

Groupwise Shape Correspondence Refinement with a Region of Interest Focus

Pierre Galmiche¹ and Hyewon Seo¹

¹ ICube Laboratory, CNRS–University of Strasbourg

Abstract

While collections of scan shapes are becoming more prevalent in many real-world applications, finding accurate and dense correspondences across multiple shapes remains a challenging task. In this work, we introduce a new approach for refining non-rigid correspondences among a collection of 3D shapes undergoing non-rigid deformation. Our approach incorporates a Region Of Interest (ROI) into the refinement process, which is specified by the user on one shape within the collection. Based on the functional map framework and more specifically on the notion of cycle-consistency, our formulation improves the overall matching consistency while prioritizing that of the region of interest. Specifically, the initial pairwise correspondences are refined by first defining the localized harmonics that are confined within the transferred ROI on each shape, and subsequently applying the CCLB (Canonical Consistent Latent Basis) framework both on the global and the localized harmonics. This leads to an enhanced matching accuracy for both the ROIs and the overall shapes across the collection. We evaluate our method on various synthetic and real scan datasets, in comparison with the state-of-the-art techniques.

CCS Concepts

• **Computing methodologies** → *Shape analysis*; • **Theory of computation** → *Computational geometry*;

1. Introduction

With the recent evolution in the shape capture technology, the use of captured geometry and the appearance of real-world shapes has emerged as important modelling tool in computer graphics, virtual reality, and medicine. A common prerequisite to such reconstructive shape modelling is the need for the full vertex correspondence among all captured shapes, for which a vast amount of solutions exist. Different strategies have been developed in order to circumvent the exponential complexity of the solution space and rapidly converge to the final or a number of acceptable solutions. Methods adopting the paradigm of combinatorial optimization on global and structural shape descriptors [ATCO*10, SSGD03, BFF*07], optimal deformation of template models [BRLB14], voting or clustering schemes to achieve collective consistency [CZ08, ATCO*10], RANSAC (Ransom sample consensus) based algorithms to *hypothesize-and-check* [TBW*09, TBW*11], use of statistical atlas to benefit from the known commonality and variability of target shapes [ACP03, ASK*05], and embedding based approaches [EK03, KLF11, KLCF10] where the shapes are mapped to a canonical domain before computing the correspondences in the new representation.

One of the successful embedding-based approaches has been established by the functional map framework. Originally presented in Ovsjanikov et al [OBCS*12] and considerably extended in sev-

eral follow-up works [RMOW20, RCB*15, AGB*21, HAGO19], it represents correspondences between shapes as a linear transformation or *mapping* between their respective function spaces. This allows most constraints (descriptor preservation, operator commutativity, or landmark correspondences) to be linear, making the map inference very efficient. Although an optimized functional map facilitates the transfer of real-valued functions between shapes, the conversion between functional and point-wise maps poses a challenging problem that is prone to errors, for which several solutions have been proposed, such as the map recovery as a probability density estimation [RMC17], and the extension of functional basis by using functional products [NMR*18]. Another inherent problem is related to the choice of the function basis and its size, which often relies on heuristics to achieve a favorable trade-off between the number of basis function and the precision of the mapping. To address this, iterative refinement methods such as ZoomOut [MRR*19, RMOW20] or EFMR (Effective Functional Map Refinement) [RMWO21] have been developed. These methods progressively up-sample the functional basis, by optimizing the point-wise and the functional map alternatively. In this work, we deploy the notion of localized harmonics [MRCB17] confined in a region of interest (ROI) to maintain effective number of functional basis while preserving high-frequency accuracy.

Recently developed functional map approaches can also deal

with partiality [RCB*15, APO21] or refine correspondences in a shape collection [HAGO19, HRWO20]. In particular, cycle consistency frameworks [HWG14, WHG13, GLT*20, MRC*21] propose to refine maps in a collective manner, by ensuring that the composite maps, starting and ending on a same shape leads to the identity mapping. While computationally efficient, such a groupwise refinement requires good initial pairwise correspondences forming cycles, thereby limiting its application to noise-free and full shape collections. We note that shape collections with different partial coverage except for an optional reference shape are prevalent in many real-world applications. In medical imaging, for instance, a reference shape is obtained during the preoperative phase, and partial data during subsequent operations or treatments.

Briefly, this work introduces a new method for refining the groupwise nonrigid correspondence with ROI focus. To the best of our knowledge, our work is the first effort to incorporate local consistency into the groupwise refinement. Our specific contributions are as follows:

- We enhance the cycle consistency framework [HAGO19] by incorporating local frequencies inherent to regions of interest (ROI) into the refinement process. This enables us to ensure matching consistency on both global and local scales.
- While establishing the necessary cyclic maps within the cycle-consistency framework, we initialize maps in both directions, even when dealing with real-world shapes with different coverage. To ensure the presence of bidirectional maps in the case of partial correspondences, we adapt the 'Bijective ZoomOut' [RMOW20] refinement technique to handle partial cases.

The rest of this paper is organized as follows. We start by reviewing preliminaries of the functional map framework for non-rigid shape correspondence computation (Section 2). In (Section 3) we present our refinement method, which is equipped with Bidirectional Partial ZoomOut and incorporates basis extensions with a ROI focus. Next, we provide details on the dataset and the implementation, in Section 4 and Section 5. After illustrating and discussing the results we obtained on several real-world and synthetic shape collections (Section 6), we conclude the paper (Section 7).

2. Preliminaries

Shapes and functional representation. Assume that we are given a collection of n 3D shapes $\mathcal{S} = \{S_i\}_{i=0}^{n-1}$, represented as triangular meshes. To each shape S_i , we associate a positive semi-definite Laplacian matrix $L_i = A_i^{-1}W_i$ by using the standard cotangent weight scheme [PP93], where W_i is the cotangent weight matrix and A_i is the diagonal lumped area matrix. The basis functions $\{\phi_{S_i}^k\}_{k \geq 1}$ composed of the eigenfunctions of L_i form an orthonormal basis of the square integrable functions on S_i , denoted as $\mathcal{L}^2(S_i)$, allowing to represent any function $f \in \mathcal{L}^2(S_i)$ using a Fourier series: $f = \sum_{k \geq 1} \langle f, \phi_{S_i}^k \rangle \phi_{S_i}^k$.

Functional correspondences. The functional map framework [OBCS*12] aims at finding correspondences between shapes S_i and S_j as a linear transformation $\tau : \mathcal{L}^2(S_i) \rightarrow \mathcal{L}^2(S_j)$ between their respective function spaces. As a linear operator, τ admits a matrix representation $C = (c_{kl})$ with coefficient determined as follows. The action of τ on a function $f \in \mathcal{L}^2(S_i)$ can be expressed as :

$$\begin{aligned} \tau f &= \tau \sum_{k \geq 1} \langle f, \phi_{S_i}^k \rangle \phi_{S_i}^k = \sum_{k \geq 1} \langle f, \phi_{S_i}^k \rangle \tau \phi_{S_i}^k \\ &= \sum_{k, l \geq 1} \underbrace{\langle \tau \phi_{S_i}^k, \phi_{S_j}^l \rangle}_{c_{kl}} \phi_{S_j}^l. \end{aligned} \quad (1)$$

Truncating the Fourier series to the first K terms allows to obtain a particularly compact representation of the functional correspondences with a $K \times K$ matrix C , where K is chosen to be small (between 20 and 100 in practice).

A general approach to optimize such a matrix assumes the availability of corresponding functions $g_c \approx \tau f_c$, $c = 1, \dots, Q$ and asks the functional map C to conserve the spectral coefficients in a least square sense:

$$\min_C \|CA - B\|_F^2 + \alpha E_{reg}(C), \quad (2)$$

where $\mathbf{A} = (\langle \phi_{S_i}^k, f_c \rangle)$ and $\mathbf{B} = (\langle \phi_{S_j}^l, g_c \rangle)$ are $K \times Q$ matrices of Fourier coefficients of the corresponding functions, α is a scalar weight and $E_{reg}(C)$ is a regularization term.

