

Robust Shape Collection Matching and Correspondence from Shape Differences

Aharon Cohen¹ Mirela Ben-Chen¹

¹Technion - Israel Institute of Technology

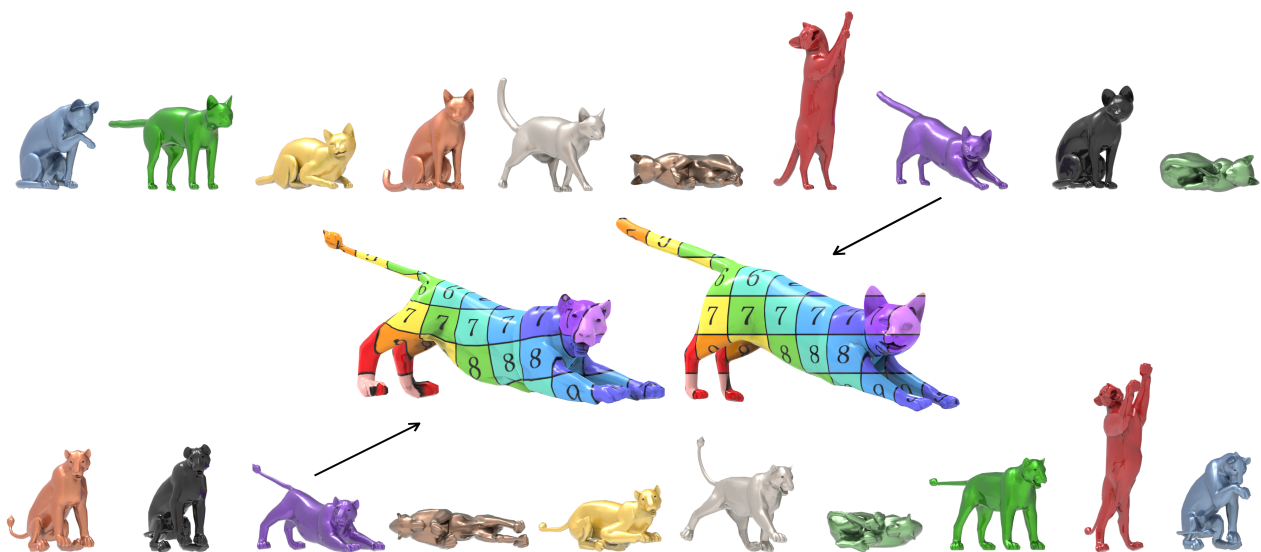


Figure 1: Given two collections of shapes in various poses (top and bottom) and the intra-collection maps, our algorithm automatically finds a matching (color coded) between the poses. Using this matching, we further find an inter-collection correspondence between two automatically-chosen shapes, one from each collection (center). We can therefore automatically extract a high-quality non-isometric correspondence solely from the two collections, without any user input.

Abstract

We propose a method to automatically match two shape collections with a similar shape space structure, e.g. two characters in similar poses, and compute the inter-maps between the collections. Given the intra-maps in each collection, we extract the corresponding shape difference operators, and use them to construct an embedding of the shape space of each collection. We then align the two shape spaces, and use the knowledge gained from the alignment to compute the inter-maps. Unlike existing approaches for collection alignment, our method is applicable to small and large collections alike, and requires no parameter tuning. Furthermore, unlike most approaches for non-isometric correspondence, our method uses solely the variation within the collection to extract the inter-maps, and therefore does not require landmarks, descriptors or any additional input. We demonstrate that we achieve high matching accuracy rates, and compute high quality maps on non-isometric shapes, which compare favorably with automatic state-of-the-art methods for non-isometric shape correspondence.

CCS Concepts

• Computing methodologies → Shape analysis;

1. Introduction

A shape space contains variations of a given 3D model, for example, a sampling of an animation, or a character in different poses. Such spaces arise, e.g., when animators generate blend shapes for

standard poses (smile, frown, A-pose, T-pose, etc.) or use rigged models for generating walking or running cycles. In many applications, a few such shape spaces are given (e.g. a walking cycle of a man and a woman), and it is required to transfer information

between them. For example, transferring shape annotations such as pose labeling, or transferring point-wise data such as texture.

Often, it is possible to automatically obtain a high quality correspondence between variations of a single model, i.e., maps within the shape collection, which we denote by *intra-maps*. On the other hand, correspondences between shapes in different collections, which we denote as *inter-maps*, are harder to compute, as these models will often have large non-isometric deformations.

Automatic computation of correspondences between non-rigid and non-isometric shapes is an active research problem. Most existing methods for computing a correspondence between non-isometric shapes require some additional semantic input, such as corresponding landmarks, and are therefore not fully automatic. We, on the other hand, leverage the cues of the *variations* within the collection as the semantic information, and use them to design a *completely automatic* method.

We draw our inspiration from a recent approach to this problem [SBC14], which similarly leverages shape variations as the semantic cues. Unfortunately, many aspects of that approach introduce technical difficulties which make the method not robust and impractical in many cases, derailing the hope for a completely automatic method. Specifically, a large sampling of shapes is required in both collections, as well as hand tuning of multiple dataset-dependent parameters. Practically, the existing approach yields low quality maps which are not on-par with state of the art methods.

While working in the same general setting, we propose very different design choices, leading to a considerably more robust system. Our method is simple to implement, completely automatic, and applicable to any mesh topology, yielding significantly better results. Additionally, our method extracts high quality point-wise maps between non-isometric shapes, generating maps which surpass state-of-the-art automatic methods for non-isometric shape correspondence. Our approach is applicable to both benchmark data-sets, e.g. FAUST, as well as rigged models available for purchase from 3D modeling websites. Finally, it is important to note that the shape collections we can handle are not limited to the same character in different poses, but can rather be sets of shapes with any variations yielding similar enough shape spaces. For instance, a collection can include *different* people in the same pose, then matched to a collection of the same people in a different pose.

1.1. Related Work

Shape Collections. Collections of shapes are useful for a variety of applications. Some examples include deformation transfer, aiming to produce shapes with desired deformations that can be induced by given deformations in another collection. Some methods include learning models for 3D shape processing and shape reconstruction [GYQ*18, HRA*19]. Another example is *creating* a collection from one shape using modal analysis of the hessian of the mesh as described in [HWAG09]. Other aspects of shape collection analysis exploit the idea that the composition of maps along cycles should be identity maps, which led to the map synchronization problem, taking maps between pairs of objects and returning improved maps that are consistent along cycles [SLHH18]. This notion is further exploited in [NBCW*11, CRA*17] in order to obtain

high quality intra-maps, that we require as an input, and which are easier to obtain than the inter-maps we compute. Other approaches use modular latent spaces, based on nonlinear embedding spaces, to find the correspondences within a collection [GSDG18].

In [ROA*13] some approaches to exploring shape collections are presented, including browsing shapes by a user-defined region of interest, taking advantage of the localization property of the *shape difference operator*. Shape analogies are also used in [ROA*13] to match corresponding shapes from two collections. However, they use a brute-force search among all the *permutations* of the shapes that best aligns them with the other collection, limiting the applicability to very small collections. The shape difference operator has been further investigated and analyzed in terms of stability [HCO18] and optimal shape collection representation [HAGO19].

Our work generalizes the method proposed by [SBC14]. In that approach, given two shape collections, each representing a shape space, it is assumed that the intra-maps are given, and they are used to compute the shape variations, represented by *shape difference operators*. These variations are then considered as points on a high dimensional manifold, and a non-linear dimensionality reduction approach is used to generate a low-dimensional embedding of the shape spaces. The shape spaces are then aligned using affine point cloud alignment. Finally, an inter-map is computed using the alignment. This approach has several severe limitations. First, a large sampling of shapes is required to reliably represent the diffusion map, and hand tuning of parameters is required to align the shape spaces using an affine map. Furthermore, the existing approach fails to obtain high accuracy rates for aligning the shape spaces, and in addition the resulting inter-maps are of low quality and are not comparable to other correspondence methods.

We, on the other hand, have made critically different design choices, leading to an algorithm that is more robust and yields maps comparable to other automatic methods. Specifically, we leverage the fact that in some cases the shape difference operators lie on a manifold which has a closed form expression for *geodesic distances*, therefore, we use these distances as a better representation of the shape space structure. Second, we use a *linear* dimensionality reduction approach, which is far less sensitive to the number of shapes in the collection. Furthermore, we use rigid alignment using a robust convex relaxation, which is parameter-free. Finally, we add a regularization to the shape correspondence formulation, and a post processing, which yield considerably better pointwise maps.

Automatic Correspondences. Computing shape correspondence using automatic algorithms, i.e. when no landmarks or user input are given, acts as a benchmark for the second part of our algorithm, aiming to obtain the correspondences between shapes in different collections. Blended intrinsic maps (BIM) [KLF11] is an automatic method to recover the point-to-point map between two given shapes. It often yields excellent results, however it is restricted to genus zero shapes. Bijective and continuous ICP (BCICP) [RPWO18] tackles this problem as well and is able to produce correspondences without landmarks. We compare our algorithm, which uses the shape differences within the collection as the additional information required for automatically extracting an inter-collection map, to these methods and demonstrate that we achieve a better performance than both.

