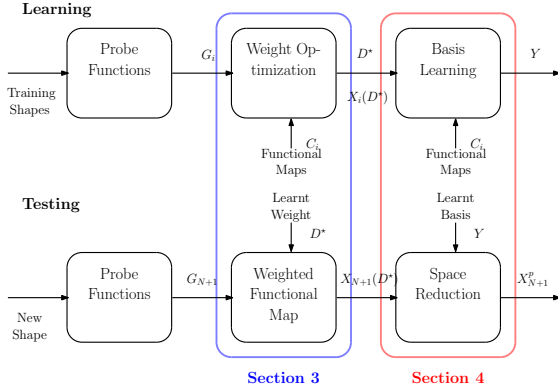


Fig. 2: Pipeline of the proposed algorithm with the corresponding section. Top: learning from a given collection. Bottom: processing a new shape.

3.1 Weighting the probe functions

As described above, our main idea is to introduce a set of weights on the functional correspondences to measure their usefulness in finding a relevant map by using a diagonal matrix D . For a given weight matrix D , the linear constraints given by the probe functions become $CG_0D = G_iD$. We can then define the function $X_i(D)$, which maps a given sets of weight to the corresponding functional map, via the solution of the optimization problem:

$$X_i(D) = \arg \min_C \|CG_0D - G_iD\|_F^2 + \alpha \|C \odot W\|_F^2 \quad (2)$$

We choose here to fix α and tune D . We could also try to tune all the parameters (α and D) but the coefficients would be defined up to a multiplicative constant and $C(D)$ may no longer be well-defined when α is equal to zero.

Since all the functional map start from the reference shape S_0 , this shape obviously plays special role in our method. Ideally we would like to take as reference the most “average” shape of the collection. Following this idea, a simple procedure is presented in [29] to find the shape of the collection which minimizes the average isometric distortion. However, in our experiments we chose the symmetric standing pose as reference.

As discussed in the previous section $X_i(D)$ is well-defined and differentiable.

3.2 Finding the best weights

Learning from a given collection. We assume we are given a collection of N nearly isometric deformations of the same object with known functional maps C_i . The optimal weights D^* are the ones that produce an approximation $X_i(D)$ that is closest to the ground truth C_i . Thus, we want to solve the following optimization problem:

$$D^* \in \arg \min_D \sum_{i=1}^N \|X_i(D) - C_i\|, \quad (3)$$

where the sum is over the set of given training maps C_i .

Note that the choice of the norm is important. We would like the functional map $X_i(D)$ to match C_i over as-large-as possible functional subspace. This is equivalent to minimizing the rank of the difference $X_i(D) - C_i$. Thus, the important quantities are the singular values of the differences $X_i(D) - C_i$.

The naive choice of the squared Frobenius norm is not well suited for our problem since $\|A\|_F^2 = \|\sigma(A)\|_2^2$ where $\sigma(A)$ is the vector containing the singular values of the matrix A . Therefore this norm would give a large weight to the biggest singular values, which correspond to the worst-matched functional subspaces. Among these subspaces is the space of antisymmetric functions that we have no hope of mapping since the probe functions give us very little information about this subspace. At the same time, the small singular values have little influence on the minimization whereas they are the ones we would like to optimize.

The choice of the norm. To tackle the rank minimization problem, we choose the following norm which is a regularization of the l^0 -norm:

$$\|A\|_\epsilon = \sum_{i=1}^N \frac{\sigma(A)_i^2}{\sigma(A)_i^2 + \epsilon}. \quad (4)$$

Note that the problem (3) is differentiable as long as $\|\cdot\|_\epsilon$ is differentiable. The gradient can be computed efficiently using the Jacobian matrix of the singular values as expressed in [25]. In practice, we solve this optimization problem using a standard L-BFGS algorithm.

The choice of ϵ can have a big impact on the results. In fact since we are using a gradient descent method the big singular values are in the flat part of $\|\cdot\|_\epsilon$ therefore their gradient will be granted a small weight. On the contrary the singular values in the slope will have a big influence on the minimization. So with an ϵ too small only a few singular values will be minimized but with an ϵ too big a lot of singular values will be minimized not well.