Recovering a point-wise map from a given functional map can be expressed as the solution of the following optimization problem:

$$T_{ij}(p) = \arg \min_{q \in S_j} \|C_{ji} \Phi_{S_j}^k(q)^T - \Phi_{S_i}^k(p)^T\|_2, \quad \forall p \in S_i, \quad (3)$$

which can be efficiently solved by searching for each row of Φ_{S_i} its nearest-neighbor in the space of rows of Φ_{S_j} , transformed by C_{ji} . Indeed, any point $p \in S_j$ can be represented by the corresponding delta function $\delta_p \in \mathcal{L}^2(S_j)$, expressed in Φ_{S_j} as

$$\Phi_{S_j}(p) = [\phi_{S_j}^1(p), \phi_{S_j}^2(p), \dots, \phi_{S_j}^K(p)].$$

Conversely, a functional map C_{ji} can be obtained from a point-to-point map $T_{ij} : S_i \rightarrow S_j$ using the following equation:

$$C_{ji} = \Phi_{S_i}^\dagger \Pi_{ij} \Phi_{S_j}, \quad (4)$$

where Π_{ij} is the matrix representation of T_{ij} such that $\Pi_{ij}(index(p), index(q)) = 1$ if $T_{ij}(p) = q$ and 0 otherwise, with $p \in S_i$ and $q \in S_j$.

Partial correspondence. Finding partial correspondences between two shapes, where only a subset of points or regions have a corresponding counterpart due to missing parts in either shape, is known as a highly challenging problem. present in many real data applications due to occlusion or partial views. In the functional map framework, partial-to-full correspondences have been represented using maps with a slanted diagonal structure. As demonstrated in [RCB*15], the slope of the functional map is proportional to the ratio of surface area of the two shapes $\frac{A(S_j)}{A(S_i)}$, $S_j \subset S_i$. Based on this property, they proposed a partial shape matching approach that regularizes the functional maps by enforcing them to have a slanted-diagonal structure. This work has been extended to cluttered non-rigid correspondences [CRM*16] and to multiple shape settings [LRB*20]. While showing encouraging results in challenging settings, the explicit representation of the part entails complicated alternating optimizations over the part indicator

functions in the spatial domain and the functional map in the spectral domain. To avoid this complex optimization scheme, Litany et al. [LRBB17] proposed a fully spectral partial shape matching (FSPM) approach where they absorbed the spatial mask into a new basis for the part function space with following optimization problem:

$$\min_{Q \in S(K,r)} \text{off}(Q^T \Lambda_S Q) + \mu \|A_r - Q^T B\|_{2,1}. \quad (5)$$

Here $S(K, r)$ is the Stiefel manifold of orthogonal $k \times r$ matrices, $\text{off}(A) = \sum_{i \neq j} a_{ij}^2$ and $A_r = W_r A$ with $W_r = (I_{r \times r} \ 0_{r \times K-r})$ denotes the $r \times K$ matrix containing the first rows of A . r is the rank of the partial functional map. This optimization computes a new basis, localized to a latent part of the full shape (data term) and that minimizes Dirichlet energy (smooth term) in order to obtain bases that behave as laplacian eigenfunctions. Such methods rely on hand-crafted features such as SHOT [TSDS10], which are sensitive to the mesh discretization and can lead to failure when shapes have significantly different discretizations.

Group-wise refinement. In [HWG14], Huang et al. introduce the Functional Map Networks (FMN) to compute consistent functional maps within heterogeneous shape collection. A FMN with n shapes is represented as a connected and directed graph $\mathcal{G}(\mathcal{F}, \mathcal{E})$, where $F_i \in \mathcal{F}$ represents the functional vector spaces of shape S_i , and an edge $(i, j) \in \mathcal{E}$ represents a functional map C_{ij} . Given a FMN \mathcal{G} , Wang et al. [WHG13] proposed to extract a set of consistent latent basis (CLB) $\{Y_i\}_{i=0}^{n-1}$ such that $C_{ij} Y_i \approx Y_j$, $\forall i, j$ through the following optimization:

$$Y = \arg \min_Y \sum_{(i,j) \in \mathcal{E}} \omega_{ij} \|C_{ij} Y_i - Y_j\|_F^2, \quad (6)$$

$$s.t. Y^T Y = nI^k, Y = [Y_0; Y_1; \dots; Y_{n-1}],$$

where ω_{ij} is a scalar reflecting the relative confidence assigned to each functional map of the network. Indeed, since initial maps usually vary in quality, one would prefer to assign larger weights to good maps and lower ones to noisy maps. This way, the optimized latent basis are less perturbed by potential noisy maps in the shape collection.

The latent basis Y_i can be seen as a functional map from some latent shape (playing the role of the average shape) to each S_i . Huang et al. [HAGO19] extended the latent bases by introducing an extra normalization to the CLB construction in order to incorporate a metric information to the geometric information contained in the CLB. The latent bases thus obtained are named Canonical Consistent Latent Basis or CCLB. The use of such consistent latent basis has several advantages. First, it directly represents the ‘average shape’ of the collection in the functional domain, avoiding the explicit construction of a template or average shape, which can be computationally expensive. Second, the latent spaces can be used to factor a functional map between two shapes S_i and S_j by using the latent bases composition $C_{ij} = Y_j Y_i^{-1}$. Finally, one can efficiently add a previously unseen shape S_k to the latent representation by composing the functional map C_{ik} to the latent basis Y_i of the shape S_i . The corresponding formula $Y_k = C_{ik} Y_i$, improves the scalability of the latent basis construction.

3. Our Approach

Given a collection of n 3D shapes denoted by $\mathcal{S} = \{S_i\}_{i=0}^{n-1}$ and an initial correspondence among them, our goal is to refine the initial non-rigid correspondences among them with a focus on a specific region of interest (ROI). The ROI is defined once by the user by manually selecting faces on one shape that we call as *reference shape*. Without loss of generality, let S_0 be the reference shape and $S_i \in \mathcal{S}$ ($i = 1, \dots, n-1$) represent the other shape instances. In general, each shape instance $S_i \in \mathcal{S}$ ($i = 1, \dots, n-1$) has a different coverage of a shape undergoing some elastic deformation and contains topological noise. Fig. 1 illustrates an example of such a shape collection. The ROI defined on S_0 is colored in cyan, and its border in yellow for the sake of visibility.




Figure 1: Front (a) and back view (b) of a shape collection $\mathcal{S} = \{S_i\}_{i=0}^{n-1}$. Left: Reference shape S_0 and its ROI in cyan. Right: Deformed instance shapes S_i with different coverages.

To improve the quality of correspondences among multiple shapes, we adapt cycle-consistency scheme by Huang et al. [HAGO19], requiring initial functional maps forming cycles. In the rest of this paper, we use the FMN represented in Fig. 2 to perform our correspondence refinement, composed of maps $\{C_{0i}\}$ between




Figure 2: Correspondence graph represents the correspondences that we use and refine in our groupwise refinement.

the reference shape and each of the instance shapes, as well as $\{C_{ij}\}$ between consecutive instance shapes. Unlike in Huang et al., however, our shapes have different coverage, necessitating the computation of partial correspondences. Partial correspondences usually being computed in one direction (from the partial shape to the full shape), we need to compute the missing correspondences (from the full shape to the partial one) to perform the groupwise refinement. To this end, we first invert the partial-to-full pointwise maps to obtain full-to-part one. We then adapt the ‘Bijjective ZoomOut’ [RMOW20] refinement to the partial setting so as to refine the both maps in a bidirectional manner.

By using this approach, we can obtain initial correspondences that form cycles, which serve as the basis for our groupwise refinement method. We extend ‘canonical consistent latent basis (CCLB)’ scheme by Huang et al [HAGO19], by incorporating an ROI focus. The ROI is a user-specified area on a shape in the collection, emphasizing the need for precise correspondences on that region. Instead of using a sufficiently large number of functional basis, we extend the basis by the use of localized harmonics (LMH) defined on the ROI. This approach allows to preserve high-frequency accuracy while maintaining a limited number of functional basis. It is worth noting that the number of manifold harmonics plays an important role in the functional map framework. As noted in Melzi et al, [MRR*19], spectral methods can be both computationally expensive and unstable when using an excessive number of harmonics, whereas reducing the number of harmonics leads to imprecise maps that lack high-frequency details. Although the concept of LMH has been previously introduced [MRCB17] in the literature, we demonstrate how the notion of LHM can be integrated in the groupwise refinement process. Details of our method are described in what follows.

3.1. Initialization of pairwise correspondences

We first compute pointwise initial correspondences T_{i0} from each of the instance shapes S_i to the reference shape S_0 , and $T_{i+1,i}$ between consecutive instance shape pairs, i.e. from S_{i+1} to S_i . Two different approaches are adopted depending on the presence of partiality in the collection:


(1) Collections of full shapes: For the synthetic dataset without partiality, we utilize the classical functional map approach (Equation 3) with three regularization terms: the Laplacian commutativity term, the orientation preserving term, and the multiplicative operator term.