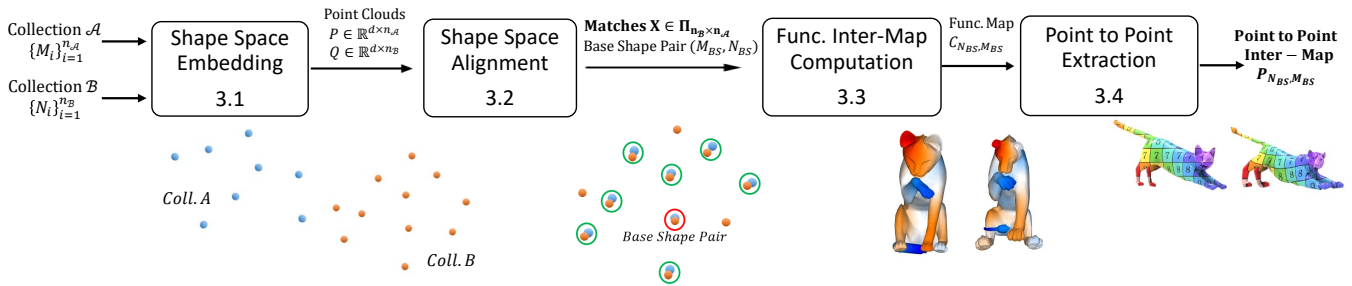


Figure 2: Block diagram of our algorithm, see the outline in Section 1.2.

Matching Problems and Procrustes Analysis. In order to solve the matching problem, i.e. finding the matching pairs from the two collections, one can make use of the distance matrix between the shape differences of each collection. The distances represent the differences between the variations in the collection. One approach to solve the matching problem is to use the method presented in [KKBL15] in order to find the best permutation on the distance matrices' rows and columns such that they have the same structure. The retrieved permutation defines the pairing of shapes we wish to obtain. Unfortunately, this approach is not feasible for large data-sets where the distance matrices are large. An improved method has been proposed in [DML17] (see also references within for additional approaches to quadratic matching).

Instead of aligning the distance matrices, it is also possible to generate a low dimensional embedding of the point clouds, and use Procrustes matching (PM) for the registration. We chose to take this approach, first using linear dimensionality reduction to generate a point cloud representation of the shape collections, and then convex semidefinite programming (SDP) relaxation [MDK*16] for the Procrustes matching of the point clouds. This approach has a few advantages. First, linear dimensionality reduction requires less parameters. Second, the SDP relaxation scheme is tight leading to close to optimal results, and finally, it is very efficient, enabling us to handle larger collection sizes.

1.2. Method Outline

Our algorithm follows the general structure suggested in [SBC14], yet with some critically different design choices. We first describe the general pipeline (see Figure 2), and then explain each component separately in Section 3.

Input: Two shape collections that represent two shape spaces of different models, not necessarily with the same number of shapes, and the intra-maps within each collection.

Output: (1) An injective function from the smaller collection to the larger one, which represents the *matching* pairs, and (2) a point-to-point *inter-map* for any shape in one collection to any shape in the other.

Algorithm:

1. Use the input intra-maps to construct the shape differences between the input shapes in the same collection (Section 2.2).

2. Use the shape differences to construct a low-dimensional shape space embedding for both collections (Section 3.1).
3. Align the two shape spaces using Procrustes analysis to obtain the matching pairs and automatically determine an *optimal base shape* in each collection (Section 3.2).
4. Use the matches and the base shapes to compute a functional inter-map for the base shape pair (Section 3.3).
5. Recover a point-to-point inter-map from the functional map (Section 3.4).

1.3. Contributions

Our main contributions are:

- Matching corresponding pairs from two shape collections with a high rate of accuracy.
- Automatically obtaining a high-quality non-isometric inter-map between any shape in one collection to any shape in the other.

2. Background

2.1. Functional Maps

Functional maps [OBCS*12] have been widely used in many geometry processing applications, especially for shape correspondence. The main observation in this framework is that we can look at how *functions* on one shape are transformed to the other, rather than finding a point-to-point map between the shapes. The space of functions can often be well represented using a compact, yet informative, *functional basis*, allowing us to represent a map in this space as a *change of basis* linear operator, i.e., as a compact matrix.

Given two surfaces M and N , with a map $T : N \rightarrow M$ between them, a map between function spaces $F : L^2(M) \rightarrow L^2(N)$ is induced. Here, $L^2(\cdot)$ represents the set of square integrable real valued functions defined on the surface. F is called the functional map, mapping functions defined on M to functions defined on N , i.e. $g = F(f) = f \circ T$ where $f : M \rightarrow \mathbb{R}$ and $g : N \rightarrow \mathbb{R}$. F is a linear transformation between function spaces, and given reduced bases Ψ_M, Ψ_N of dimensions k_M, k_N , for M, N , respectively, is represented as a matrix $C \in \mathbb{R}^{k_N \times k_M}$.

2.2. Shape Difference Operators

Given two shapes and the map between them, the shape difference operator [ROA*13] captures the variations between the shapes. It

constitutes the main tool in our algorithm, allowing us to compare differences between shapes. This operator stores information about how a function on one of the shapes should be modified, such that the *inner product* of any two functions defined on the first surface equals to the inner product of the mapped functions on the other surface. In that way, it captures information about how and where one shape differs from the other. In addition, by modifying the inner product, we can quantify different types of differences, or distortions, between the shapes.

Given two shapes M and N , and a functional map $F : L^2(M) \rightarrow L^2(N)$, let $h_M : L^2(M) \times L^2(M) \rightarrow \mathbb{R}$ and $h_N : L^2(N) \times L^2(N) \rightarrow \mathbb{R}$ be inner products defined on M and N respectively, acting on two functions on the shape. Then there exists [ROA*13, Thm 2] a *unique* linear operator $D_{h_M, h_N} : L^2(M) \rightarrow L^2(M)$, denoted as the *shape difference operator*, satisfying $h_M(f, D_{h_M, h_N}(g)) = h_N(F(f), F(g))$ for any two functions $f, g : M \rightarrow \mathbb{R}$. Note that this operator depends only on the chosen inner products on M and N and the functional map F . Moreover, it is a linear self-map on the space of functions over M . Thus, the operators $D_{h_M, h_{N_1}}, D_{h_M, h_{N_2}}$ have the same domain and range $L^2(M)$, allowing us to compare the operators even if $N_1 \neq N_2$. Two inner products that are of interest, are the *area-based* inner product $h^a(f, g) = \int_M f(x)g(x)d\mu(x)$ and the *conformal* inner product $h^c(f, g) = \int_M (\nabla f(x), \nabla g(x))_x d\mu(x)$. Intuitively, the first encodes variations in area due to the map, and the second encodes variations in angles.

Given a choice of reduced basis Ψ_M, Ψ_N , the shape difference operators are represented as matrices of dimensions $k_M \times k_M$. Often, the eigenvectors of the Laplace-Beltrami operator corresponding to the lowest eigenvalues are chosen for the basis Ψ , as this choice leads to a multi-scale basis, which can represent well smooth functions using only a small number of basis functions. In this case, the explicit expressions for the shape difference operators are [ROA*13, Eq. (4)]:

$$V_{M,N} \triangleq D_{h_M^a, h_N^a} = F^T F \quad (1)$$

$$R_{M,N} \triangleq D_{h_M^c, h_N^c} = (D^M)^{-1} F^T D^N F, \quad (2)$$

where F is the functional map represented in the bases Ψ_M, Ψ_N , and D^M is a diagonal matrix of the lowest non-zero k_M eigenvalues of the Laplace-Beltrami operator of shape M .

3. Algorithm

3.1. Shape Space Embedding

Given two shape collections $\mathcal{A} = \{M_i\}_1^{n_A}$, $\mathcal{B} = \{N_i\}_1^{n_B}$, we assume that each collection includes shapes sampled from a single shape space. We further assume, that the samplings are “compatible” in the sense that they include similar shape variations. Then, given a choice of two *base* shapes, one in each collection, $M_{BS} \in \mathcal{A}, N_{BS} \in \mathcal{B}$, we compute for each collection separately the shape difference operators with respect to the base shapes. These operators encode information about the variability of the collection, and the *distances* between them encode the structure of the shape space.

We first select a *random* shape in each collection as the *base shape* for the following part. Later we show that the base shape can indeed be chosen randomly without affecting the performance of the algorithm.

Base Shape




Figure 3: A collection of facial expressions (top) with its shape differences’ principal components and their explained variance (bottom). The sign is arbitrary, thus both dark blue and dark orange indicate high variation, mainly around the mouth and eyes area.

3.1.1. Shape difference operators computation

Once we have chosen a random base shape for each collection, we compute the shape difference operators for all the shapes in the collection with respect to this base shape. As shown in the previous section, shape difference operators can be computed using either the area-based inner product (Eq. (1)) or the conformal inner product (Eq. (2)). We later show how the choice of inner product depends on the nature of the collection. However, for typical datasets, we choose the conformal inner product since the variations within the collection are mostly non-conformal.

Given the intra-map between shapes within the collection, we compute the functional map between the base shape and any other shape in the set using k_1 eigenvectors of the Laplace-Beltrami operator of the base shape, and k_2 eigenvectors for the other shape. A collection of n shapes yields a set of n shape difference matrices of size $k_1 \times k_1$. Each of them represents the modification that needs to be done on functions on the base shape such that inner products on the base shape equal to inner products on the other shape.

Figure 3 shows the principal components of the shape difference operators for a given collection using the area-based inner product. The principal components are visualized by testing how they act on a constant function, thus visualizing the main variations between the shapes and their location. The percentage of the total variance explained by each principal component is shown as well, demonstrating the significance of each component in the collection.