We chose the parameter ϵ such that at the initialization 80 percent of the singular values satisfy $\frac{\sigma_i^2}{\sigma_i^2 + \epsilon} \leq 0.9$.

4 Basis function extraction

Since the probe functions can give redundant information in some shape parts and incomplete information in others, our functional map will map some subspaces of $L^2(S_0)$ with more confidence than others. Using a collection of shapes we would like to extract the most stable subspaces.

For this purpose we propose to use the learned optimal weights D and the resulting estimated functional maps $X_i(D)$ and to identify stably mapped functional subspaces by comparing $X_i(D)$ to the reference maps C_i . The output will be Y an orthonormal basis of $L^2(S_0)$ ordered with decreasing confidence. As we demonstrate in Section 5, in most cases this order remains stable even for maps that are estimated to previously unseen shapes.

4.1 Identifying stable subspaces

The most well-mapped function $y_0 \in L^2(S_0)$ is such that $X_i y_0$ is the closest to $C_i y_0$ for all i . Such function is solution of the problem:

$$y_0 \in \arg \min_{y \in L^2(S_0), \|y\|=1} \sum_{i=1}^N \|(X_i - C_i)y\|_F^2$$

We can then iteratively define an orthonormal basis of $L^2(S_0)$ ordered by decreasing accuracy in the mapping, by solving the following problem:

$$y_{n+1} \in \arg \min_{y \in L^2(S_0), (y, y_j)=0 \forall j \leq n} \sum_{i=1}^N \|(X_i - C_i)y\|_F^2$$

Such a basis can be efficiently computed by considering the singular value decomposition of the matrix:

$$X = \begin{pmatrix} X_1 - C_1 \\ \dots \\ X_N - C_N \end{pmatrix} = U \Sigma V^t.$$

It is well-known that y_j must be equal to singular vectors corresponding to the j^{th} smallest singular value of X . We can, therefore, form a new orthonormal basis Y of $L^2(S_0)$ composed of the singular vectors of X by increasing singular values. This allows us to quantify the quality of the mapping of a functional subspace just by looking at the singular values of X : the smaller the singular values are, the better the mapping.

4.2 Functional map to a test shape using a reduced basis

Now if we are given an extra shape S_{N+1} that does not belong to the training set, we first compute its probe functions and store them in a matrix G_{N+1} . We then compute the functional map X_{N+1} by using the previously solved for weight matrix D . Finally, since we know that X_{N+1} contains some badly mapped subspaces (for example the antisymmetric functions), by using Y_p the p first column of Y , we compute the reduced map X_{N+1}^p

$$X_{N+1}^p = X_{N+1} Y_p : L^2(S_0) \cap L^2(\text{Im}(Y_p)) \rightarrow L^2(S_{N+1})$$

5 Experimental Results

5.1 Functional correspondences

The probe functions used to solve the problem in Eq. (2) are given by various descriptors computed on each shape:

- Heat Kernel Signature [31]

- Wave Kernel Signature [2] at three different variances
- Gaussian and Mean Curvature
- Logarithm of the absolute value of Gaussian and Mean Curvature
- Mesh Saliency [18]

The HKS, WKS and Mesh Saliency are computed at various scales to ensure a wide variability. The curvatures are processed in order to obtain a family of functions. Since the curvatures can have very high peaks we take the logarithm of their absolute value to put more weight on the small curvatures areas. The family of functions is then created by considering the solution of the Heat Diffusion Equation at various times when each function is used as initial heat distribution over the surface.

5.2 Isometric Shape Matching

TOSCA Dataset. We have evaluated our method on the shape matching benchmark TOSCA [6]. For each shape class we use all the available shapes for training, except one for testing and we choose the standard standing pose as shape S_0 . We compare three ways of weighting the probe functions: a unique weight for all the functions, a weight per category of descriptors and one weight per probe function.

For all the experiments we express all functions in the basis given by the first 50 eigenfunctions of the Laplace-Beltrami operator. We compute 50 probe functions divided in 9 categories (WKS is divided in three categories with three different variances) of descriptors. We take 5 functions per category except for the Mesh Saliency where 10 functions are computed. Since all the shapes in TOSCA have an internal symmetry, we cannot hope to recover the entire functional map, and thus Eq. (4) is a reasonable choice of norm.

