

CONSISTENT ZOOMOUT: Efficient Spectral Map Synchronization

Ruqi Huang¹, Jing Ren², Peter Wonka², Maks Ovsjanikov¹

¹LIX, École Polytechnique, CNRS ²KAUST

Abstract

In this paper, we propose a novel method, which we call CONSISTENT ZOOMOUT, for efficiently refining correspondences among deformable 3D shape collections, while promoting the resulting map consistency. Our formulation is closely related to a recent unidirectional spectral refinement framework, but naturally integrates map consistency constraints into the refinement. Beyond that, we show further that our formulation can be adapted to recover the underlying isometry among near-isometric shape collections with a theoretical guarantee, which is absent in the other spectral map synchronization frameworks. We demonstrate that our method improves the accuracy compared to the competing methods when synchronizing correspondences in both near-isometric and heterogeneous shape collections, but also significantly outperforms the baselines in terms of map consistency.

1. Introduction

Finding correspondences across 3D shapes is a fundamental task in geometry processing and computer graphics that enables a wide range of applications. Typically, users favor correspondences differently according to their task at hand. For instance, conformal (angle preserving) maps are suitable for texture transfer [APL15], orientation-preserving maps are desired for surface or volume parameterization [RPPSH17], and maps of low metric distortion are critical in shape retrieval [BBK06] and statistical shape analysis [ASK*05, HSS*09].

Despite the large variability of specific preferences in practice, there exist some universal principles in seeking high-quality maps between shapes. One essential requirement is that a map should be close to bijective, which allows transferring information across shapes consistently, without directional bias. More generally, map consistency is desirable in the context of shape collections. However, integrating the consistency or bijectivity constraint into objectives tailored to practical applications is often non-trivial and can lead to difficult optimization problems.

In this paper, we consider deformable 3D shape matching where the goal is to find near-isometric maps between shapes. To this end, many works have been proposed to promote consistency (i.e., bijectivity) in the pairwise setting [ERGB16, HO17, RPWO18, ESBC19]. Moreover, it has long been observed that consistency can play an important role in map denoising [HG13], in the context of shape collections. Stimulated by this observation, various data-driven approaches [NBCW*11, WHG13, CRA*17, HLW*19] have been proposed for synchronizing shapes in collections. Among these approaches, a noticeable recent trend is the use of spectral techniques, especially through the notion of functional maps [OBCS*12]. By using multi-scale spectral bases, func-




Figure 1: ZOOMOUT can fail in the presence of poor initialization quality or large non-isometric deformations, while our method handles these challenging cases by jointly refining maps in a collection.

tional maps encode correspondences as linear operators (small-sized matrices in practice) that transfer real-valued functions across shapes. The algebraic nature of functional maps makes this framework well-suited for map manipulation, such as composition and inversion, which in turn facilitates formulating map consistency constraints.

It is worth noting, however, that functional maps are usually expressed in low-dimensional spectral bases. Apart from the efficiency consideration, directly encoding and refining functional maps in the high-dimensional spectral spaces can be unstable and error-prone, especially in the presence of noisy initial constraints. On the other hand, the reduced dimensionality can impede recovering high-quality maps from the functional representation, due to

the loss of details encoded in the medium and high-frequency parts of the spectral bases.

The dimensionality issue has long been a bottleneck in shape matching via functional maps. A recent framework tackles this issue by proposing ZOOMOUT refinement [MRR*19], which iteratively updates maps in the spectral domain and finally allows to recover accurate point-wise correspondences from high-dimensional functional maps. Nevertheless, ZOOMOUT refinement is *unidirectional*, making it difficult to incorporate bijectivity constraints, let alone map consistency in a shape collection.

In this paper, we formulate a novel multi-scale synchronization approach, which can be seen as an extension of both ZOOMOUT to promote consistency in a map collection, and an extension of the basic functional maps synchronization [WHG13] to promote consistency on every leading principal sub-block of the functional map matrices. Beyond that, we further show that our formulation can be adapted to recover the underlying isometry among near-isometric shape collections with a theoretical guarantee, which is absent in the other spectral map synchronization frameworks.