Now that we have the set of shape difference operators, our goal is to embed them in a low dimensional space such that the information of the distances between the operators is best preserved.

3.1.2. Distance matrix construction

To construct a distance matrix for a given collection, we treat separately the area-based and conformal shape differences, as they have a different structure as operators.

Area based. The area-based shape difference operator, given by Eq. (1), is a *symmetric positive-definite* (SPD) matrix, as long as

the functional map F is an invertible matrix. We therefore use the manifold of SPD matrices to compute the *geodesic distances* between the shape difference operators. The manifold of SPD matrices has a unique shortest geodesic curve between any two points, and its length has a closed form expression. Specifically, given two SPD matrices, V_1, V_2 , their geodesic distance on the SPD manifold is [Bha09, Eq. (6.14)]:

$$d_{SPD}^2(V_1, V_2) = \sum_{i=1}^n \log^2 \left(\lambda_i(V_1 V_2^{-1}) \right), \quad (3)$$

where $\lambda_i(V)$ is the i^{th} eigenvalue of the matrix V .

Conformal. The conformal shape difference matrices, given by Eq. (2), are not SPD, and we therefore use the Euclidean distances between the vectorized matrices.

Figure 4 shows the distance matrices obtained for the Sumner data-set (Figure 1), where the shapes are ordered correspondingly for better visualization. The distance matrix for the set of cats (collection \mathcal{A}) has a similar structure to the distance matrix computed for the set of lions (collection \mathcal{B}). In both sets we use the shape differences operators derived for the conformal inner product, since the variations within each collection are mainly non-conformal. Hence, in this case, the conformal shape differences capture more reliably the information about the variations compared to the area-based shape differences.

3.1.3. Dimensionality reduction

To solve the matching problem, we first apply multidimensional scaling (MDS) [Mea92] on the distance matrix obtained in order to embed the operators in a d -dimensional space, the *shape space embedding*. The dimension d is determined such that the energy accumulated in this d -dimensional space is at least β of the total energy, for both shape collections. The energy of each dimension is defined as the variance explained by this dimension using PCA. For instance, choosing $\beta = 0.95$ means that we use at least 95% of the total energy for both collections. In this case, we lose only up to 5% of the information regarding the variations, but store it in a space of dimension in the range of 3-8 for typical data-sets.

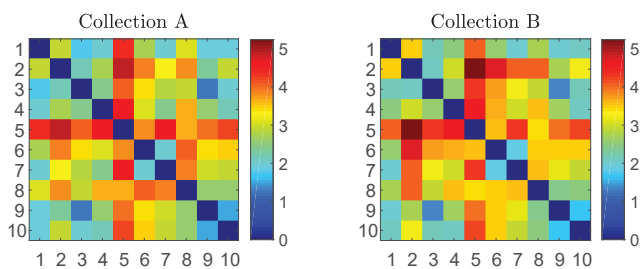


Figure 4: Distance matrices of the collections from Figure 1, for the conformal shape differences (both collections are ordered correspondingly for better visualization). Note that the cats (collection \mathcal{A}) and the lions (collection \mathcal{B}) have similar distance matrices, which we use for matching the collections.

3.2. Shape Space Alignment

After constructing the shape space embedding of each collection, our goal is to *align* them such that we can match shapes from the two collections. We treat the two shape space embeddings as two point clouds, and assume that the distance matrices of the two collections are similar, so that we can use *rigid alignment*. In this setting, the alignment problem is known as the *Procrustes matching problem*, formulated as follows.

We are given two point clouds P and Q of dimension d and n_1, n_2 points respectively, $P \in \mathbb{R}^{d \times n_1}, Q \in \mathbb{R}^{d \times n_2}$, where we assume, without loss of generality, that $n_1 \leq n_2$. Our goal is to find a linear isometry (an orthogonal transformation) $R \in \mathcal{O}(d)$ and a permutation $X \in \Pi_{n_2 \times n_1}$, minimizing the distance between the point clouds:

$$\begin{aligned} d(P, Q) &= \min_{X, R} \|RP - QX\|_F^2 \\ \text{s.t.} & \\ X &\in \Pi_{n_2 \times n_1}, R \in \mathcal{O}(d). \end{aligned} \quad (4)$$

Since $n_1 \leq n_2$, this formulation matches every point in P to exactly one point in Q , where some points in Q can remain unmatched.

This is in general a difficult, non-convex problem, however, recently, a very effective convex relaxation to the problem has been proposed [MDK*16], which we leverage to find the alignment.

Let X^*, R^* be the optimal solutions of the optimization problem (4), then the alignment error $d(P^*, Q^*)$ provides a quantitative measure of the success of the alignment process. We define the *normalized alignment error*:

$$\hat{e} = \frac{\|R^*P - QX^*\|_F}{\|P\|_F + \|Q\|_F}, \quad (5)$$

which is invariant to the dimensions of the point clouds. We will later use this error to evaluate the matching accuracy of our results.

Optimal Base Shape Selection. The point clouds $\tilde{P} = R^*P \in \mathbb{R}^{d \times n_1}$ and $\tilde{Q} = QX^* \in \mathbb{R}^{d \times n_1}$ include only pairs of matching shapes *after* the alignment. The closest pair of points in \tilde{P}, \tilde{Q} are the two shapes which are most likely to have been matched correctly. Therefore, we select these as the new base shapes M_{BS}, N_{BS} , and the shape difference operators are recomputed with respect to them, yielding a better input for the next part of the algorithm.

3.3. Functional Inter-Map Computation

The next goal of the algorithm is to compute the functional map between the newly chosen base shapes of the two collections, thus retrieving the inter-map correspondence.

Shape analogies. We use the shape analogies constraint presented in [SBC14], taking advantage of the similarity between the shape difference matrices of matching pairs. More precisely, let $\{(M_i, N_i) \mid M_i \in \mathcal{A}, N_i \in \mathcal{B}, i = 1, \dots, n\}$, be the $n = \min\{n_{\mathcal{A}}, n_{\mathcal{B}}\}$ matched pairs extracted by the shape space alignment. Recall that (M_{BS}, N_{BS}) are the pair of base shapes of the two collections, defined as the closest matching pair. Our goal is to find a matrix C , the functional map between the base shapes.

Our main assumption is that the shape difference operators of

matching shapes in both collections, i.e., V_{M_{BS}, M_i} and V_{N_{BS}, N_i} act similarly on functions. Hence, first applying a map between the collections, and then applying the shape difference operator, should yield a similar function to first applying the shape difference and then mapping the function to the other collection.

Regularization. Differently from [SBC14], we observe that without a regularization term the optimization results are unstable, leading to poor functional maps. This is especially evident for areas on the shapes with none or few variations, where the shape differences do not hold any information. To handle this problem, we adopt the regularization term that forces the functional map to commute with the Laplace-Beltrami operator [OBCS*12].

Optimization problem. Our optimization problem is therefore:

$$\min_{C \in \mathbb{R}^{k_B \times k_A}} \sum_{i=1}^n \left(\|CV_{BS,i}^A - V_{BS,i}^B C\|_F^2 + \|CR_{BS,i}^A - R_{BS,i}^B C\|_F^2 \right) + \alpha \|C\Delta_{BS}^A - \Delta_{BS}^B C\|_F^2, \quad (6)$$

where $k_A = k_B = k_1$, n is the number of matching pairs, $V_{BS,i}^A$ is the area based shape difference from the base shape in \mathcal{A} to $M_i \in \mathcal{A}$, Δ_M^A is a diagonal matrix with the eigenvalues of the Laplace-Beltrami operator of shape $M \in \mathcal{A}$, and α is a parameter controlling the regularizer.

To solve this optimization problem we observe that the objective is linear in the elements of C , resulting in a homogeneous linear system that is solved using SVD.

As proposed in [OBCS*12, SBC14], we apply iterative refinement as post-processing to the minimizer of Eq. (6) in order to refine the solution so that it better represents a point-to-point map. The refinement process is effectively ICP in eigenspace, where the solution of the optimization problem is used as initialization.

Finally, we compute the functional map, F_{N_j, M_i} between any two shapes in the two collections, $M_i \in \mathcal{A}$ and $N_j \in \mathcal{B}$ using the following composition:

$$F_{N_j, M_i} = F_{j, BS}^B C_{N_{BS}, M_{BS}} F_{BS, i}^A, \quad (7)$$

where the functional map $C_{N_{BS}, M_{BS}} : L^2(M_{BS}) \rightarrow L^2(N_{BS})$ is the computed inter-map, i.e. the minimizer of Eq. (6), and the functional maps $F_{BS, i}^A : L^2(M_i) \rightarrow L^2(M_{BS})$ and $F_{j, BS}^B : L^2(N_{BS}) \rightarrow L^2(N_j)$ are the input intra-maps within the collections.

In Figure 5 we visualize the computed functional map and demonstrate it with a function transferred from the apple (center) to the orange (right) for a data-set of rigged shapes (top and bottom). Note that a high quality map can be derived for the areas where the shape differences hold the information about the shapes, i.e. areas with more variations such as the "face" of the apple and the orange. However, the map might be distorted in areas where there is no information in the shape differences, such as the side of the apple and orange, since there are no variations within the set in these regions. Our last step is to recover the point-to-point map, which additionally alleviates these issues (Fig. 15).