(2) Collections with partial shapes: We use the ‘Fully Spectral Partial Matching’ (FSPM) [LRBB17] to compute partial correspondences between the shape. We use the part-to-full FSPM because using the part-to-part FSPM produced worse results in some cases while being computationally more expensive. Then, to generate the correspondences of opposite direction, we invert the previously computed pointwise correspondences T_{i0} and $T_{i+1,i}$. Specifically, $T_{ji}^{-1} \approx T_{ij} : S_i \rightarrow S_j$ is defined as follows:


$$\forall q \in S_i, T_{ji}^{-1}(q) = \begin{cases} p & \text{if } T_{ji}(p) = q \\ p_0 & \text{otherwise} \end{cases},$$

where $p_0 \in S_j$ is randomly chosen.

Note that inverting part-to-full correspondences leads to full-to-part correspondences, a problem that has not been addressed before. Although it may seem unnatural to compute it, it constitutes an essential component to obtain maps forming cycles, which is required by the CCLB framework our work is based on.



(a) ROI transfers by using the reference-to-instance pointwise correspondences. The reference shape S_0 with the user-specified ROI is shown on the left. On the right, the instance shapes (S_1, \dots, S_8) are arranged from left to right, and top to bottom.



(b) Colormap transfers between consecutive instance shapes by using inverse pointwise correspondences.

Figure 3: ROI pointwise transfers (a) and inverse pointwise correspondences (b) obtained by using the FSPM [LRBB17] on the woman’s torso data.

3.2. Pairwise correspondence refinement

The initial pairwise correspondences obtained from Section 3.1 often need further refinement, especially in case of partial shape collections. As shown in Fig. 3 the result of the initial correspondences is sparse and noisy. The ROI transferred through the initial maps T_{i0} (Fig. 3a) show noisy borders, and the inverse pointwise maps $T_{i,i+1}$ (Fig. 3b) reveals the sparsity. To refine the initial correspondences, we once again adopt two different approaches according to the presence of partiality in the collection. For the full-to-full cases, the classical iterative up-sampling method called ‘ZoomOut’ [MRR*19] is utilized. In the partial setting, we adapt ‘Bijjective ZoomOut’ [RMOW20] to the partial case, which we describe in what follows.

Partial Bijjective ZoomOut. In the case of shapes with partiality, we extend the Bijjective ZoomOut approach to the partial setting. Throughout this work, S_0 is considered as the full shape and S_i as the partial one for the reference-to-instance maps. Similarly, for the consecutive maps, we designate S_i as the full shape and S_{i+1} as the partial one assuming that their coverage is comparable, although they form part-to-part relations in a strict sense.

Algorithm 1 outlines our method for bidirectional partial corre-

spondence refinement. We first define a new frequency updating rule, based on the slanted diagonal structure in functional maps [RCB*15] relating a partial shape \mathcal{M} to a full shape \mathcal{N} . This updating rule is then used in order to refine the functional map in n_{ZM} iterations, starting from an initial frequency k_{init} and ending with the final frequency $k_{\mathcal{N}}$ on the full shape \mathcal{N} .



(a) ROI transfers by using the refined (reference-to-instance) pointwise correspondences. The reference shape (S_0) with the user-specified ROI is shown on the left. On the right, the instance shapes (S_1, \dots, S_8) are arranged from left to right, and top to bottom.



(b) Colormap transfers between consecutive instance shapes by using the refined inverse pointwise correspondences.

Figure 4: Pairwise correspondences after the refinement on the woman's torso data. (a) ROI pointwise transfers; and (b) inverse pointwise correspondences.

This refinement process is applied to both reference-to-instance (from S_0 to each S_i) and instance-to-instance (from S_i to S_{i+1}) correspondences. The refined correspondences yield denser and less noisy results, as shown in Fig. 4. The ROIs transferred through the pairwise refined maps T_{i0} (Fig. 4a) exhibit smooth borders, and the inverse pointwise maps $T_{i,i+1}$ (Fig. 3b) become denser. However, there are still some shifted ROIs, which can be observed in the first instance shape on the top row.

3.3. Group-wise correspondence refinement

With the bidirectional map refinement implemented, we are now prepared to perform groupwise refinement on both global and local scales within the ROIs. To achieve this goal, we adapt the cycle-consistency refinement scheme [HAGO19] to our case, as presented in Algorithm 2. It consists of five steps: We first transfer the ROI defined on the reference shape S_0 to the instance shapes S_i (step 1), and construct k' harmonics localized on each ROI by using the Localized Manifold Harmonics (LMH) [MRCB17] (step 2). We then convert each input point-wise maps T_{ji} into two different functional maps by using Equation 4. The functional maps

Algorithm 1: Bidirectional Partial ZoomOut

Input : A partial shape \mathcal{M} , a full shape \mathcal{N} and their corresponding eigenfunctions $\{\phi_{k_{\mathcal{M}}}^{\mathcal{M}}\}$ $\{\phi_{k_{\mathcal{N}}}^{\mathcal{N}}\}$, initial pointwise maps $\Pi_{\mathcal{M}\mathcal{N}}$ and $\Pi_{\mathcal{N}\mathcal{M}}$, number of iterations n_{ZM} , initial frequency k_{init}

Output : Refined pointwise maps $\Pi_{\mathcal{M}\mathcal{N}}$ and $\Pi_{\mathcal{N}\mathcal{M}}$ and functional maps $C_{\mathcal{N}\mathcal{M}}$ and $C_{\mathcal{M}\mathcal{N}}$

(1) $r \leftarrow \max_{i=1}^{k_{\mathcal{M}}} \{i | \lambda_i^{\mathcal{M}} < \max_{j=1}^{k_{\mathcal{N}}} \lambda_j^{\mathcal{N}}\}$; $slope \leftarrow \frac{r}{k_{\mathcal{N}}}$;
 $increment \leftarrow \max(\lceil (r - k_{init} + 1) / n_{ZM} \rceil, 0)$;

(2) $k_{\mathcal{M}} \leftarrow k_{init}$;

for $i \leftarrow 1$ **to** n_{ZM} **do**

(3) $k_{\mathcal{M}} \leftarrow k_{\mathcal{M}} + increment$, $k_{\mathcal{N}} \leftarrow \lceil k_{\mathcal{M}} \times slope \rceil$;

(4) Apply 'Bijective ZoomOut' iteration [RMOW20] with $\Pi_{\mathcal{N}\mathcal{M}}$, $\Pi_{\mathcal{M}\mathcal{N}}$, $\phi_{k_{\mathcal{M}}}^{\mathcal{M}}$, and $\phi_{k_{\mathcal{N}}}^{\mathcal{N}}$ as input to obtain refined maps $\Pi_{\mathcal{M}\mathcal{N}}$, $\Pi_{\mathcal{N}\mathcal{M}}$, $C_{\mathcal{N}\mathcal{M}}$ and $C_{\mathcal{M}\mathcal{N}}$

end for

C_{ij}^{MH} maps the K manifold harmonics Φ_{S_i} to the K manifold harmonics Φ_{S_j} , while C_{ij}^{LMH} maps the K' localized harmonics Ψ_{S_i} to the K' localized harmonics Ψ_{S_j} (step 3). The next step consists of creating two distinct canonical consistent latent basis Y_i and Z_i for each shape S_i . By using the CCLB [HAGO19] approach, a consistent basis Y_i of size $K \times M$, $M \leq K$ is extracted from the functional maps C_{ij}^{MH} . In the same manner, localized consistent basis Z_i of size $K' \times M'$, $M' \leq K'$ is extracted from the functional maps C_{ij}^{LMH} (step 4). Finally, we convert the latent basis into point-wise maps by using the nearest neighbors strategy proposed by Huang et al. [HRWO20] (step 5). We provide a detailed description of each step below.

Algorithm 2: Groupwise refinement with ROI

Input : FMN $\mathcal{G}(\mathcal{E}, \mathcal{F})$, with shapes $S = \{S_i\}_{i=0}^{n-1}$ and edges $(i, j) \in \mathcal{E}$, corresponding initial pointwise maps $\{\Pi_{ji}\}$, spectral embeddings sizes K, K', M, M'

Output : Refined point-wise maps $\{\Pi_{ij}\}_{(i,j) \in \mathcal{E}}$, Latent Basis $\{Y_i\}_{i=0}^{n-1}$, ROI localized Latent Basis $\{Z_i\}_{i=0}^{n-1}$

(1) Transfer reference shape (S_0) ROI to each instance shapes S_i ($i = 1, \dots, n-1$).