Based on our formulation, we propose a practical CONSISTENT ZOOMOUT method, which gracefully incorporates map consistency constraints into a simple iterative scheme. Fig. 1 shows shapes from the SRHEC'07 [GBP07] dataset, for which given the same initial maps, ZOOMOUT can fail in the presence of large non-isometric deformations, while by jointly refining maps in the collection, our CONSISTENT ZOOMOUT finds more accurate and consistent correspondences.

Our CONSISTENT ZOOMOUT can be applied to refine maps both between a pair of shapes and in the context of shape collections. We show that our pipeline achieves similar or better accuracy than that of competing methods in both near-isometric and heterogeneous shape collections, and especially outperforms the baselines in terms of consistency.

2. Related Work

Non-rigid shape matching or alignment is a core topic in geometry processing [VKZHCO11, TCL*13, BCBB16]. An exhaustive review is beyond the scope of this paper. We review most closely related techniques, focusing especially on matching heterogeneous shape collections, map synchronization, and adaptive basis selection in functional map estimation.

Maps in Shape Collections While most non-rigid shape matching techniques concentrate on the pairwise setting, several methods have been proposed to find correspondences in the context of shape collections. Such approaches typically leverage cycle-consistency constraints to improve maps computed between each pair of shapes. Early methods in this domain concentrated on detecting and correcting inconsistent cycles to enforce consistency constraints [Hub02, ZKP10, NBCW*11]. This can be expensive as potentially many cycles might need to be tested, and may fail in resolving local errors. More recent approaches have tried to exploit the link between cycle consistency and low-rank properties of matrices that encode maps in the entire collection [HG13, WS13, LZD17] leading to relatively simple formulations that enjoy optimality guaran-

tees, e.g. [HG13]. To alleviate the computational burden of handling potentially very large matrices various approaches have been proposed using alternating minimization [ZZD15], sparse modeling [CRA*17], reweighted least squares [CMG13, HLBH17], iterative factorization [ARFF18, BTGT19], and relaxations based on eigen-decomposition [PKS13, SHSS16] among others.

Most closely related to ours are methods in this domain that use the functional map representation [WHG13, HWG14, HAGO19, SVBC19]. This representation is particularly well-suited for map synchronization first because it allows to encode maps as small size matrices in a reduced basis, and second because it naturally enables an interpretation of maps as information carriers, allowing both a rich set of regularization constraints coming from the pairwise setting [OCB*17] and enforcing consistency through map re-routing via functional “latent shapes” [WHG13, HAGO19]. As a result, existing functional map-based techniques for map synchronization are particularly efficient due to the use of a compact map representation. At the same time, their expressive power, and thus the accuracy of the recovered point-wise maps, are limited as only low-frequency functions can be transferred.

In this paper we show how this limitation can be lifted using a recent spectral upscaling method [MMR*19], while still benefiting from the efficiency enabled by the functional map representation.

Basis Selection for Functional Maps. Our approach is also related to basis synchronization within the functional map framework. In the pairwise setting, this has been exploited for matching both complete [KBB*13, KBBV15, KGB16] and partial [LRBB17] shapes. These works aim to compute bases that both diagonalize the respective Laplacians and satisfy given (e.g., descriptor preservation) constraints. Such approaches can be particularly powerful for non-isometric shape pairs where the standard Laplace-Beltrami eigenfunctions are especially unstable. We also note a recent unpublished work in this domain [AL19] that proposes to simultaneously optimize for the basis and solve for a functional map between a pair of shapes in a joint problem using the Alternating Direction Method of Multipliers method.

Our work is different in that we jointly optimize for the bases across all shapes by computing a consistent latent basis in the collection. However, unlike other such approaches, e.g., [WHG13], crucially we also modify the size of the latent basis through progressive spectral upsampling [MMR*19], which allows us to capture higher frequency functional spaces and as we show, leads to accurate and consistent maps even in challenging scenarios.