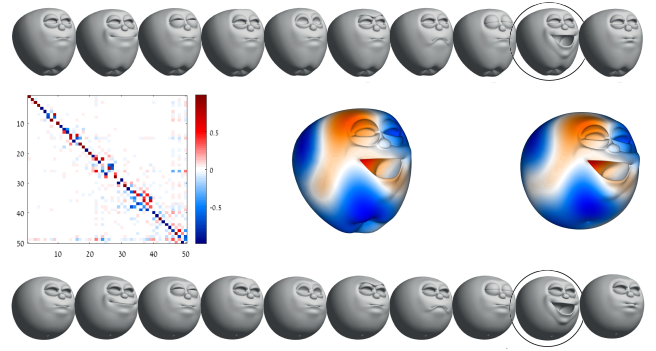


Figure 5: The computed functional map is of high quality in areas with informative shape differences (the "face"), and distorted where there is less variation (the sides). When recovering the point-to-point map these distortions are fixed (Fig. 15).

3.4. Recovering a Point-to-Point Inter-Map

The final step of our algorithm is to produce a point-to-point map using the functional map we computed. As we aim for a high quality map that can be used to transport textures, we use a recent map reconstruction approach [EBC17] to obtain a vertex-to-point map.

Some data-sets require post-processing of the point-to-point map using Reversible Harmonic Maps (RHM) [ESBC19]. This post-processing scheme is especially effective in areas on the shape where there are few or no variations within the collection, such that the shape difference operators contain no information for these areas. As a result, in these regions, the shape analogy constraints are not effective, resulting, locally, in a poor map. Post-processing using a map smoothing algorithm allows us to smoothly interpolate the map in such regions. To make a fair comparison, we also apply post-processing using RHM to all the methods that we compare to. Both [EBC17] and [ESBC19] are publicly available.

4. Analysis

In this section we evaluate our design choices, and investigate the properties of our approach. We show that the initial choice of base shapes can indeed be random, without affecting the algorithm's performance; discuss the considerations for choosing different shape difference operators; and show our algorithm's ability to handle collections of different size.

4.1. Resilience to Base Shape Choice

We show that our method is resilient to the choice of base shape, thus we can select it randomly, without requiring additional input.

In [SBC14] the result strongly depended on the choice of base shape in each collection, i.e. a bad base shape choice led to poor matching and poor functional maps. Our algorithm, on the other hand, is able to match corresponding pairs successfully, even when the choice of base shape is random. With a random choice of base shape in each collection, the chosen shapes (in most cases) do not represent a matching pair. Yet, the first part of our algorithm, aiming to match pairs of the two collections, is barely affected. The

Table 1: Resilience to Base Shape Choice, for corresponding and non-corresponding base shape pairs.

Result/Data-Set	corresponding base shape pair				non-corresponding base shape pair			
	Blend Shapes	Sumner	Rigged	FAUST	Blend Shapes	Sumner	Rigged	FAUST
Average matching accuracy	99.88%	100%	100%	87.9%	99.25%	100%	100%	85.5%
Average normalized alignment error	0.1542	0.1950	0.0566	0.3383	0.2012	0.3119	0.1481	0.3555
Rate of corresponding base shapes after alignment	40/40	10/10	10/10	86/100	39/40	10/10	10/10	85/100

second part of the algorithm, producing the functional map, is affected by the choice of base shape, however it benefits from the result of the first part by using the computed matched pairs.

Shape space alignment. To demonstrate the invariance of our approach to the choice of base shape, we test all possible base shape choices on a few datasets, and compare with the ground truth. We test both the case that the base shapes correspond, i.e., represent the same pose (Table 1, middle), and the case where the base shapes do not correspond (Table 1, right). For each possible choice of base shapes, we run the algorithm and measure a few performance indicators. *Matching accuracy* represents the percent of shapes that were matched correctly to the corresponding shape in the second collection, and it is averaged over all choices of base shape pair. The *normalized alignment error*, defined in Eq. (5), indicates how well the two point clouds are aligned after applying the rigid transformation. It is also averaged over all experiments. Finally, as discussed previously, after alignment we choose the pair with the smallest distance as the new base shape pair. The *rate of corresponding base shapes after alignment* indicates in how many cases the chosen base shape pair after alignment indeed represents a corresponding pair. For the FAUST data set we averaged the result over 10 subset pairs, each subset containing 10 shapes referring to the same person.

Note that for tested data-sets, the choice of base shapes, and whether or not the base shapes correspond, does not significantly affect the average matching accuracy nor the rate of corresponding base shapes chosen after alignment. The reason is that the principal components derived from PCA do not significantly depend on the base shape choice, causing only small variations in the shape space, that our alignment procedure can handle. Although the average normalized alignment error is larger for non-corresponding base shapes, the matching accuracy is hardly affected, meaning that we can still align successfully the two point clouds. These results also demonstrate that in most cases the new chosen base shape pair is a corresponding pair.

Inter-map computation. In the second part of the algorithm, it is more important to have corresponding base shapes. Figure 6 shows that the functional map computation has preferable results for corresponding base shapes, resulting in a less noisy functional map matrix, which is more similar to the ground truth. This motivates our choice to leverage the alignment results from the first part of the algorithm, and choose a new pair of base shapes which are more likely to match for computation of the functional inter-map.

4.2. Area-based vs. Conformal Shape Differences

The choice of using area-based or conformal shape difference operators depends on the nature of the data-set. For collections in which the variations between the shapes are area preserving but not conformal (e.g. the same person in different poses, Figure 7), it is beneficial to use the conformal shape difference operator, as it better captures the variations in such a collection. Similarly, for collections with conformal variations but not area preserving (e.g. different people in the same pose, Figure 8) we would rather use the area-based shape differences.

Figure 9 demonstrates how the choice of shape difference type affects the performance of the algorithm, in the case of data-sets in which each collection has non-conformal variations and in the case of non-area-preserving variations. We show the *matching accuracy* (top) and the *normalized alignment error* (bottom left), defined in Eq. (5), as a function of the dimension of the shape space. Note that for non-conformal variations, using the conformal shape difference operators yields better results, using a lower shape space dimension, and vice versa for the non-area-preserving case.

Figure 9 (bottom, right) shows the *cumulative energy*, defined as the sum of variances explained up to a certain dimension by using PCA. We demonstrate that for a collection with mostly area variations (same pose), the area based shape difference explains the data

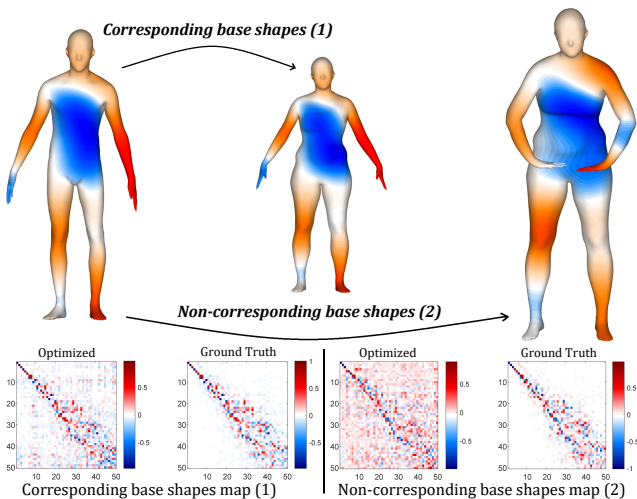


Figure 6: Functional maps and their visualization for corresponding and non-corresponding base shapes. Corresponding base shapes lead to a higher quality functional map (e.g. arms area).

using less principal components than the conformal shape difference. For the data-set with the non-conformal variations, the number of principal components required to explain the data is similar for conformal and area-based shape differences.

Note that even if the distances between the shape differences are similar in both collections, the low-dimensional embedding can be unstable, as the principal components are only defined up to sign. Furthermore, even for *identical* shape differences, the order of principal components that have a similar explained variance is arbitrary. However, as we seek a rigid transformation that matches the two sets, our method is robust to these instabilities. If, however, we take a considerably smaller number of dimensions, some principal components might appear in one collection and not the other, leading to the non-monotonic performance graphs in Figure 9.

4.3. Collections of different size

For collections of different size, some shapes in the bigger collection do not have a match in the smaller collection. Our algorithm naturally handles this case, since we find a permutation matrix that selects points from the larger collection's point cloud such that they align better with the rotated smaller point cloud (see Eq. (4)). We analyze how the excess shapes in the bigger collection affect our results. We split the answer into two cases.

4.3.1. Small collections

For small collections, e.g. containing five shapes, adding an excess shape can significantly affect the shape space of the resulting extended collection. This happens since the variations of the added shape are not always explained by a linear combination of the existing principal components. Hence, the newly computed principal components are significantly different, leading to very different shape space embeddings for the two collections, yielding

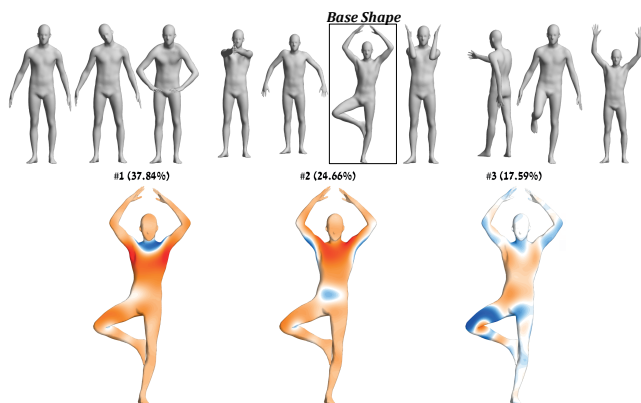


Figure 7: A collection with non-conformal variations (top) and their first 3 conformal shape differences' principal components (bottom) with their explained variance. Variations are located at joints, where motion occurs, mainly in the shoulders (first and second principal components) and the right knee (third principal component).