(2) For each shape S_i , compute K' LMH on the ROI of each shape.

(3) For $(i, j) \in \mathcal{E}$, construct the LMH and MH functional maps C_{ij}^{LMH} and C_{ij}^{MH} by converting the pointwise maps Π_{ji} by using Equation 4.

(4) Compute M and M' global and localized latent bases Y_i and Z_i on each shape S_i by using CCLB [HAGO19].

(5) Recover refined pointwise maps Π_{ji} for each edge (i, j) by using a nearest neighbors strategy between the latent functional embeddings.

ROI transfer (step 1). We transfer the region of interest (ROI) de-

defined on the reference shape S_0 to each instance shape S_i . In most cases, the transfer is done by assigning to each point $p \in S_i$ the value of its corresponding point $q = \Pi_{i,0}(p)$ on S_0 . The resulting ROI on the target shape S_i is represented by a binary mask (1 inside the ROI and 0 elsewhere). Depending on the dataset considered, however, such pointwise transfer may lead to reduced or misaligned ROIs on S_i . To address this, we also incorporate the functional transfer of the ROI. Because of the low pass filter representation of the harmonics, the ROI by functional transfer is smoother and can represent a wider region than expected. Note that conversion from functional maps to pointwise maps is a challenging optimization problem in itself.

LMH computation (step 2). In this step, we compute K' LMH Ψ_i on each shape S_i . The generated harmonics are defined on the whole shape and are compatible with the standard manifold harmonics. Depending on the ROI considered and the ROI transfer method, the LMH parameters η and μ are adjusted to promote the orthogonality of the obtained harmonics with respect to the manifold harmonics and their locality on the input region of interest.

Point-wise to functional map conversion (step 3). We use the Equation (4), $C_{ji} = \Phi_{S_i}^\dagger \Pi_{ij} \Phi_{S_j}$, to convert the initial point-wise maps $\{\Pi_{ij}\}_{(i,j) \in \mathcal{E}}$ to global and local functional maps. Applying this conversion with Φ_{S_i} and Φ_{S_j} yields global functional maps C_{ij}^{MH} , while utilizing the localized harmonics Ψ_{S_i} and Ψ_{S_j} generates local functional maps C_{ij}^{LMH} .

Global and local consistency with CCLB (step 4). To ensure the cycle-consistency of the current local and global functional correspondences, we construct the CCLB [HAGO19] using the functional maps obtained in step 3 as input. The global latent bases Y_i obtained with C_{ij}^{MH} provide a compact and coherent representation of the shapes on a global scale. Additionally, the localized latent bases Z_i offer a consistent representation only focused on the ROI.

Pointwise recovery (step 5). Point-wise maps are then recovered from the CCLB Y_i and Z_i by using the pointwise map conversion procedure as in Huang et al [HRWO20]. Denote the CCLB as $\{Y_i^k\}_{i=0}^{n-1}$ and $\{Z_i^k\}_{i=0}^{n-1}$, then the problem is formulated as :

$$T_{ij}(p) = \arg \min_{q \in S_j} \|\Gamma_j(q)^T - \Gamma_i(p)^T\|_2, \quad \forall p \in S_i, \quad (7)$$

with $\Gamma_i = [\Phi_i^k Y_i^m, \Psi_i^{k'} Z_i^{m'}]$ and $\Gamma_j = [\Phi_j^k Y_j^m, \Psi_j^{k'} Z_j^{m'}]$, where $[\cdot, \cdot]$ represents the concatenation of the two arrays. This way, the pointwise conversion boils down to the nearest neighbor search for each row of Γ_i among the rows of Γ_j . By concatenating consistent local $\Psi_i^{k'} Z_i^{m'}$ (cLMH) and global $\Phi_i^k Y_i^m$ (cMH) harmonics, we incorporate higher localized frequencies on ROI to the nearest neighbors search, thereby improving the accuracy of the conversion.

4. Dataset

Woman's torso. Our first dataset is a set of female torsos acquired from a patient undergoing post-operative breast radiotherapy. It is composed of one reference surface reconstructed from the CT image, along with multiples scans of the frontal torso surface acquired at each radiotherapy session, which were spaced approximately two weeks apart. Each scan captures different patient poses, respiratory

phase, and varying degrees of breast inflation due to radiation therapy. Additionally, the scanned regions also differ from one another, making a particularly challenging setting to the problem.

Animated face. A subset of a 3D animated head model has been extracted from a 3D facial expression dataset, CoMA [RBSB18]. The head model represents a sequence of 3D facial scans captured from a real person demonstrating an expression ("mouth side" in this work). In addition to the raw scans composed of a variable number of vertices with noises and several holes, we also utilized registered meshes in our experiments. These meshes, obtained through a model-based optimization [LBB*17], share the share topology as the template mesh (consisting of 5023 vertices and 9976 triangles) and contain holes for the mouth and the neck.

Organic shapes from SHREC. A number of organic shapes, shark, ant, and boy, have been chosen from the SCHREC11 and SHREC15 datasets [LGB*11] [GDB*15]. We utilize the multiple poses available for each entity as a shape collection.

5. Implementation

Preprocessing. We remesh the surface meshes of the woman's torso dataset to approximately 10K vertices by using the VCG (Visualization and Computer Graphics Library) surface reconstruction filter in Meshlab [CCC*08]. Meshlab was also used to ensure that real world meshes (Woman's torso and Animated faces) were single connected components. For every shape in all datasets, we compute Laplacian eigenfunctions by using the point cloud method proposed by Sharp and Crane [SC20].

Pairwise shape matching. Two different methods have been used for the initial map depending on the dataset. For the real-world datasets with partiality (such as woman's torso), we used FSPM [LRBB17]. We perform 10 iterations of FSPM part-to-full optimization, with the dense and sparse penalties set to 0.5 and 30, respectively, and the maximum number of ICP iterations to 25. As in [LRBB17], we select the top 10% of the best correspondences and use a farthest point strategy to construct delta functions as sparse but well spread descriptors to further refine the correspondences. 400 delta functions with radius 7 have been used for this step. Finally, we used a 512 dimensional SHOT [TSDS10] descriptor (instead of a 352 as proposed in the paper) to better deal with the noise.

As for the others (animated face and animals), we make use of the python implementation of the functional map optimization [Mag22] to compute initial functional maps with 60 eigenfunctions, 5 subsample steps, 100 Weight Kernel Signature (WKS) descriptors and the following weights to the energy terms: $w_{desc} = w_{dcomm} = w_{orient} = 1$, and $w_{lap} = 1e-1$.

Bidirectional pairwise refinement. In case of datasets with partiality, we perform 5 iterations of the bidirectional partial ZoomOut (Sec. 3.2) initialized with functional maps of size 40×40 obtained through FSPM. The output is rectangular functional maps of size $r \times 60$, where r is determined by using the partiality between the two input shapes in a similar fashion to [RCB*15].

Cycle-consistent refinement. We implemented the computation of

LMH in python to obtain localized harmonics on the various shape collections. The LMH parameters used are summarized in Table 1.

	μ	η
Woman's torso	$1e^{-2}$	$1e^1$
Organic shapes (SHREC)	$1e^5$	$1e^{10}$
Animated face	$1e^{-1}$	$1e^5$

Table 1: Locality (μ) and orthogonality (η) LMH parameters used in our experiments.

To compute the canonical consistent latent bases (CCLB) Y_i and Z_i , we use the "pyFM" [Mag22] implementation using the initial functional maps containing the manifold harmonics, and the functional maps defined only on the LMH, respectively.

6. Results

6.1. Initial ROI transfer

The initial ROI transfer itself can pose challenges depending on the nature of the data collection. As expected, the real-world data showed more noticeable errors overall, leaving a greater need for further improvements during the refinement steps. Fig. 5 shows the initial ROI transfer we obtained by our refined pairwise correspondences on the woman's torso dataset, real-world raw scans exhibiting a considerable level of partiality. The transferred ROIs show errors ranging from moderate to noticeable.




Figure 5: Initial ROI transfer on the woman's torso dataset: FSPM followed by the bidirectional refinement has been used to transfer ROI from S_0 (leftmost) to each S_i , in a pairwise manner.