3. Preliminaries

In this paper, we assume to be given a collection of 3D shapes $\mathcal{S} = \{S_i\}_{i=1}^n$ represented as triangle meshes. To each shape S_i we associate a Laplacian matrix \mathcal{L}_i discretized using the standard cotangent weight scheme [PP93, MDSB03], s.t. $\mathcal{L}_i = M_i^{-1}L_i$, where L_i is the cotangent weight matrix and M_i is the diagonal lumped area matrix. We denote by Λ_i the diagonal matrix of the first k eigenvalues of \mathcal{L}_i and by $\Phi_i^k = [\phi_i^1, \phi_i^2, \dots, \phi_i^k]$ the matrix of the corresponding eigenfunctions, where ϕ^l is the l^{th} eigenfunction (e.g., Fig. 2 shows the first few eigenfunctions of the airplanes in the top-middle). This

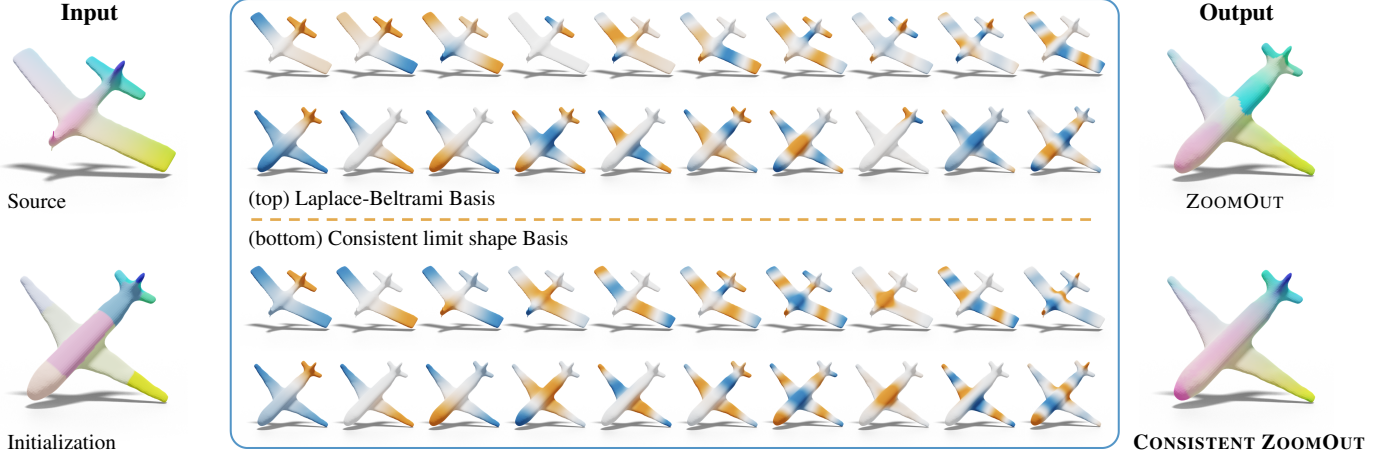


Figure 2: Left: we show two airplanes with different style. The patch initialization from 11 landmarks is visualized via color transfer. Middle: we show two sets of basis: the standard Laplace-Beltrami basis shown on the top, and our synchronized basis shown on the bottom. We can see that after the synchronization, the basis of the two shapes aligns much better to each other. Right: the refined maps using different basis sets visualized via color transfer. Our result (shown on the right bottom) is more accurate and smoother than the map refined by ZOOMOUT using the standard Laplace-Beltrami basis.

eigen-decomposition admits $L_i \Phi_i^k = M_i \Phi_i^k \Lambda_i$, and the eigenfunctions are orthonormal with respect to M_i : $(\Phi_i^k)^T M_i \Phi_i^k = I_k$, the identity matrix of dimension k .