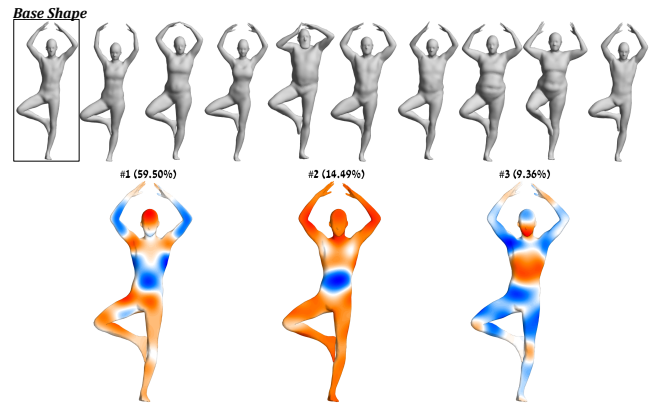


Figure 8: A collection with non-area-preserving variations (top) and their first 3 area based shape differences' principal components (bottom) with their explained variance. Variations occur at areas where people differ, mostly hips, torso and abdominal area. E.g., the second principal component corresponds to greater abdominal area at the expense of thicker arms (or vice versa).

poorer matching accuracy after aligning the two point clouds. Figure 10 (left) shows the matching accuracy for various small data-sets. Starting from collections of five shapes each, we added one shape at a time to one of the collections and tested the matching accuracy averaged over all possible choices of base shapes for the two collections. Finally, the bigger collection had twice as many shapes as the original. It can indeed be seen that the matching accuracy is significantly affected and decreases quickly. The notation *A, B* and *C* for the Blend Shapes data-set refers to three different subsets from this data-set. Note that each collection in the FAUST data-set contains up to ten shapes, thus it is treated as a small collection.

4.3.2. Large collections

In the case of large collections, containing twenty shapes for instance, it is likely that the variations of the added shapes can be described using a linear combination of the existing principal components. In this case, since we use linear dimensionality reduction, the shape space is not significantly affected, and adding shapes is equivalent to adding points to the point cloud without moving the original points. Figure 10 (right) demonstrates this, as the match-

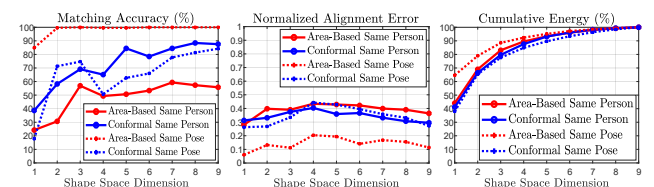


Figure 9: Comparison of our results for two collection types, with two shape difference operators. Note that for a collection with same-person-different-poses (Figure 7), the conformal shape differences show better performance. On the other hand, for a collection with same-pose-different-people (Figure 8) the area-based shape differences are better. See the text for details.

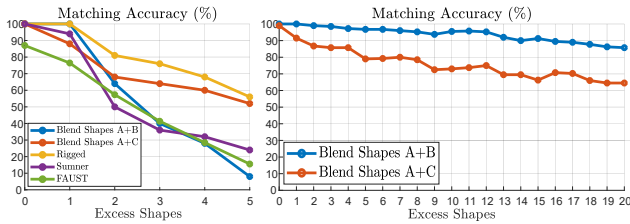


Figure 10: Matching accuracy for collections of different size. Note the considerable effect on the matching accuracy for small collections (left) in contrast to large collections (right).

ing accuracy is less affected in this case. Here we start with original collections of size twenty and add up to twenty more shapes to one of the collections resulting in a collection which has twice as many shapes. Despite this, the matching accuracy decreases only to 65 – 85%, depending on the data-set. Also note that due to the linear dimensionality reduction, we do not have a strict limit on the largest collection size. However, the point cloud alignment timing will be the performance bottle neck for very large data-sets.

5. Experimental Results

In this section we show some results of our algorithm, compare to the previous shape collection matching method, and to other automatic correspondence methods. We show that our method automatically computes high-quality maps, which compare favorably to existing automatic methods, given even two small collections or short animations, and even for less-than-perfect matching results.

5.1. Implementation details

Parameters. Our algorithm does not require data-dependent parameter tuning and we fixed the parameters to $k_1 = 50$ eigenfunctions on the base shape and $k_2 = 3k_1 = 150$ eigenfunctions on the other shape when computing the shape difference operators resulting in 50×50 matrices. The matching part has been done with conformal shape differences, unless stated otherwise. We set $\beta = 0.95$ to determine the shape space dimension and $\alpha = 0.1$ for the Laplacian commutativity in the energy we minimize (after normalizing both terms to have the same Frobenius norm). We run the ICP refinement of the functional map until the biggest change in the functional map matrix elements is smaller than 10^{-10} . We set the number of iterations in the RHM post-processing to 200, as recommended by the authors.

Timing. Our timings are comparable to other automatic methods for computing correspondences. Implemented with MATLAB, on a desktop machine with an Intel Core i7 @3.4GHz processor, the first part of the algorithm matches the pairs of the two collections of 10 shapes each, with shapes of $5K - 8K$ vertices in 45 seconds. The second part of the algorithm extracts the point-to-point inter-map in 130 seconds. The RHM post-processing takes 220 seconds on our machine. For most data-sets, it is possible to use smaller functional map matrices, e.g. $k_1 = 30$ and $k_2 = 3k_1 = 90$ without significantly decreasing the performance and obtaining both matches and an inter-map in approximately 75 seconds.

Limitations. We assume that the collections will lead to a similar point cloud structure of the shape space embedding, i.e., the variations within each collection should be similar. We also require non-symmetric shape space embeddings to allow unambiguous rigid alignment. For a very high dimensional shape space embedding (12 and greater) alignment using PM-SDP [MDK*16] can be slow.

5.2. Comparison: Matching Corresponding Shapes

In the previous Section, Table 1 showed our algorithm’s matching accuracy on four different data-sets, regardless of the base shape choice and demonstrated the robustness of our algorithm on several data-sets. For the FAUST data-set, it was enough to mistake only two pairs to decrease the matching accuracy to 80%, since every collection contains ten shapes.

Now we compare our algorithm to the only method known to us having the same goal of matching corresponding pairs from two collections. We test our algorithm and the method in [SBC14] on the Blend Shapes data-set with varying collection size. The results appear in Table 2 and show that in contrast to [SBC14], our method does *not* require large data-sets for successful matching and we get perfect matching for any collection size. Our results were averaged over 100 random initial base shapes choices since our algorithm selects a base shape randomly.

Table 2: Blend Shapes data-set: Matching accuracy vs [SBC14]

Method \ Collection Size	20	25	35	40
SBC14	35%	36%	63%	90%
Ours	99.47%	99.78%	99.54%	99.41%

Figure 11 shows the quality of our functional map compared to that of [SBC14], for the case of 40 shapes. Note that our map is better, even though the matching accuracy of [SBC14] is 90%. It is clear that using the regularization term of the Laplacian commutativity in Eq. (6) significantly improves the results, allowing us to obtain a functional map very similar to the ground truth.

5.3. Comparison: Point-to-Point Inter-Map Computation

Now we evaluate the computation of the point-to-point inter-map between the two base shapes, one from each collection. For the qualitative comparison, we show the texture on the target mesh and the ground truth, if it exists, of the source mesh. We compare our method to other automatic methods: BIM [KLF11] and BCICP [RPWO18]. For all methods, we show the results with and without post-processing using RHM for map refinement. For the quantitative evaluation we measure the map smoothness through its conformal and area distortions, and its semantic accuracy, using the distance to the ground truth correspondence, when given.

5.3.1. Quality Metrics

Conformal distortion. We use the definition given in [HG00, Eq. (3)] for the conformal distortion of a single triangle $f_{M_{BS}} \in \mathcal{F}_{M_{BS}}$: $\kappa(f) = \frac{\sigma_{M_{BS}}}{\sigma_{N_{BS}}} + \frac{\sigma_{N_{BS}}}{\sigma_{M_{BS}}}$ where $\sigma_{M_{BS}} \geq \sigma_{N_{BS}}$ are the singular values of

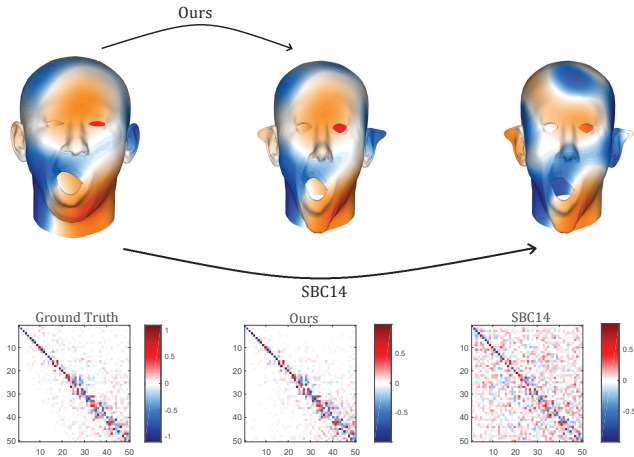


Figure 11: Comparing the functional maps obtained using our algorithm and [SBC14] using function transport (top), and comparison to the ground truth matrix (bottom). Our algorithm obtains a high-quality functional map, very similar to the ground truth.

the linear transformation which maps $f_{M_{BS}}$ from M_{BS} to N_{BS} . We subtract 2 such that the minimal conformal distortion is zero and visualize the result as a cumulative graph showing the percentage of triangles with less than a certain distortion value.