The experiments follow the pipeline shown in Figure 2. First we learn the optimal weights and extract the ordered basis using the training set of shapes. Second we are given an unknown shape. We use the optimal weights to compute the functional map and the extracted basis to suppress the badly mapped function subspaces. The L-BFGS algorithm used to solve the optimization problem in Eq. (3) is initialized with the naive functional maps solution of (1) with $\alpha = 10^{-3}$. As several methods using functional maps [22, 24] have been proved to be more efficient than the state-of-the-art methods, we compare all of the functional maps and subspaces computed with our method to the baseline “naive” map, obtained using the identity matrix D , which correspond to the original method described in [22].

Performance The proposed approach was implemented in MATLAB. Note that the number of vertices of each shape has no effect on the performance since all the functions are expressed in a reduced function basis. The most time consuming task in our pipeline is the training part which requires to solve a difficult non-linear optimization problem (3). The processing cost is dominated by the computation of the gradient of the energy, which is done by solving two linear




Fig. 3: Left: Optimal weights for different strategies after training with 9 cats. Right: Effect of the different weights on the distribution of the singular value of the difference $X_{N+1}(D) - C_{N+1}$.

systems for each shape of the training set at each iteration. However, the contributions of each shape to the gradient are independent so this can be done in parallel. The learning process with a training set of 10 shapes took about 45 min on an Intel i7 processor without parallelization.

Optimal Weights Figure 3 (left) shows the weights obtained after solving the problem in Eq. (3) with a training set composed of 9 cats. To demonstrate the importance of weighting the probe functions on the quality of the functional map, we study the distribution of the singular values of the difference $X_{N+1}(D) - C_{N+1}$ for the different learned weights. In Figure 3 (right) each curve depicts the percentage of singular values below the threshold given on the x -axis. For the perfect map, all singular values would be zero. As can be seen, the functional maps with the optimal weights have a bigger concentration of small singular values than the naive functional map. Therefore there exists a bigger functional subspace on which these functional maps provide a good approximation of the ground truth. Note that the naive map has no small singular values and is indeed a very bad approximation.

Stable subspaces From the naive maps and functional maps with optimal weights, we extract four function bases ordered by decreasing stability. The most stable functions for each case are shown in Figure 4. Even the most stable functions from the naive maps are not mapped very accurately since they are very bad approximation of the truth. For the other bases the functions seem consistent with the information we would expect from descriptors as HKS and WKS: a distinction between flat area (body) and salient area (legs, tail, head). Note that even with only one weight we are able to retrieve meaningful stable areas.

We also evaluate the extracted functional basis by computing the difference between the ground truth map and our approximation on the unseen shape:

$$\epsilon_i = \|(X_{N+1}(D) - C_{N+1})y_i\|_2, \quad (5)$$

where y_i is the i^{th} function of the extracted basis. We compare this error for the three weighting strategies with the naive map in Figure 5 (left). The extracted

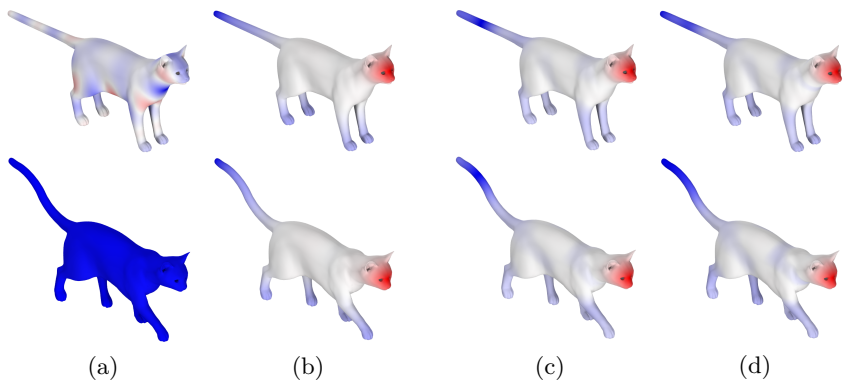


Fig. 4: Visualization of the first component of the extracted basis. First row: the reference shape. Second row: transfer using different functional maps: (4a) naive map, (4b) unique weight, (4c) one weight per category of function, (4d) one weight per functions.

basis was ordered by decreasing quality on the training set. Note that this order is still preserved on the unknown shape for the 25 first functions of the basis. Most of all we are able to identify the worst mapped subspaces, which can be safely removed.