The raw scans of animated faces show another difficult case of the initial transfer. The chosen ROI (lips), constituting high-frequency features of the face, present a challenging region for accurate transfer, particularly due to the presence of large nonlinear deformations among shape instances. This becomes even more evident when comparing the transfer results on the synthetic and raw scan versions of the animated face, as shown in Fig. 6 and 7. The ROI is conserved well for the synthetic dataset (Fig. 6), both the functional transfer and the pointwise transfer results. The considered shapes are more or less similar, and the use of global harmonics is sufficient. In case of the raw scans (Fig. 7), on the other hand, the functional transfer completely failed, due to the lack of sufficiently high frequencies to represent and transfer the lips region. It confirms that the global nature of the manifold harmonics is not suited for the partial case, resulting in the correspondences of poor quality. This issue has been mitigated by using the product manifold transfer, as depicted in Fig. 8.




Figure 6: The initial ROI transfer from the reference (the leftmost shapes) to each of the instance shapes, on the synthetic dataset of the animated face.




Figure 7: The initial ROI transfer from the reference (the leftmost shapes) to each of the instance shapes, on the raw scan dataset of the animated face.



Figure 8: Functional transfer on the raw scan dataset of the animated face by using the product harmonics [VLR* 17].

6.2. Consistent LMH (cLMH)

Fig. 9 and 10 show the original LMH [MRCB17] (on the left side) and our consistent LMH (on the right), for the woman's torso data and the raw scans of animated faces, respectively. Since LMH harmonics are defined independently on each shape, they are not consistent across the entire collection although some shapes possess similar frequencies. We can observe, for instance, different harmonics in the $k=4$ row of Fig. 9, where the second and third harmonics are of opposite sign and different from the fourth shape. With our approach, the obtained harmonics are both localized and consistent across all shapes. In Fig. 10 is shown the different harmonics obtained with a latent space reduced to $M' = 6$ or $M' = 10$. The resulting harmonics differ according to the final size of the latent space. Indeed, with $M' = 6$, the first harmonics exhibit lower frequencies, whereas with $M' = 10$, the first harmonics present higher frequencies. In both cases, we observe that the resulting harmonics exhibit consistency across all shapes in the collection, while also being concentrated on the ROI (lips).

k'	Localized Harmonics [MRCB17]	m'	Consistent Localized Harmonics (ours) (10 from 20)
0		0	
2		2	
4		4	
6		6	

Figure 9: Localized manifold harmonics (LMH) $\Psi_{S_i}^k$, $k' = 0, 2, 4, 6$ (left) and consistent LMH $\Psi_{S_i}^k Z_i$, $m' = 0, 2, 4, 6$ (right) generated by using our approach on the woman's torso dataset.

	Localized Harmonics [MRCB17]	cLMH 6 from 20 (ours)	cLMH 10 from 20 (ours)
$k=0$			
$k=2$			
$k=4$			

Figure 10: Localized manifold harmonics Ψ_k , $k = 0, 2, 4$ on the raw scans of the animated face. Each cell of the tabular shows the k_{th} localised harmonics on each shape of the collection. The left column shows the harmonics computed on the initial ROI by using the standard [MRCB17] approach. The middle and the right columns present the consistent harmonics using our approach with a reduction from 20 LMH to respectively 6 and 10 cLMH harmonics.

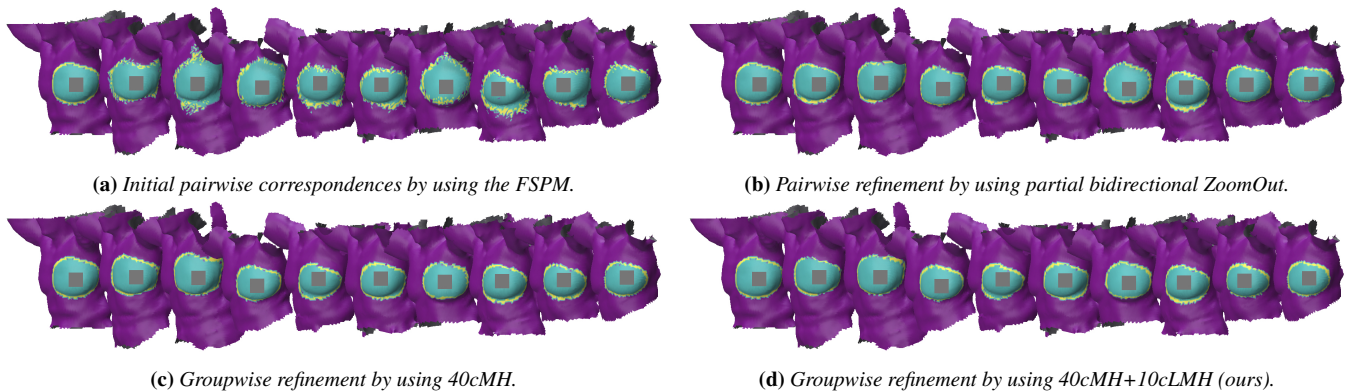


Figure 11: Pointwise ROI transfer results on the woman's torso.

(a)	40MH	50MH
Functional standard		
Extension with product harmonics		
(b)	30cMH	40cMH
Without cLMH		
With 10 cLMH (ours)		

Figure 12: ROI functional map transfer on the raw scans of the animated face. (a) The initial pairwise transfer by using the standard harmonics (1st row) fails to preserve regions. We therefore use its extension with the product harmonics [VLR*17] (2nd row) to obtain a favorable initialization, resulting in transfers that are more localized on the lips region. (b) Groupwise refinement by using cMH alone (1st row) does not conserve the ROI well, while the incorporation of merely 10 cLMH (2nd row) demonstrates effective region transfers.

6.3. ROI transfer with consistent LMH

Womans' torso. Fig. 11 shows the pointwise transfer of the ROI on the woman's torso collection by using different approaches. Fig. 11a and 11b present the ROI transfer results obtained by pairwise approaches. We can observe that the initial partial correspondences obtained by the FSPM approach lead to noisy transfer, often with some region shifts. The 2nd and the 8th shapes in Fig. 11a present a significant downward shift, for instance. The use of the bidirectional refinement 11b leads to smoother ROI transfers overall. Additionally, certain regions, such as the 7th and the 8th shape, are better positioned compared to the FSPM transfer. Fig. 11c shows the ROI pointwise transfer obtained with the cycle-consistent approach [HAGO19] with 30 consistent MH.

Fig. 11d depicts the pointwise ROI transfer result using our approach, by extending the latent bases with localised consistent harmonics. Incorporating cLMH in addition to the cMH leads to a ROI closer to the initial one, as can be observed on the second and third shapes, among others. Fig. 13 further shows the improved ROI transfer result by incorporating the cLMH in addition to the cMH. The inclusion of consistent LMH contributed to improved localization of the breast ROI, resulting in a rounder shape, as demonstrated in the second column.

Raw scans of animated face. Fig. 12 shows the ROI transfer results on the raw scans of animated face by using different harmonics. The initial pairwise transfer by using the standard harmonics (Fig. 12a, 1st row) fails to preserve regions. We therefore use its extension with the product harmonics [VLR*17] (Fig. 12a, 2nd

Reference	ROI transfer results		
	30cMH	40cMH	30cMH+10cLMH (ours)

Figure 13: ROI pointwise transfer from the reference to a shape instance, with 30 (2nd) and 40 harmonics (3rd and 4th). Adding 10 cMH to 30 cMH (40 cMH, 3rd) does not really improve the quality of the transfer, whereas adding 10 cLMH (4th) does.

row) to obtain a favorable initialization, resulting in transfers that are more localized on the lips region. The standard approach fails to transfer the region while the extended approach, while not very precise, leads to transfer that are more localized on the lips. Fig. 12b shows the groupwise refinement results with different numbers of consistent Manifold Harmonics (cMH) and Localized Manifold Harmonics (cLMH). We can observe that the functional transfer by using cMH alone ('without cLMH' on 1st row) does not conserve the ROI well, while the incorporation of merely 10 cLMH (2nd row) demonstrates effective region transfers.

Others. Fig. 14 and 15 show our results on four shape collections: shark (Fig. 14a), ant (Fig. 14b), synthetic animated face (Fig. 15a),

Table 2: Mean relative pointwise cyclic error (% , normalized by the bounding box size) on four shape collections in Fig. 14 and 15, arranged in the reading order. The compared methods are the refined pairwise correspondences (Pairwise), the groupwise refinement (Groupwise), and our approach (Ours). The methods coupled with the "ROI" suffix present the error computed only on the vertices on the ROI.