Given two shapes S_1, S_2 and a point-wise map $T_{21} : S_2 \rightarrow S_1$, we denote by Π_{21} a binary matrix encoding T_{21} , such that, $\Pi_{21}(p, q) = 1$ if $T_{21}(p) = q$ and 0 otherwise. The two notations of point-wise maps are used interchangeably hereafter. Based on the spectral bases, functional maps can encode point-wise maps into compact matrices, namely:

$$C_{12} = (\Phi_2^k)^\dagger \Pi_{21} \Phi_1^k, \quad (1)$$

which gives rise to a matrix representation of size $k \times k$ (here \dagger indicates the pseudo inverse). A functional map transfers real-valued functions across shapes by translating the coefficients expressed in the eigenbasis. Namely, let \mathbf{f} be a coefficient vector representing a function on S_1 in basis Φ_1^k , then the corresponding function on S_2 is simply given as a coefficient vector $C_{12} \mathbf{f}$ in basis Φ_2^k .

On the other hand, for any vertex $p \in S_2$, the corresponding delta function δ_p is expressed in Φ_2^k as $\Phi_2^k(p) = [\phi_2^1(p), \phi_2^2(p), \dots, \phi_2^k(p)]$, i.e., the p^{th} row of Φ_2^k . From this point of view, Φ_2^k can be regarded as a spectral embedding of S_2 , which represents each p as a k -dimensional vector and the functional map C_{12} plays a role of aligning spectral embeddings across different shapes. Recovering the point-wise map from a given functional map can therefore be formulated as solving the following optimization problem:

$$T_{21}(p) = \arg \min_{q \in S_1} \|C_{12} \Phi_1^k(q)^T - \Phi_2^k(p)^T\|_2, \forall p \in S_2 \quad (2)$$

This recovery can be efficiently done through a nearest-neighbor query for each row of Φ_2^k in the space of rows of Φ_1^k , transformed by C_{12} .

Now, given a shape collection \mathcal{S} , a functional map network

(FMN) $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ represents a set of functional maps relating shapes in \mathcal{S} . The i^{th} node in \mathcal{V} corresponds to shape S_i , and $(i, j) \in \mathcal{E}$ if and only if S_i and S_j are connected by a functional map C_{ij} . The consistency of \mathcal{G} is then defined as the deviation of the composition of functional maps along a cycle in \mathcal{G} from the identity I – ideally, $C_{i_1, i_2} C_{i_2, i_3} \dots C_{i_k, i_1} = I, \forall \text{ cycle } \{i_1, i_2, \dots, i_k, i_1\} \in \mathcal{G}$. In [WHG13], the authors propose to extract a set of consistent latent bases Y_i for each S_i , such that $C_{ij} Y_i \approx Y_j, \forall (i, j) \in \mathcal{E}$. Note that Y_i can be both thought of as functions on S_i and as functional maps from some underlying latent object to S_i . From this point of view, C_{ij} can be factored out as $Y_j Y_i^{-1}$, i.e., a functional map composition from S_i to the latent shape, and then to S_j . This latent shape is formally analyzed in [HAGO19], where the authors show how to endow it with a geometric structure (measure and metric), resulting in a well-defined functional central shape, or so-called *limit shape*. As shown in [HAGO19], although the limit shape might not admit an actual embedding in the ambient space, it is fully characterized by a *canonical consistent latent basis*, which, again, can be treated as functional maps from the limit shape to each S_i .

4. Formulation

In this section, we first formulate our basic *multi-scale map synchronization* approach as a variational optimization problem, then we show how this formulation can be extended to more strongly promote near-isometries in shape collections. Our formulation is closely related to that of ZOOMOUT, but effectively alleviates the directional bias of the former, giving rise to the CONSISTENT ZOOMOUT pipeline described in Section 5.

Multi-scale shape collection synchronization We consider the following multi-scale optimization problem:

$$\min_{\mathcal{G} \in \mathcal{P}} E_{\text{col}}(\mathcal{G}), \quad (P_{\text{col}})$$

where

$$E_{\text{col}}(\mathcal{G}) = \sum_{\forall \{i_1, i_2, \dots, i_p\} \in \Gamma} \sum_k \frac{1}{k} \|C_{i_{p-1}i_p}^k \cdots C_{i_2i_1}^k C_{i_p i_1}^k - I^k\|_F^2. \quad (3)$$

Here C_{ij}^k is the functional map of dimension $k \times k$, and I^k is the identity matrix of the same dimension, and Γ is the set of cycles in \mathcal{G} . We denote by \mathcal{P} the set of functional maps that can be induced by some point-wise map (see Eq. (1)), and $\mathcal{G} \in \mathcal{P}$ implies all the functional maps in \mathcal{G} belong to \mathcal{P} .