Despite only estimating the functional maps on a subspace of the full functional space, we converted them to point-to-point correspondences using the method described in [22]. Figure 5 right compares the quality of the correspondences before and after reducing the space dimension from 50 to 25. We obtain better results with our learned weight than with the naive map. For the weighted maps, the reductions perform better or similarly than the full maps. Thus our basis extraction manages to identify correctly the most stable subspaces. For the naive map our reduction space strategy fails as there is no well-mapped subspace.

5.3 Non-Isometric Shape Matching

Until now we have assumed the deformation to be nearly isometric. Our algorithm to find the optimal set of probe functions and the extraction of the most stable subspaces do not contain any explicit knowledge of the type of deformation. In fact, this assumption is only used to construct the least-square problem (1). Which means that our framework can be adapted to any kind of deformation model as long as we have a consistent way of computing a functional map from probe functions.

To test this case, we consider a man or a gorilla and 12 women in different poses from the TOSCA dataset. The ground truth functional maps are computed using a thousand user-picked correspondences. The meshes have different number of vertices and different connectivity. We train our method on 10 poses and use the last for testing.




Fig. 5: Left: Accuracy of the extracted function basis measured with Eq. (5). Right: Comparison between full map (plain lines) and reduced map (dash lines). Symmetric correspondences are considered correct.

The resulting problem is very noisy for two reasons. First, the ground truth functional maps are computed from sparse correspondences, and therefore can be inaccurate on some functional subspace. However, the use of a collection of shape for training allows us to remove this noise. Second, since all the probe functions used are designed for isometric deformation, few are going to contain useful information.

In order to introduce a wide variability of functions we pre-compute on each shape 50 basis functions, 310 probe functions and we put a weight on each probe function. The descent algorithm is initialized as in the previous experiment. Figure 7 (left) shows the most stable functions learned from the training maps. Each of these functions is also mapped to a previously unseen pose using the “ground truth” map converted to a functional map. Note that the functions are badly transferred, due to the incompatibility of the LB basis and the noise in the input maps. Compare this with Figure 7 (right) where the probe functions have been weighted using our method. The stable functions indicate the head, the hands and the feet to be the most stable area. Besides, these functions are correctly mapped on a new shape using *a computed* functional map with the learned weights. Clearly, the fact that we use a collection helps removing the noise of the input data. Thus, our method is able to correctly identify the most stable functional subspace with mild assumptions on the underlying deformation.

6 Conclusion

In this work, we presented a method to learn the most informative descriptors for non-rigid shape matching, from a given set of shape correspondences. Our method is purely intrinsic and allows us to obtain high quality *consistent* correspondences to new, unseen shapes, and to identify the most reliably mapped functional subspaces. The approach is flexible and can potentially be applied to scenarios that lack a good theoretical deformation model, as demonstrated by meaningful even non isometric deformation. One of its weaknesses is a relatively high cost for the training, and in the future, we plan to explore more efficient optimization strategies.