Shape	Shark					Ant				Animated face					Boy				
	S1	S2	S3	S4	S5	A1	A2	A3	A4	F1	F2	F3	F4	F5	B1	B2	B3	B4	B5
Pairwise	34.62	22.79	18.27	15.69	30.96	15.22	36.07	16.95	16.16	2.91	2.53	1.57	2.97	2.73	2.91	7.25	8.34	2.55	2.73
Groupwise	1.15	1.04	0.88	1.57	2.10	4.28	11.10	3.72	4.04	0.11	0.12	0.12	0.17	0.18	1.50	1.61	1.12	1.20	1.39
Ours	1.11	0.98	0.81	1.46	1.19	4.23	11.09	3.65	3.99	0.21	0.19	0.19	0.25	0.31	1.50	1.56	1.01	1.23	1.37
Pairwise ROI	18.44	18.18	15.09	7.36	19.41	2.79	42.00	2.44	3.09	9.76	9.14	6.70	10.65	7.77	0.33	2.28	5.45	0.38	0.44
Groupwise ROI	1.57	1.62	1.48	1.99	5.02	3.09	2.86	3.74	3.49	0.31	0.32	0.41	0.45	0.52	0.23	0.33	0.32	0.28	0.43
Ours ROI	1.44	1.29	1.04	1.11	1.04	2.32	1.99	2.29	2.48	0.23	0.24	0.26	0.46	0.46	0.21	0.28	0.29	0.30	0.33

and boy (Fig. 15b). For each shape collection, we show the functional ROI transfer (left), pointwise ROI transfer (middle), and the colormap transfer (right). The reference shape is on the leftmost side within each cell, with the ROI highlighted in yellow. The results obtained by using the refined pairwise correspondences (top), their groupwise refinement, [HAGO19] (middle), and our method (bottom) are illustrated.

As expected, the use of the cLMH results in a more localized functional transfer of ROI. This is especially clear in the case of shark (Fig. 14a) and synthetic animated faces (Fig. 15a). For instance, the initial and groupwise correspondences computed on the shark show the functional transfer has succeeded on the tail region but failed on the fin. Moreover, the transfers lack concentration on the ROI, exhibiting non-zero function values across the entire shapes. In contrast, our method achieves an improved functional transfer, focusing on both regions while maintaining low values in other areas. On the synthetic animated face (Fig. 15a), we observe a degraded transfer when using the groupwise correspondences compared to the initial ones. This is contributable to the lack of sufficiently high frequencies to represent the lips when using truncated Manifold Harmonics (MH) alone. In contrast, the proposed method leads to a focused and localized functional ROI transfer solely on the lips, rather than being distributed around its surroundings. In the case of the ant (Fig. 14b) and the boy (Fig. 15b), the initial or groupwise correspondences already achieve accurate positioning of the region of interest in the functional transfers. This is because the ROIs are located in regions of sufficiently low frequency, thus they can be represented well using global harmonics. The improved performance of the consistent LMH can also be seen in the pointwise transfers. On the shark (Fig. 14a), the pointwise transfer by our approach has successfully included the top of the fin within the ROI, in contrast to the groupwise approach. Similarly, the noises have been corrected around the base of antennas on the ant (Fig. 14b), and on the chin of the synthetic animated face (Fig. 15a), and on the shoulder of the boy (Fig. 15b). Our approach persistently focuses on the local region, even in cases where the global map is erroneous, as can be observed on some legs of the ants (Fig. 14b). Table 2 presents the mean vertex-wise cycle error on each shape of the four shape collections. The error is computed as the averaged geodesic distance between a vertex and its counterpart obtained through the best cyclic correspondence transfer, e.g. from S_0 to S_1, \dots , and back to S_0 . As expected, our method improves over other methods under comparison, especially in the ROI. As a final note, we observe that in many cases, the use of consistent cLMH alone leads to compa-

table results to combining it with cMH in terms of ROI transfer, demonstrating the effectiveness of our strategy.

6.4. Limitations

Our method still has some limitations. To ensure the effective construction of the CCLB, our method relies on initial maps of sufficiently good quality, which may be difficult to obtain in some challenging, real-world cases. Additionally, the LMH parameters have been adjusted heuristically in our experiments. Developing an algorithm to automate this process would broaden its usability across various applications and user groups. While the pointwise recovery from the functional maps has been performed using a nearest neighbor strategy in this paper, our consistent LMH are not tied to this choice. Other pointwise conversions, such as the precise maps [EBC17] or the product harmonics [VLR*17] can benefit from this consistent functional representation of the ROI.

7. Conclusions

We have presented a new method for refining non-rigid correspondences within a collection of 3D shapes. It is achieved by introducing the notion of localised harmonics on a user-specified region of interest and incorporating them in the Canonical Consistent Latent Basis construction. We have demonstrated their use for the consistent functional transfer of the region of interest (ROI) as well as their positive impact in the pointwise map reconstruction from the functional maps. Compared to existing works, our method demonstrates enhanced robustness to noise, and can handle partiality commonly present in real-world data. We believe that the ability to focus on a specific region, especially when dealing with a collection of partial shapes, is important for many applications in medicine, biology and engineering. Given the direct relevance to the construction of the 'limit shape' [HAGO19] through the CCLB framework, our approach can also be used for the analysis of shape collections, particularly focusing on the ROI. We intend to further investigate this in the future.

Acknowledgements

Pierre Galmiche has been supported by the French ministry of higher education. This work has been partially supported by the French national project Human4D (ANR-19-CE23-0020).