Our formulation can be viewed as a natural extension to that of [WHG13] in two major ways. Firstly, in that work, the authors promote consistency for some specific fixed value of k . Instead, rather than using a single arbitrary fixed k our formulation is of multi-scale, encouraging *each* principal leading sub-block of the functional maps to be consistent. Secondly, the optimization of E_{col} is subject to $\mathcal{G} \in \mathcal{P}$, which introduces a very strong regularization. On the other hand, [WHG13] refines cycle-consistency purely on the spectral domain (via functional maps), and only converts to point-wise maps at the last step.

Near-isometric shape collection synchronization While the formulation in Eq. (1) already provides a strong regularization on the functional map, we further extend it to more strongly promote near-isometric maps. First, we recall that in [MRR*19], the authors consider that the following optimization problem for recovering the isometry between S_1, S_2 :

$$\min_{C_{12} \in \mathcal{P}} E_{\text{zm}}(C_{12}), \text{ where } E_{\text{zm}}(C_{12}) = \sum_k \frac{1}{k} \|(C_{12}^k)^T C_{12}^k - I^k\|_F^2, \quad (P_{\text{zm}})$$

As demonstrated in [MRR*19], the global optimizer of Problem (P_{zm}) with zero error corresponds to the underlying isometry between S_1, S_2 , and ZOOMOUT refinement is proven to solve it efficiently. At the same time, however, this formulation is based on a specific fixed direction between a source and target shape, which can bias the results and lead to non-bijective maps if functional maps in both directions are optimized independently.

Now we show how to remove this bias not only for a given shape pair, but more generally in the context of shape collections and functional map networks. For this we observe that since isometry is defined locally (pairwise), we can regularize each C_{ij} by $E_{\text{zm}}(C_{ij})$, and on the other hand inject global consistency by incorporating $E_{\text{col}}(\mathcal{G})$ as defined above. This leads to the following problem:

$$\min_{\mathcal{G} \in \mathcal{P}} E_{\text{col-iso}}(\mathcal{G}), \quad (P_{\text{col-iso}})$$

where

$$E_{\text{col-iso}}(\mathcal{G}) = \sum_{(i,j) \in \mathcal{E}} E_{\text{zm}}(C_{ij}) + E_{\text{zm}}(C_{ji}) + E_{\text{col}}(\mathcal{G}), \quad (4)$$

The following theorem (see proof in Appendix A) suggests that the global optimizer of $E_{\text{col-iso}}$ is stronger than that of E_{zm} , since it promotes maps that are both isometric and consistent.

Theorem 4.1 *Given a shape collection \mathcal{S} , in which all shapes*

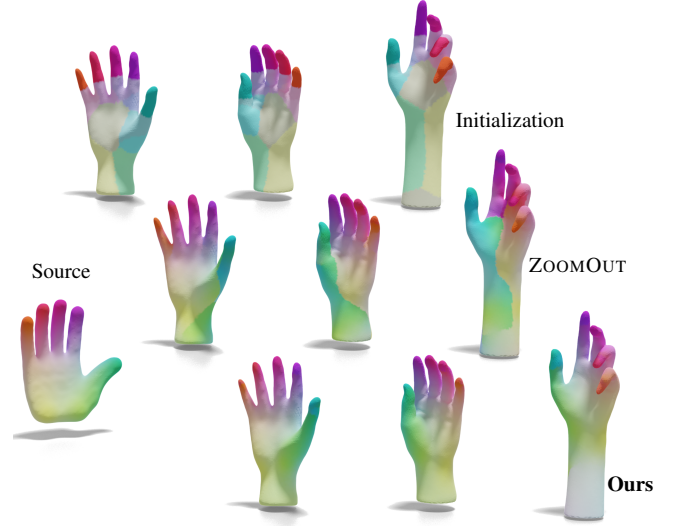


Figure 3: We apply our method to refine pairwise maps. Due to the large geometric difference between the target shapes and the source shape, ZOOMOUT produces noisy maps with evident discontinuities, while our results are more accurate and smoother.