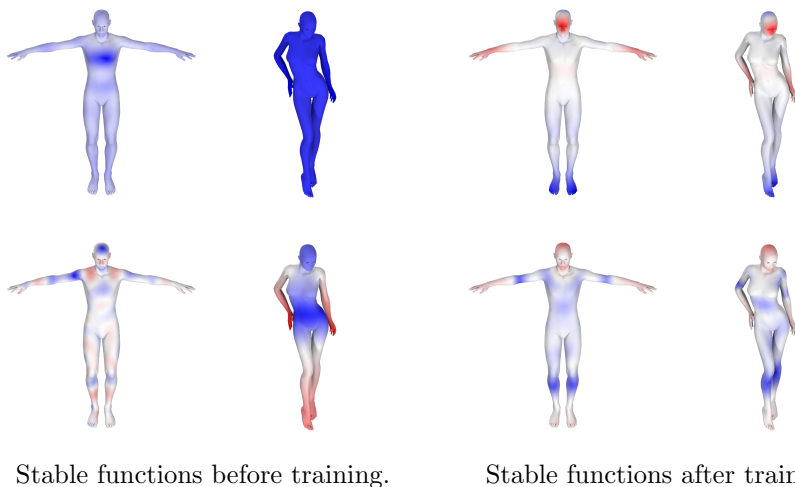


Fig. 6: Visualization of the first two components of the extracted basis on the reference shape (man) then mapped to the unseen shape (woman) using the functional map without optimization (left) and with the learned weight (right).

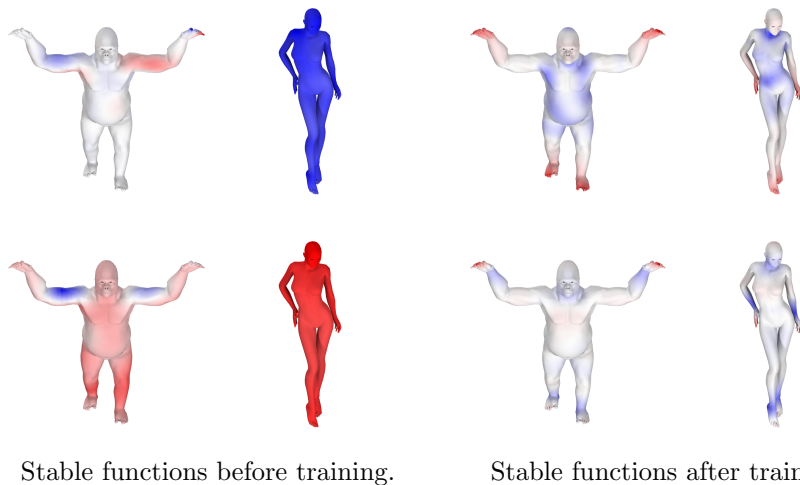


Fig. 7: Visualization of the first two components of the extracted basis on the reference shape (gorilla) then mapped to the unseen shape (woman) using the functional map without optimization (left) and with the learned weight (right).

Acknowledgements The authors would like to acknowledge the support of the FUI project “TANDEM 2”, the French Direction Générale de l’Armement (DGA), the CNRS chaire d’excellence, Marie Curie CIG-334283-HRGP, a Google Faculty Research Award, and the “EANOI” joint FWF n. I1148 / ANR-12-IS01-0003 project.

References

1. 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). vol. 24, pp. 408–416. ACM (2005)
2. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: Computer Vision Workshops (ICCV Workshops). pp. 1626–1633. IEEE (2011)
3. Azencot, O., Ben-Chen, M., Chazal, F., Ovsjanikov, M.: An operator approach to tangent vector field processing. *Computer Graphics Forum* 32(5), 73–82 (2013)
4. Bookstein, F.L.: Landmark methods for forms without landmarks: morphometrics of group differences in outline shape. *Medical image analysis* 1(3), 225–243 (1997)
5. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103(5) (2006)
6. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Numerical geometry of non-rigid shapes. Springer (2008)
7. Chen, X., Saparov, A., Pang, B., Funkhouser, T.: Schelling points on 3d surface meshes. *ACM Trans. Graph. (TOG)* 31(4), 29 (2012)
8. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Training models of shape from sets of examples. In: BMVC92, pp. 9–18 (1992)
9. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models-their training and application. *Computer vision and image understanding* 61(1), 38–59 (1995)
10. Dryden, I.L., Mardia, K.V.: Statistical shape analysis, vol. 4. John Wiley & Sons New York (1998)
11. Gao, X., Su, Y., Li, X., Tao, D.: A review of active appearance models. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 40(2), 145–158 (2010)
12. Hasler, N., Stoll, C., Sunkel, M., Rosenhahn, B., Seidel, H.P.: A statistical model of human pose and body shape. In: *Computer Graphics Forum*. vol. 28, pp. 337–346. Wiley Online Library (2009)
13. Huang, Q.X., Adams, B., Wicke, M., Guibas, L.J.: Non-rigid registration under isometric deformations. *CGF (Proc. SGP)* 27(5), 1449–1457 (2008)
14. Huang, Q., Wang, F., Guibas, L.: Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.* 33(4), 36:1–36:11 (jul 2014)
15. Kalogerakis, E., Hertzmann, A., Singh, K.: Learning 3d mesh segmentation and labeling. *ACM Trans. Graph.* 29(4), 102 (2010)
16. Kilian, M., Mitra, N.J., Pottmann, H.: Geometric modeling in shape space. In: *ACM Transactions on Graphics (TOG)*. vol. 26, p. 64. ACM (2007)
17. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *ACM TOG (Proc. SIGGRAPH)* 30(4) (2011)
18. Lee, C.H., Varshney, A., Jacobs, D.W.: Mesh saliency. In: *ACM Transactions on Graphics (TOG)*. vol. 24, pp. 659–666. ACM (2005)