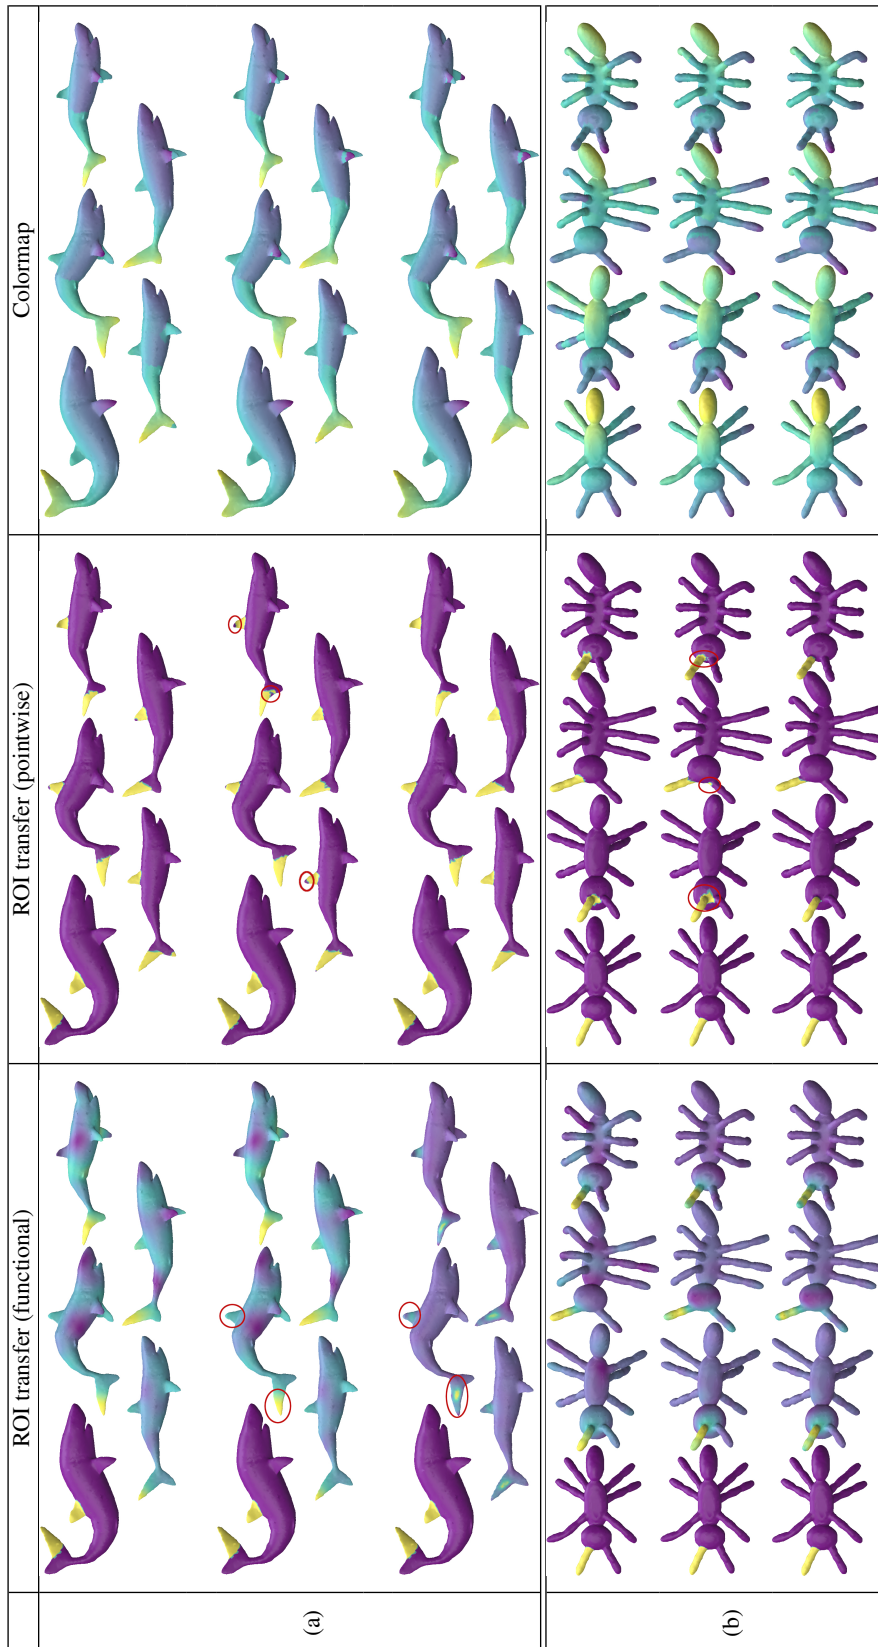


Figure 14: Functional (left) and pointwise (middle) transfers of the ROI and colormap (right) obtained on the shark and the ant collections. Each cell illustrates the results of the refined pairwise correspondences (top), the groupwise refinement (middle), and our approach (bottom).

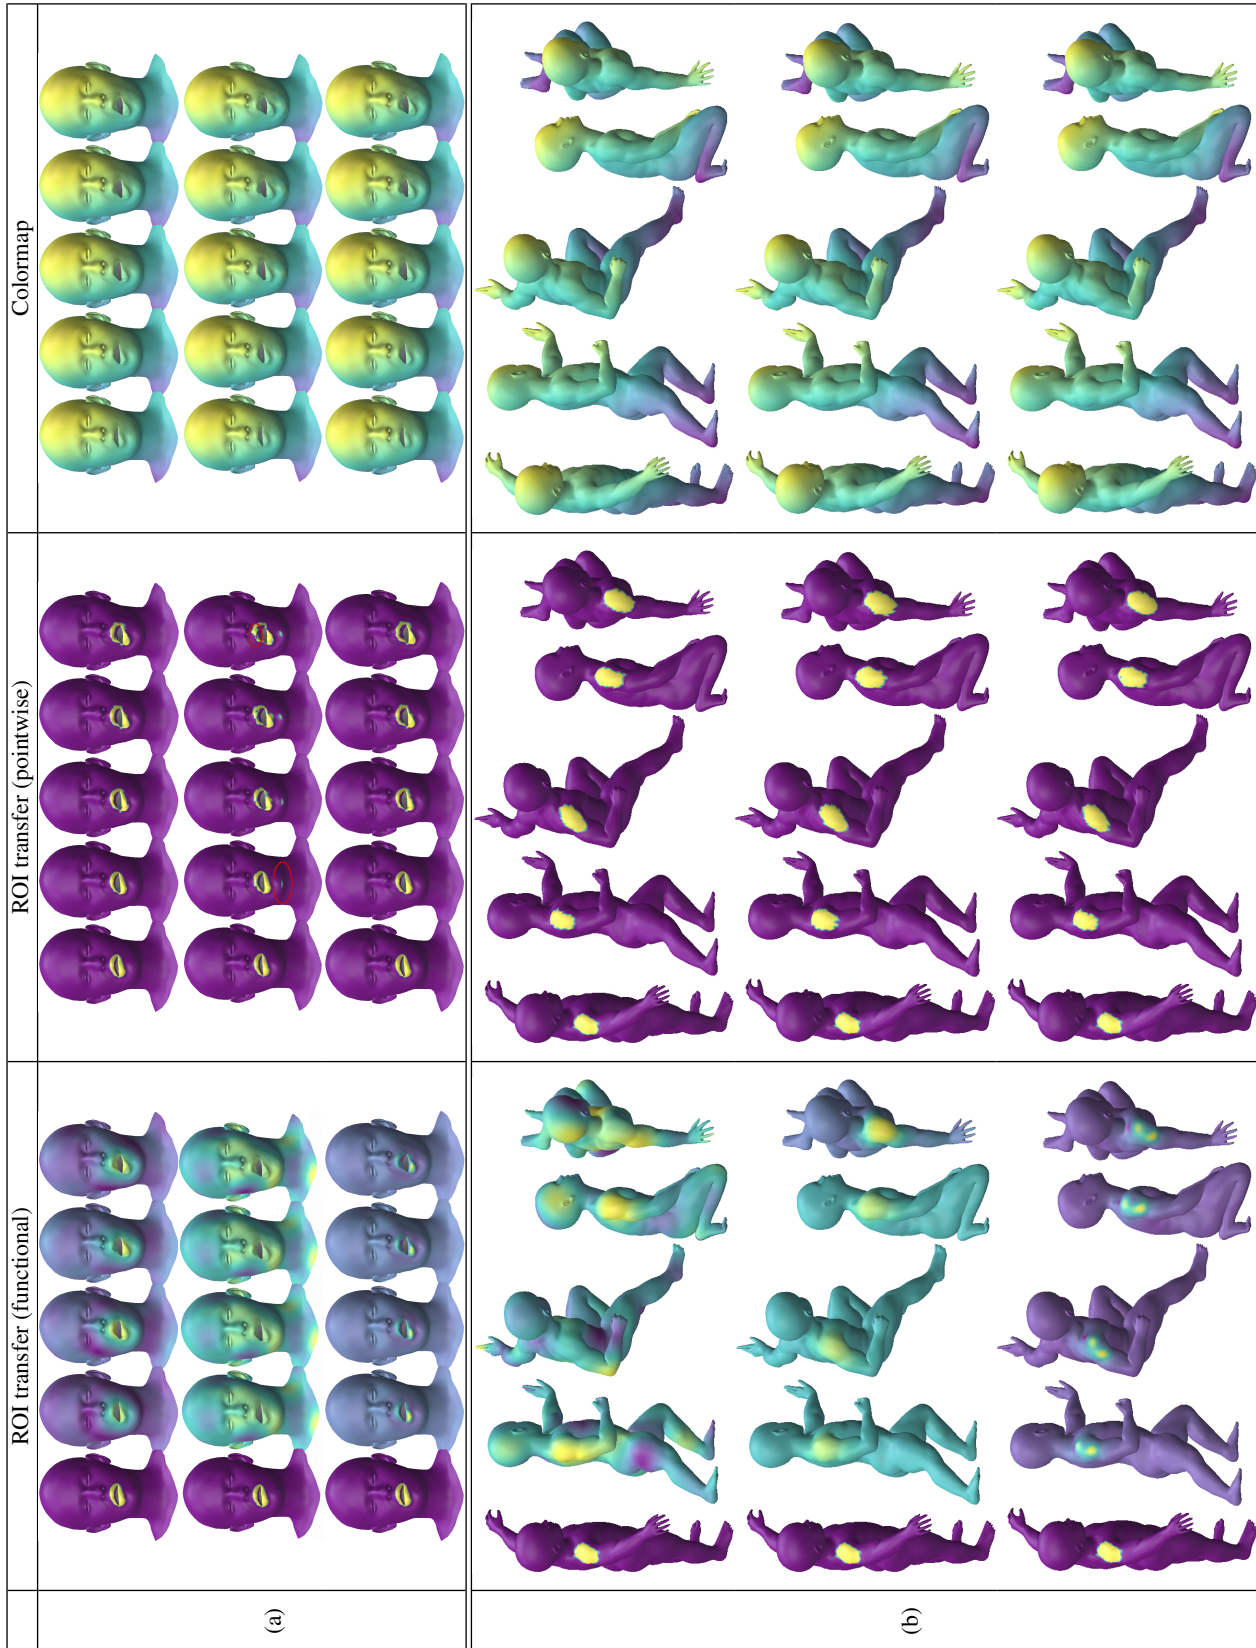


Figure 15: Functional (left) and pointwise (middle) transfers of the ROI and colormap (right) obtained on the animated face and the boy collections. Each cell illustrates the results of the refined pairwise correspondences (top), the groupwise correspondences (middle), and our approach (bottom).