Area distortion. We measure the distortion in the area of each triangle caused by the map from M_{BS} to N_{BS} . For each triangle we compute $\left| \log \left(\frac{A(P_{N_{BS}, M_{BS}} f_{M_{BS}})}{A(f_{N_{BS}})} \right) \right|$ where $A(f_{N_{BS}})$ denotes the area of a triangle $f_{N_{BS}} \in \mathcal{F}_{N_{BS}}$ and $A(P_{N_{BS}, M_{BS}} f_{M_{BS}})$ the area of that triangle mapped by $P_{N_{BS}, M_{BS}}$, the extracted point-to-point inter-map. We visualize the result as a cumulative graph showing the percentage of triangles with less than a certain area distortion value.

Distance from ground truth. When a ground truth map is given, we measure the distance from the ground truth using the protocol suggested by [KLF11, Section 8.2]. For every mapped vertex, we measure its geodesic distance from the ground truth location, relative to the square root of the total area of N_{BS} , and visualize the percent of vertices whose distortion is less than a given value.

5.3.2. Setup

Choice of pairs. We always compute the map between the automatically computed base shapes after alignment. Since the initial choice of base shapes is random, the pair with the smallest distance after alignment may vary, resulting in a different pair of shapes (one from each collection) that we map between. Therefore, We average the results obtained by running the algorithm n times, where n equals the (smaller) collection size. For n experiments, we compute the correspondence for the computed base shape pair, where each base shape pair can be considered at most once. This procedure gives a set of $m \leq n$ pairs, on which we run the other methods and average the results. As shown in Table 1, in most cases we indeed get a corresponding base shape pair.

Parameters for comparison methods. For BIM, we use the default parameters (no need to set them manually). When applying BCICP we use the parameters recommended by the authors [RPWO18]: $k_1 = k_2 = 50$ (the number of eigenfunctions used for each shape), the weight for the orientation-preserving term was set to 0.1 and we used 10 iterations for the BCICP refinement step. For the FAUST data-set we set the time-scale parameter for computing the WKS descriptors to 100 and the skip size for the computed descriptors to 10. For all the other data-sets, we use the parameters suggested for the non-isometric TOSCA data-set, i.e. the time-scale parameter for the WKS descriptors was set to 50.

Mesh normalization. While our method does not require all the meshes to have the same surface area, other methods do, hence we normalized all the meshes to the same area.

5.3.3. Blend Shapes data-set

We use our method to compute the correspondences for the Blend Shapes data-set (using 10 shapes) (Figure 12). For this data-set, our algorithm did not benefit from RHM refinement while the other methods did. It is seen both qualitatively and quantitatively that our method achieves the best performance while the others fail to compute a high-quality map or resolve symmetry ambiguities (BCICP flips the upper part of the head while maintaining the correct orientation of the bottom part).

5.3.4. FAUST data-set

For the FAUST data-set, we use the ten subsets that represent ten

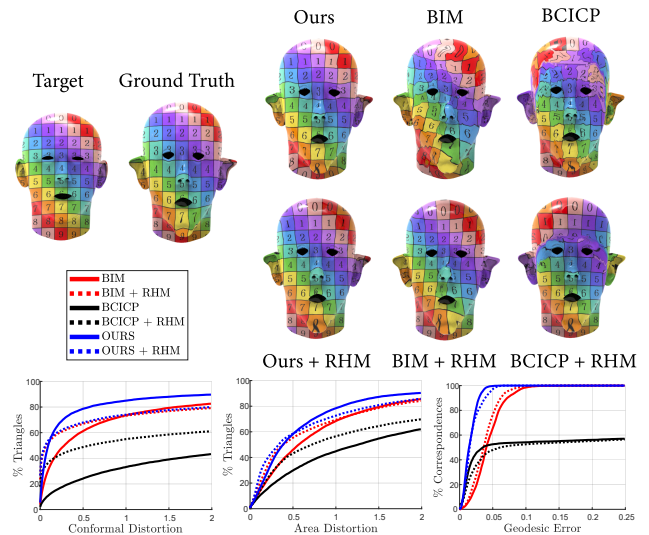
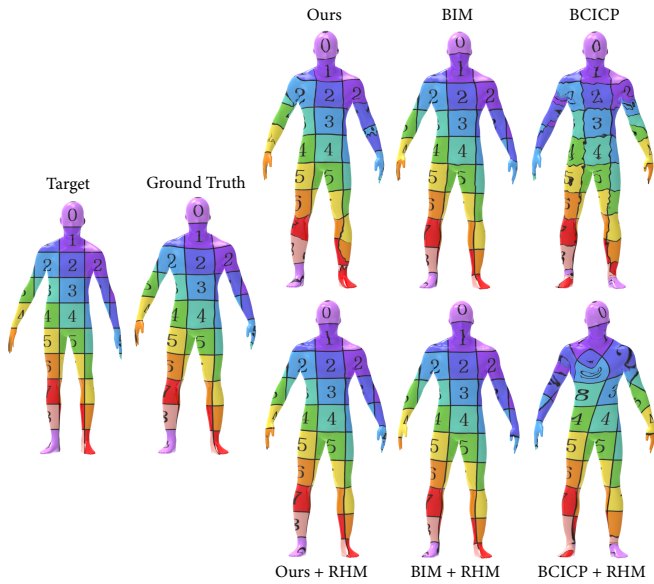
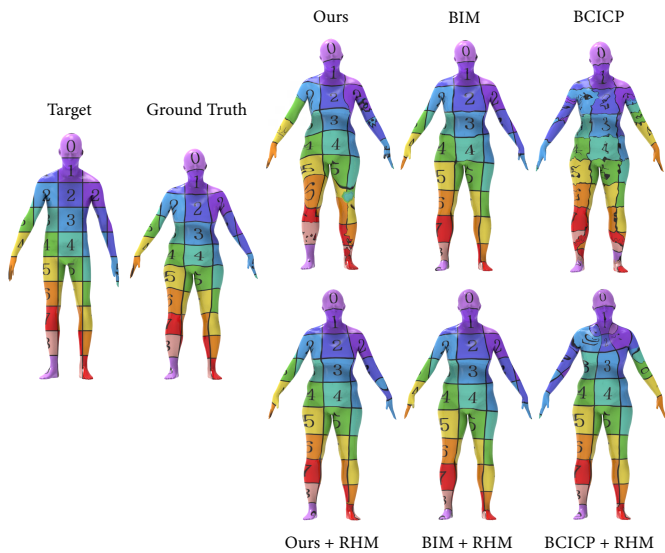


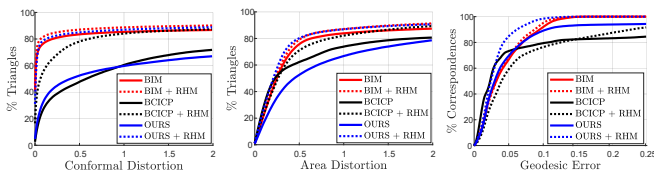
Figure 12: Blend Shapes, qualitative (top) and quantitative (bottom) comparison to other automatic methods BIM [KLF11] and BCICP [RPWO18], with and without final refinement using RHM. Our method yields high-quality maps, while others suffer from distortions or do not resolve symmetry ambiguities.



(a) Our algorithm refined with RHM obtains the highest quality map: Note that the feet and hands areas are closer to the ground truth than BIM. BCICP cannot completely resolve symmetry ambiguity, flipping the hands.



(b) Despite imperfections of our results on the hands and feet, our map refined with RHM leads to a smaller geodesic error.



(c) Quantitative results.

Figure 13: FAUST - qualitative and quantitative comparison: Our method and other automatic methods - BIM and BCICP, with and without final refinement using RHM.

different people as collections and use all the subset pairs to evaluate our algorithm (45 combinations in total). For the chosen base shapes in each subset pair we compute the resulting map using the comparison methods. When RHM is not used, BIM obtains preferable results. However, after the refinement our method achieves similar conformal and area distortion, yet the smallest distance from the ground truth (Figure 13 (c)). Note, that BIM is only applicable to genus zero shapes, whereas our method (including the RHM post-processing) is applicable to any mesh topology.

Qualitatively, Figure 13 (a) shows that our algorithm refined with RHM obtains the highest quality map. Note that the feet and hands areas are closer to the ground truth than BIM. On the other hand, BCICP cannot completely resolve the symmetry ambiguity, flipping the right arm with the left one, while maintaining the correct orientation for the other parts. In Figure 13 (b), despite imperfections of our method on the hands and feet, our map refined with RHM still obtains the smallest geodesic error. We conjecture that BCICP has difficulties resolving symmetry ambiguities. In our case, the base shape pairs computed by our algorithm are, in most cases, corresponding pairs, yielding more symmetric shapes for both the source and the target than in averaging over all random pairs in this data-set (most of the shapes in this data-set are extrinsically symmetric). Moreover, BCICP appears to be very sensitive to the parameters of the WKS descriptors, which need to be tuned for each data-set separately.