share the same non-repeated eigenvalues of Laplacian matrices. A functional map network $\mathcal{G} \in \mathcal{P}$ satisfies $E_{\text{col-iso}}(\mathcal{G}) = 0$ if and only if the corresponding point-wise maps are all isometries, and are the consistent with respect to all cycles in \mathcal{G} , i.e., $\Pi_{i_1, i_p} \Pi_{i_p, i_{p-1}} \cdots \Pi_{i_2, i_1} = I, \forall \{i_1, i_2, \dots, i_p\}$ forming a cycle.

5. CONSISTENT ZOOMOUT Pipeline

In this section, we propose our CONSISTENT ZOOMOUT pipeline for synchronizing shape collections. In particular, we design two versions tailored for solving Problem ($P_{\text{col-iso}}$) (CONSISTENT ZOOMOUT-iso) and Problem (P_{col}) (CONSISTENT ZOOMOUT) approximately. Obviously the latter is a relaxed version of the former, so we start with the isometric case.

5.1. CONSISTENT ZOOMOUT-iso

At a high-level, we follow the same spectral upsampling scheme as ZOOMOUT – instead of optimizing the energy $E_{\text{col-iso}}$ across all scales, we first consider scale k , optimizing for the point-wise maps among the shape collection, and then use them as an initial guess for optimizing at a higher scale $k+1$.

However, optimizing Problem ($P_{\text{col-iso}}$) at scale k is still challenging since an ideal solution should simultaneously satisfy: 1) each $C_{ij}^k \in \mathcal{G}^k$ is near orthonormal; 2) \mathcal{G}^k admits cycle-consistency, where the high-order terms are heavily entangled; 3) each C_{ij}^k corresponds to some point-wise map. To address these challenges, we propose the following three steps for approximately optimizing Problem ($P_{\text{col-iso}}$).

Restricting the search space First, we relax the original optimization problem by restricting the search space to the following subspace:

$$\Omega = \{(C_{ij}^k, C_{ji}^k) \in \mathcal{P} \text{ and } C_{ij}^k = (C_{ji}^k)^T, \forall k \text{ and } (i, j) \in \mathcal{E}\}$$

ALGORITHM 1: CONSISTENT ZOOMOUT Iteration

- Input** : An FMN \mathcal{G}^k relating a shape collection \mathcal{S} , in which $C_{ij} \in \mathbb{R}^{k \times k}$.
- Output**: A refined FMN \mathcal{G}^{k+1}
- (1) For each $(i, j) \in \mathcal{E}$, if $\|C_{ij}^k - I^k\|_F \leq \|C_{ji}^k - I^k\|_F$, set $C_{ji}^k = (C_{ij}^k)^T$, otherwise set $C_{ij}^k = (C_{ji}^k)^T$
For CONSISTENT ZOOMOUT-iso only.
 - (2) Compute the weight matrix ω for \mathcal{G}^k via the consistency voting scheme from [NBCW*11]
 - (3) Construct a set of latent bases via Eq. (6) with ω , and then apply the canonicalization proposed in [HAGO19] to obtain the CCLB $\{Y_i^k\}$ (see Appendix C for details)
 - (4) Let $\Psi_i = \Phi_i^k Y_i^k, \forall i$, and recover the point-wise maps $T_{ij} : S_i \rightarrow S_j, \forall (i, j) \in \mathcal{G}$ via Eq. (5)
 - (5) Encode the computed point-wise maps T_{ij} via Eq. (1) using Φ_i^{k+1} . This results in a new FMN \mathcal{G}^{k+1} with $C_{ij} \in \mathbb{R}^{(k+1) \times (k+1)}$

Remark that in Ω , the orthonormality terms are equivalent to those of pairwise consistency $-\forall (i, j) \in \mathcal{E}, \|(C_{ij}^k)^T C_{ij}^k - I^k\|_F^2 = \|C_{ji}^k C_{ij}^k - I^k\|_F^2$.