References

- [ACPO3] ALLEN B., CURLESS B., POPOVIĆ Z.: The space of human body shapes: Reconstruction and parameterization from range scans. *ACM Trans. Graph.* 22, 3 (jul 2003), 587–594. 1
- [AGB*21] ANTONSANTI P.-L., GLAUNÈS J., BENSEGHIR T., JUGNON V., KALTENMARK I.: Partial matching in the space of varifolds, 2021. 1
- [APO21] ATTAIKI S., PAI G., OVSJANIKOV M.: DPFM: deep partial functional maps. *CoRR abs/2110.09994* (2021). 2
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: Scape: Shape completion and animation of people. *ACM Trans. Graph.* 24, 3 (jul 2005), 408–416. 1
- [ATCO*10] AU O., TAI C.-L., COHEN-OR D., ZHENG Y., FU H.: Electors voting for fast automatic shape correspondence. *Comput. Graph. Forum* 29 (05 2010), 645–654. 1
- [BFF*07] BIASOTTI S., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., MARINI S., PATANÉ G., SPAGNUOLO M.: 3D Shape Description and Matching Based on Properties of Real Functions. In *Eurographics 2007 - Tutorials* (2007), Myszkowski K., Havran V., (Eds.), The Eurographics Association. 1
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (Piscataway, NJ, USA, June 2014), IEEE. 1
- [CCC*08] CIGNONI P., CALLIERI M., CORSINI M., DELLEPIANE M., GANOVELLI F., RANZUGLIA G.: MeshLab: an Open-Source Mesh Processing Tool. In *Eurographics Italian Chapter Conference* (2008), Scarano V., Chiara R. D., Erra U., (Eds.), The Eurographics Association. 6
- [CRM*16] COSMO L., RODOLA E., MASCI J., TORSSELLO A., BRONSTEIN M. M.: Matching deformable objects in clutter. In *2016 Fourth International Conference on 3D Vision (3DV)* (2016), IEEE, pp. 1–10. 2
- [CZ08] CHANG W., ZWICKER M.: Automatic registration for articulated shapes. *Comput. Graph. Forum* 27 (07 2008), 1459–1468. 1
- [EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. *Computer Graphics Forum* 36, 5 (2017), 165–174. 10
- [EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25, 10 (2003), 1285–1295. 1
- [GDB*15] GODIL A., DUTAĞACI H., BUSTOS B., CHOI S., DONG S., FURUYA T., LI H., LINK N., MORIYAMA A., MERUANE R., OHBUCHI R., PAULUS D., SCHRECK T., SEIB V., SIPIRAN I., YIN H., ZHANG C.: Shrec'15: Range scans based 3d shape retrieval. 6
- [GLT*20] GAO M., LÄHNER Z., THUNBERG J., CREMERS D., BERNARD F.: Isometric multi-shape matching, 2020. Number: arXiv:2012.02689. 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 *Symposium of Geometry Processing 2019* (Milan, Italy, July 2019), vol. 38, pp. 187–202. 1, 2, 3, 4, 5, 6, 9, 10
- [HRWO20] HUANG R., REN J., WONKA P., OVSJANIKOV M.: Consistent zoomout: Efficient spectral map synchronization. *Computer Graphics Forum* 39, 5 (2020), 265–278. 2, 5, 6
- [HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.* 33, 4 (jul 2014). 2, 3
- [KLCF10] KIM V., LIPMAN Y., CHEN X., FUNKHOUSER T.: Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum (Symposium on Geometry Processing)* 29, 5 (July 2010). 1
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Trans. Graph.* 30, 4 (jul 2011). 1
- [LBB*17] LI T., BOLKART T., BLACK M. J., LI H., ROMERO J.: Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.* 36, 6 (nov 2017). 6
- [LGB*11] LIAN Z., GODIL A., BUSTOS B., DAOUDI M., HERMANS J., KAWAMURA S., KURITA Y., LAVOUÉ G., NGUYEN H. V., OHBUCHI R., OHKITA Y., OHISHI Y., PORIKLI F., REUTER M., SIPIRAN I., SMEETS D., SUETENS P., TABIA H., VANDERMEULEN D.: SHREC '11 Track: Shape Retrieval on Non-rigid 3D Watertight Meshes. In *Eurographics Workshop on 3D Object Retrieval* (2011), Laga H., Schreck T., Ferreira A., Godil A., Pratikakis I., Veltkamp R., (Eds.), The Eurographics Association. 6
- [LRB*20] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D.: Non-rigid puzzles. *CoRR abs/2011.13076* (2020). 2
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Fully Spectral Partial Shape Matching. *Computer Graphics Forum* (2017). 3, 4, 6
- [Mag22] MAGNET R.: pyFM - Python bindings for functional maps. <https://github.com/RobinMagnet/pyFM>, 2022. 6, 7
- [MRC*21] MARIN R., RAMPINI A., CASTELLANI U., RODOLÀ E., OVSJANIKOV M., MELZI S.: Spectral Shape Recovery and Analysis Via Data-driven Connections. *International Journal of Computer Vision* 129, 10 (Oct. 2021), 2745–2760. 2
- [MRCB17] MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M. M.: Localized manifold harmonics for spectral shape analysis, 2017. Number: arXiv:1707.02596. 1, 4, 5, 7, 8
- [MRR*19] MELZI S., REN J., RODOLÀ E., WONKA P., OVSJANIKOV M.: Zoomout: Spectral upsampling for efficient shape correspondence. *CoRR abs/1904.07865* (2019). 1, 4
- [NMR*18] NOGNENG D., MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M., OVSJANIKOV M.: Improved functional mappings via product preservation. *Computer Graphics Forum* 37, 2 (2018), 179–190. 1
- [OBBS*12] 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 (jul 2012). 1, 2
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 1 (1993), 15–36. 2
- [RBSB18] RANJAN A., BOLKART T., SANYAL S., BLACK M. J.: Generating 3D faces using convolutional mesh autoencoders. In *European Conference on Computer Vision (ECCV)* (2018), pp. 725–741. 6
- [RCB*15] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSSELLO A., CREMERS D.: Partial functional correspondence, 2015. Number: arXiv:1506.05274. 1, 2, 5, 6
- [RMC17] RODOLÀ E., MOELLER M., CREMERS D.: Regularized point-wise map recovery from functional correspondence. *Computer Graphics Forum* 36, 8 (2017), 700–711. 1
- [RMOW20] REN J., MELZI S., OVSJANIKOV M., WONKA P.: Maptree: Recovering multiple solutions in the space of maps. *ACM Trans. Graph.* 39, 6 (nov 2020). 1, 2, 4, 5
- [RMWO21] REN J., MELZI S., WONKA P., OVSJANIKOV M.: Discrete optimization for shape matching. *Computer Graphics Forum* 40, 5 (2021), 81–96. 1
- [SC20] SHARP N., CRANE K.: A laplacian for nonmanifold triangle meshes. *Computer Graphics Forum* 39, 5 (2020), 69–80. 6
- [SSGD03] SUNDAR H., SILVER D., GAGVANI N., DICKINSON S.: Skeleton based shape matching and retrieval. In *2003 Shape Modeling International*. (2003), pp. 130–139. 1
- [TBW*09] TEVS A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.: Isometric registration of ambiguous and partial data. In

- 2009 *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009)*, 20-25 June 2009, Miami, Florida, USA (2009), IEEE Computer Society, pp. 1185–1192. [1](#)
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.-P.: Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum* (2011). [1](#)
- [TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proceedings of the 11th European Conference on Computer Vision Conference on Computer Vision: Part III* (Berlin, Heidelberg, 2010), ECCV'10, Springer-Verlag, p. 356–369. [3](#), [6](#)
- [VLR*17] VESTNER M., LITMAN R., RODOLA E., BRONSTEIN A., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), IEEE, pp. 6681–6690. [7](#), [9](#), [10](#)
- [WHG13] WANG F., HUANG Q., GUIBAS L. J.: Image co-segmentation via consistent functional maps. In *2013 IEEE International Conference on Computer Vision* (2013), IEEE, pp. 849–856. [2](#), [3](#)