5.3.5. Summer data-set

The Summer data-set does not have ground truth correspondences, hence we show only conformal and area distortion as well as a qualitative visualization. For this data-set, after RHM refinement we get conformal distortion similar to that of BIM and a smaller area distortion. The qualitative results in Figure 14 show that we indeed get high-quality maps. Since every time we run the algorithm we might get two different shapes that we put in correspondence, we get a different choice of base shape pair than in Figure 1, demonstrating the results on another pair.

5.3.6. Rigged fruit data-set

We acquired this dataset on TurboSquid, and the ground truth map is not available. In this data-set all the shapes are symmetric, hence, the maps obtained cannot distinguish between left and right. However, note that BIM also flips the map upside-down, which is fixed using RHM refinement. In this data-set, when no RHM refinement is applied, our method yields the best map since it is not as noisy as BCICP and does not flip up and down as BIM. After RHM refinement good results are obtained for all methods (see figure 15).

5.4. Inter-Map Computation using Composition of Maps

We now demonstrate that it is possible to obtain the inter-map from any shape in collection A to any shape in collection B using the map composition in Eq. 7, exploiting the inter-map for the corresponding base shapes and the given intra-maps. Figure 16 shows

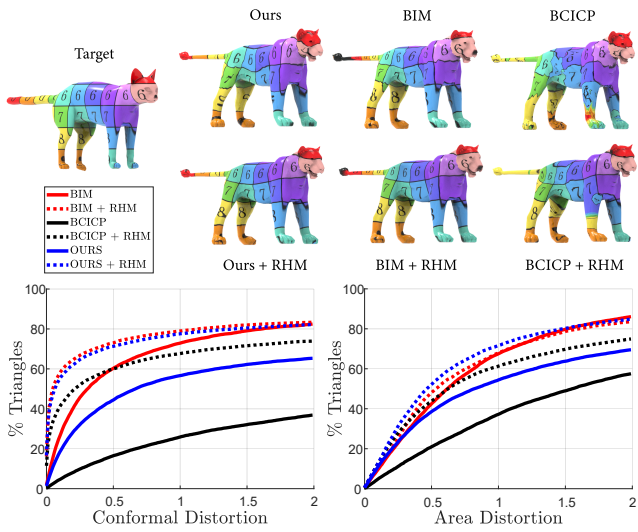


Figure 14: Summer, qualitative and quantitative comparison with automatic methods BIM and BCICP, with and without RHM refinement. Note that we achieve conformal distortion similar to BIM, yet a better area distortion.

the case of an inter-map for *non*-corresponding shapes. First, we show the case of an inter-map computed *directly* for the two shapes (as if they were chosen as base shapes). As explained in section 4.1, this case of non-corresponding base shapes yields poorer maps as indeed can be seen. Second, we visualize the map obtained using Eq. 7 where the base shapes are corresponding (as in Figure 13). In this case, the computed maps are of high-quality. We also compare to the results obtained using BIM and BCICP. In all cases maps are visualized after RHM post-processing.

5.5. Inter-Map Computation with Low Matching Accuracy

After we showed that the cross collection map computation using our method leads to good results both qualitatively and quantitatively, we test the resilience of this part of the algorithm when the matching accuracy, determined in the first part of the algorithm, is low. Using the ground truth of the matches, we set them manually to obtain the desired accuracy and test the results of the map computation. The matching accuracy affects the analogies constraints (Eq. (6)), since we use also wrong matches as constraints. We show that our approach can handle wrong matching in the first phase of the algorithm and still retrieve high-quality maps.

In Figure 17 we demonstrate the experiment for FAUST dataset, all the results and visualizations include RHM refinement. Note that starting from 80% accuracy it is possible to obtain high-quality maps. In other words, the algorithm can tolerate the wrong matching of two pairs, when we use collections with ten shapes.

Figure 18 shows the same experiment for the Blend Shapes dataset (using 40 shapes). Here we visualize the maps after RHM refinement, all achieving similar conformal and area distortion and high-quality maps, Figure 18 (top). We conclude that in spite of the low matching accuracy, RHM refinement is able to recover a map very

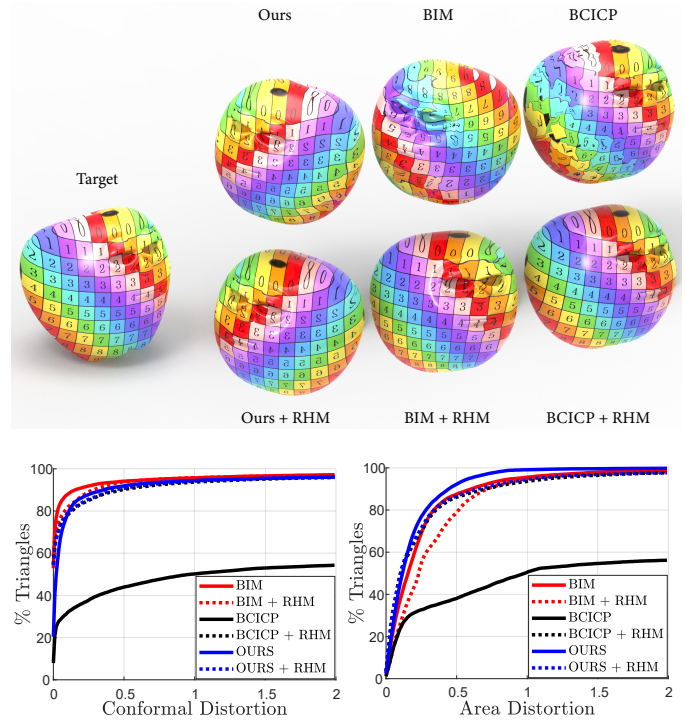


Figure 15: Rigged fruit, qualitative and quantitative comparison with BIM and BCICP, with and without RHM refinement.

similar to the ground truth. Hence, to show the differences between different matching accuracies, we plotted the conformal and area distortion for the results *without* RHM refinement, Figure 18 (bot-

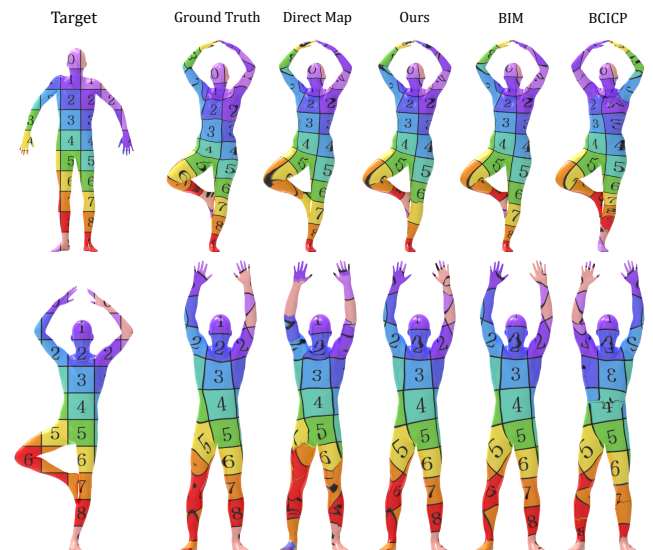


Figure 16: Map composition leads to a higher quality inter-map for non-corresponding shapes compared with computing the map directly. All maps include RHM refinement.

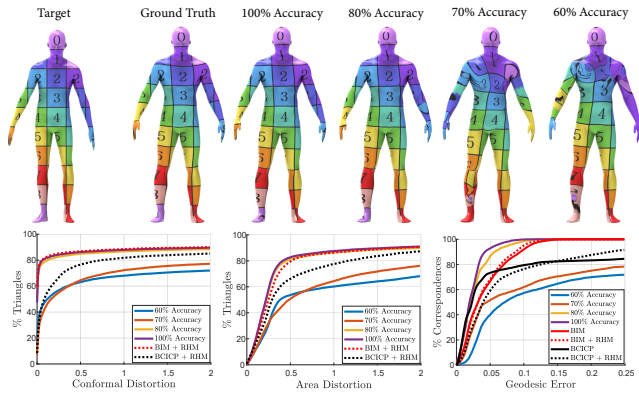


Figure 17: FAUST, qualitative (top) and quantitative (bottom) comparison for varying matching accuracy. We manually set the matches to allow varying accuracy and evaluate the computed map. Note that starting from 80% we get high-quality maps. All results include RHM as post-processing.

tom). For this data-set high-quality maps are obtained even when the matching accuracy is as low as 40%. Since this data-set has more shapes than FAUST (40 vs 10), we have more terms in the energy function (Eq. (6)) and since the variations within the shapes are localized to the face region, we recover the functional map even with a lower matching accuracy.

5.6. Inter-Map Computation for Varying Collection Size

We evaluate the computed map depending on the collection size, especially for small collections, i.e. with collection size as low as 2 or 3. We demonstrate that computing correspondences using our method is possible for such small collections, allowing us to retrieve high-quality maps even when only few shapes are given. In Figure 19 (left) we see that for the Blend Shapes data-set it is possible to get a high-quality map using any collection size varying from 2 to 40. For the case with 2 shapes, we set the matches manually since aligning point clouds with two points each has 50% success rate as they are necessarily symmetric. Results are shown *without*

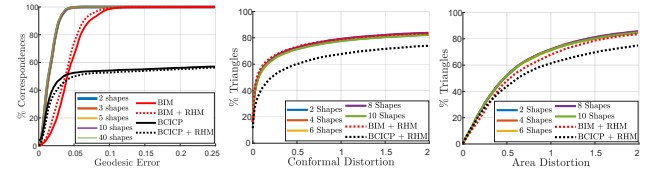


Figure 19: Varying collection size, Blend Shapes (left, without RHM) and Sumner (center and right, with RHM) datasets. We obtain a high-quality map even for collections with only 2 shapes.