Enforcing global consistency Secondly, we leverage the formulation of [HAGO19] to handle the global consistency, which extracts a set of *canonical consistent latent bases* (CCLB) $\{Y_i\}_{i=1}^n$ from \mathcal{G} . As formulated in [HAGO19], the CCLB can be seen as functional maps from a well-defined central (limit) shape that are consistent with respect to \mathcal{G} i.e., ideally $C_{ij} = Y_j Y_i^\dagger, \forall (i, j) \in \mathcal{E}$. The global consistency follows from $C_{i_{p-1}i_p} \cdots C_{i_1 i_2} C_{i_p i_1} = Y_{i_p} Y_{i_{p-1}}^\dagger Y_{i_{p-1}} \cdots Y_{i_2} Y_{i_1}^\dagger Y_{i_1} Y_{i_p}^\dagger = I, \forall \{i_1, i_2, \dots, i_p\} \in \Gamma_3$.

Solving for the point-wise maps On the one hand, the search space restriction turns the orthonormality terms into pairwise consistency terms; on the other hand, we drop all the high-order (i.e., ≥ 3 -cycle) consistency terms by incorporating the CCLB in the same spirit as [HG13]. Together, this leads to a simplified energy that consists of pairwise consistency terms along edges in \mathcal{E} , along with constraints introduced in the first two steps. For simplicity we consider a shape pair, arriving at the following optimization:

$$\begin{aligned} \min_{\Pi_{ij}, \Pi_{ji}} \quad & \|C_{ij}^k C_{ji}^k - I^k\|_F^2 + \|C_{ji}^k C_{ij}^k - I^k\|_F^2, & (P_{\text{single}}) \\ \text{s.t.} \quad & C_{ij}^k = (\Phi_j^k)^\dagger \Pi_{ji} \Phi_i^k, C_{ji}^k = (\Phi_i^k)^\dagger \Pi_{ij} \Phi_j^k, \\ & C_{ij}^k = Y_j^k (Y_i^k)^\dagger, C_{ji}^k = Y_i^k (Y_j^k)^\dagger, C_{ij}^k = (C_{ji}^k)^T. \end{aligned}$$

The two independent representations (via Eq. (1) and the CCLB decomposition) of C_{ij}^k in the constraint allow us to decouple Π_{ij} and Π_{ji} during optimization. In particular, we obtain the following theorem.

Theorem 5.1 *If $\mathcal{G}^k \in \Omega$ and admits CCLB decomposition: $C_{ij}^k = Y_i Y_j^\dagger, \forall (i, j) \in \mathcal{E}$, then under appropriate regularization, the opti-*

mal solution of Problem (P_{single}) is given by:

$$\begin{aligned} \Pi_{ij}^k &= \arg \min_{\Pi} \|\Pi \Phi_j^k Y_j^k - \Phi_i^k Y_i^k\|_F^2 \\ \Pi_{ji}^k &= \arg \min_{\Pi} \|\Pi \Phi_i^k Y_i^k - \Phi_j^k Y_j^k\|_F^2 \end{aligned}$$

Theorem 5.1 leads to the following simple procedure of computing Π_{ij} , which reduces to nearest neighbour searches of each row of $\Phi_i^k Y_i^k$ among the rows of $\Phi_j^k Y_j^k$.

$$T_{ij}(p) = \arg \min_{q \in S_j} \|\Psi_j(q)^T - \Psi_i(p)^T\|_2, \forall p \in S_i, \quad (5)$$

where $\Psi_i = \Phi_i^k Y_i^k$ and $\Psi_j = \Phi_j^k Y_j^k$. As we demonstrate in Section 5.4, this novel point-wise map conversion procedure leads to increased robustness while removing the directional bias, compared to the one used in [OCB*17, MRR*19] (c.f., Eq. (2)).