19. Lipman, Y., Funkhouser, T.: Mobius voting for surface correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28(3) (Aug 2009)
20. Litman, R., Bronstein, A.: Learning spectral descriptors for deformable shape correspondence (2013)
21. Mémoli, F.: On the use of Gromov-Hausdorff Distances for Shape Comparison. In: *Symposium on Point Based Graphics*. pp. 81–90 (2007)
22. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31(4), 30:1–30:11 (Jul 2012), <http://doi.acm.org/10.1145/2185520.2185526>
23. Ovsjanikov, M., Mériçot, Q., Mémoli, F., Guibas, L.: One point isometric matching with the heat kernel. *CGF* 29(5), 1555–1564 (2010)
24. Ovsjanikov, M., Mériçot, Q., Pătrăucean, V., Guibas, L.: Shape matching via quotient spaces. *Computer Graphics Forum* 32(5), 1–11 (2013)
25. Papadopoulo, T., Lourakis, M.I.: Estimating the jacobian of the singular value decomposition: Theory and applications. In: *Computer Vision-ECCV 2000*, pp. 554–570. Springer (2000)
26. Rodolà, E., Bulò, S.R., Windheuser, T., Vestner, M., Cremers, D.: Dense non-rigid shape correspondence using random forests. In: *Proc. CVPR* (2014)
27. Rustamov, R.M., Ovsjanikov, M., Azencot, O., Ben-Chen, M., Chazal, F., Guibas, L.: Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics (TOG)* 32(4), 72 (2013)
28. Sahilliođlu, Y., Yemez, Y.: Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum* 30(5), 1461–1470 (2011)
29. Shapira, N., Ben-Chen, M.: Cross-collection map inference by intrinsic alignment of shape spaces. In: *Computer Graphics Forum*. Wiley Online Library (2014)
30. Sumner, R.W., Zwicker, M., Gotsman, C., Popović, J.: Mesh-based inverse kinematics. In: *ACM SIGGRAPH*. pp. 488–495 (2005)
31. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Computer Graphics Forum*. vol. 28, pp. 1383–1392. Wiley Online Library (2009)
32. Tam, G.K., Cheng, Z.Q., Lai, Y.K., Langbein, F.C., Liu, Y., Marshall, D., Martin, R.R., Sun, X.F., Rosin, P.L.: Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *Visualization and Computer Graphics, IEEE Transactions on* 19(7), 1199–1217 (2013)
33. Tevs, A., Bokeloh, M., Wand, M., Schilling, A., Seidel, H.P.: Isometric registration of ambiguous and partial data. In: *Proc. CVPR*. pp. 1185–1192 (2009)
34. Van Kaick, O., Tagliasacchi, A., Sidi, O., Zhang, H., Cohen-Or, D., Wolf, L., Hamarneh, G.: Prior knowledge for part correspondence. *Computer Graphics Forum (Proc. Eurographics)* 30(2), 553–562 (2011)
35. Van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. *Computer Graphics Forum* 30(6), 1681–1707 (2011)
36. Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., Van Kaick, O., Tagliasacchi, A.: Deformation-driven shape correspondence. In: *Proc. SGP*. pp. 1431–1439 (2008)