RHM refinement. This result means that we can get the cross collection map when we have only three shapes in each collection, or two with manual matching, and still get higher-quality map than other methods, with very low timing, since for smaller collection size the time needed for our algorithm notably decreases. It is important to mention that matching accuracy remained as high as 100% in all cases of collection sizes, even for 3 shapes only (except the manual matching for the case with 2 shapes). Figure 19 (center, right) shows the same experiment for Sumner data-set. Note that the conformal and area distortion of the map are not affected when we decrease the number of shapes, neither is the matching accuracy. Our results include post-processing using RHM. As in the demonstration for the Blend Shapes data-set, we set the matches manually for the case with 2 shapes in a collection, to handle the case of symmetric shape space.

5.7. Regularization Effect

We analyze the quality of the functional map for varying α (the weight of the regularization) and varying collection size on the Blend Shapes data-set. We measure the error of the optimized functional map $C_{N_{BS}, M_{BS}}$ (after ICP refinement), and plot the error as a function of α (Figure 20). The error is the distance from the ground truth map C_{GT} , namely $\|C_{N_{BS}, M_{BS}} - C_{GT}\|_F$. Analysing the results, we see that without regularization ($\alpha = 0$), the error is negatively correlated with the collection size: more analogies yield better maps. If, however, the amount of data is limited to a small number of analogies, there exists a range of α ($0.1 \leq \alpha \leq 0.15$), for which we get the same best-optimized functional map with the same error for any collection size. Hence, a good regularization decreases the dependency on the collection size. Finally, for any collection size, as α increases the error converges to the error value with no analogies term (the black horizontal line). Hence, at least two analogies are required, and regularization alone ($\alpha \rightarrow \infty$), cannot yield a good functional map.

6. Conclusion

In this work, we presented a robust and simple to implement method for matching two shape collections with high rates of accuracy. Unlike previous approaches, our method can perfectly handle small and large collections alike, while for the latter we maintain high accuracy rates even when one of the collections contains excess shapes that do not need to be matched, thus allowing to handle various noisy sampling of the shape space.

Furthermore, a high-quality inter-map is obtained using only

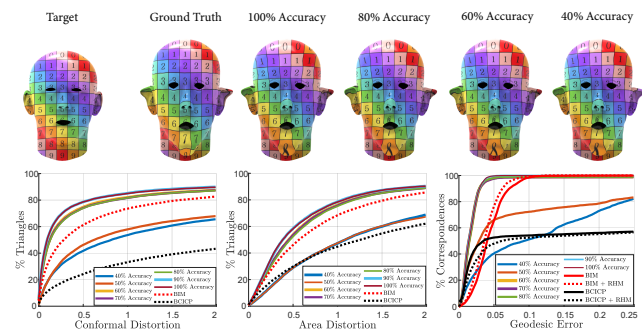


Figure 18: Blend shapes, qualitative (with RHM, top) and quantitative (without RHM, bottom) comparison for varying matching accuracy. We manually set the matches to allow varying accuracy and evaluate the computed map. We obtain a high-quality refined map with matching accuracy as low as 40%.

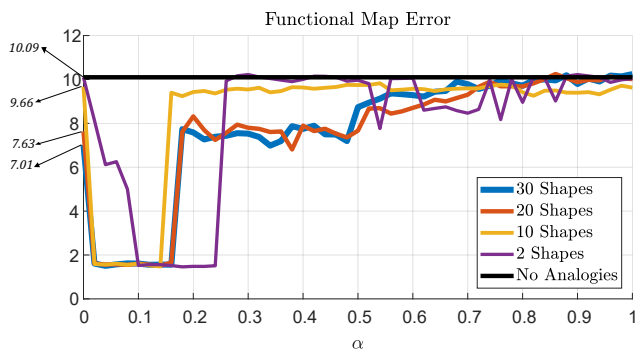


Figure 20: Functional map error for varying regularizer weight and several collection sizes. See the text for details.

the analogies as the semantic information, allowing our algorithm to serve as a fully-automatic method for computing correspondences, surpassing other state-of-the-art automatic methods for non-isometric shape correspondence. As an additional proof of its robustness, we showed that even if not every shape has a match or matching was imperfect, the second step of the algorithm, computing the inter-map, is resilient to it. Finally, the variation within the collection can be large, and both quasi-isometric or non-isometric making our approach applicable to various data-sets.

We believe that our approach can serve as an important new tool in the shape analysis and correspondence toolbox. Future work and generalizations include using other inner product metrics to tailor the shape differences for specific applications [CO19], and using better regularization constraints [RPWO19]. Moreover, it might be beneficial to use *multiple* base shapes, and to learn the shape differences operators from data. Finally, since the functional map approach is agnostic to the geometry representation, it can be interesting to apply our approach to settings where the data has an intrinsic parametric structure allowing to generate many corresponding instances, e.g. parametric CAD models.

Acknowledgments

M. Ben-Chen acknowledges funding from the European Research Council (ERC starting grant No. 714776 OPREP), and the Israel Science Foundation (grant No. 504/16).

References

- [Bha09] BHATIA R.: *Positive definite matrices*, vol. 16. Princeton university press, 2009. 5
- [CO19] CORMAN E., OVSJANIKOV M.: Functional characterization of deformation fields. *ACM Transactions on Graphics (TOG)* 38, 1 (2019), 8. 14
- [CRA*17] COSMO L., RODOLA E., ALBARELLI A., MÉMOLI F., CREMERS D.: Consistent partial matching of shape collections via sparse modeling. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 209–221. 2
- [DML17] DYM N., MARON H., LIPMAN Y.: Ds++: A flexible, scalable and provably tight relaxation for matching problems. *ACM Transactions on Graphics (TOG)* 36, 6 (2017). 3
- [EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 165–174. 6
- [ESBC19] EZUZ D., SOLOMON J., BEN-CHEN M.: Reversible harmonic maps between discrete surfaces. *ACM Transactions on Graphics* (2019). 6
- [GSDG18] GANAPATHI-SUBRAMANIAN V., DIAMANTI O., GUIBAS L. J.: Modular latent spaces for shape correspondences. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 199–210. 2
- [GYQ*18] GAO L., YANG J., QIAO Y.-L., LAI Y.-K., ROSIN P. L., XU W., XIA S.: Automatic unpaired shape deformation transfer. In *SIGGRAPH Asia 2018 Technical Papers* (2018), ACM, p. 237. 2
- [HAGO19] HUANG R., ACHLIOPTAS P., GUIBAS L., OVSJANIKOV M.: Limit shapes—a tool for understanding shape differences and variability in 3d model collections. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 187–202. 2
- [HCO18] HUANG R., CHAZAL F., OVSJANIKOV M.: On the stability of functional maps and shape difference operators. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 145–158. 2
- [HG00] HORMANN K., GREINER G.: *MIPS: An efficient global parametrization method*. Tech. rep., ERLANGEN-NUERNBERG UNIV (GERMANY) COMPUTER GRAPHICS GROUP, 2000. 9
- [HRA*19] HUANG R., RAKOTOSAONA M.-J., ACHLIOPTAS P., GUIBAS L., OVSJANIKOV M.: Operatomet: Recovering 3d shapes from difference operators. *arXiv preprint arXiv:1904.10754* (2019). 2
- [HWAG09] HUANG Q.-X., WICKE M., ADAMS B., GUIBAS L.: Shape decomposition using modal analysis. In *Computer Graphics Forum* (2009), vol. 28, Wiley Online Library, pp. 407–416. 2
- [KKBL15] KEZURER I., KOVALSKY S. Z., BASRI R., LIPMAN Y.: Tight relaxation of quadratic matching. In *Computer Graphics Forum* (2015), vol. 34, pp. 115–128. 3
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended Intrinsic Maps. *ACM Transactions on Graphics (TOG)* 30 (2011). 2, 9, 10
- [MDK*16] MARON H., DYM N., KEZURER I., KOVALSKY S., LIPMAN Y.: Point registration via efficient convex relaxation. *ACM Transactions on Graphics (TOG)* 35, 4 (2016). 3, 5, 9
- [Mea92] MEAD A.: Review of the development of multidimensional scaling methods. *Journal of the Royal Statistical Society: Series D (The Statistician)* 41, 1 (1992), 27–39. 5
- [NBCW*11] NGUYEN A., BEN-CHEN M., WELNICKA K., YE Y., GUIBAS L.: An optimization approach to improving collections of shape maps. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 1481–1491. 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). 3, 6
- [ROA*13] 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 (2013), 72. 2, 3, 4
- [RPWO18] REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (TOG)* 37, 6 (2018). 2, 9, 10
- [RPWO19] REN J., PANINE M., WONKA P., OVSJANIKOV M.: Structured regularization of functional map computations. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 39–53. 14
- [SBC14] SHAPIRA N., BEN-CHEN M.: Cross-collection map inference by intrinsic alignment of shape spaces. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 281–290. 2, 3, 5, 6, 9, 10
- [SLHH18] SUN Y., LIANG Z., HUANG X., HUANG Q.: Joint map and symmetry synchronization. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2018), pp. 251–264. 2