Overall pipeline Putting the above together, we propose our spectral upsampling scheme CONSISTENT ZOOMOUT-iso, which takes as input an FMN \mathcal{G}^k :

1. Restrict \mathcal{G}^k into Ω , such that $C_{ij}^k = (C_{ji}^k)^T, \forall (i, j) \in \mathcal{E}$.
2. Compute CCLB $\{Y_i\}$ of \mathcal{G}^k ;
3. Compute point-wise maps according to Eq. (5), and then convert to \mathcal{G}^{k+1} via Eq. (1).
4. Set $k = k + 1$ and repeat until k hits the given upper bound.

In Alg. 1, we summarize a single iteration of the practical implementation of CONSISTENT ZOOMOUT-iso, including the restriction Step (1) and the additional weighting that we use via the consistency voting [NBCW*11]. We refer reader to Section 5.3 for further details.

5.2. CONSISTENT ZOOMOUT

In fact solving Problem (P_{col}) is easier than the above one: since now the orthonormality terms are dropped, the only constraints are global consistency and compatibility between point-wise maps and functional maps. Following the same argument in Section 5.1, the algorithm of CONSISTENT ZOOMOUT is the same as Alg. 1, but skips the first step of search space restriction.

5.3. Implementation Details

Consistency voting for FMN Given a FMN \mathcal{G}^k , the consistency of the CCLB is enforced by the following optimization, proposed originally in in [WHG13]:

$$\begin{aligned} Y &= \arg \min_Y \sum_{(i,j) \in \mathcal{E}} \omega_{ij} \|C_{ij} Y_i - Y_j\|_F^2, \\ \text{s.t.} \quad & Y^T Y = n I^k, Y = [Y_1; Y_2; \dots; Y_n]. \end{aligned} \quad (6)$$

Here ω_{ij} is the weight assigned to each functional map C_{ij} , reflecting the relative confidence (the higher, the more confident). Often the initial maps vary widely in quality, especially in the presence of large deformations in the collection. It is thus natural to




Figure 4: Another example of applying CONSISTENT ZOOMOUT for Pairwise Shape Matching.

assign larger weights to the “good” maps, so that the extracted latent bases are less perturbed by the noisy maps. In this paper, following the approach proposed in [OBSCS*12], we adapt the consistency voting scheme introduced in [NBCW*11]. This scheme outputs a distance-like measure d_{ij} to each $C_{ij} \in \mathcal{G}^k$. We then obtain a weight matrix ω by letting $\omega_{ij} = \exp(-d_{ij}^2/2\sigma^2)$, where σ is the median of all d_{ij}^2 s. The consistency voting is conducted through all the iterations, since the FMN is consistently updated. On the other hand, this scheme is efficient since it only involves a linear optimization after light computation for consistency evaluation of functional maps along 3-cycles in \mathcal{G}^k .

Acceleration techniques First of all, we notice that our pipeline is highly parallelizable in the recovery of point-wise maps in each iteration, which gives a considerable boost in practice. Beyond that, several strategies are proposed in [MRR*19] to accelerate ZOOMOUT as well as improve its scalability. Our CONSISTENT ZOOMOUT benefits from the strategies immediately. Among those, the subsampling strategy is critical in our implementation. Namely, we only maintain the conversion between functional and point-wise maps among a set of subsampled vertices and only compute the dense point-wise maps in the final step as output. This strategy is especially useful in the scenarios where the shape collection is large, or the input shapes are of high complexity.

Network Topology In theory, our method can handle input \mathcal{G} with arbitrary topology as long as being connected. Though, it is observed in practice that a denser FMN is more robust to the noise present in initialization. By default, we only maintain the maps, either point-wise or functional ones, that are present in the initial \mathcal{G} through the refinement (thus the topology of \mathcal{G} is unchanged). Finally, via Eq. (5), we can output all the pairwise maps, regardless of it being present in the initialization or not.

5.4. CONSISTENT ZOOMOUT for Pairwise Shape Matching

While our framework is geared towards synchronizing maps in shape collections, it can naturally accommodate shape pairs as a special case. We have observed that the resulting method is more robust to challenging input than ZOOMOUT by promoting bijectivity and exploiting maps in both directions.

For instance, in Fig. 2, we consider refining maps between two airplanes with different styles from the SHREC’07 dataset [GBP07]. The rough patch-based initialization is shown in the bottom-left via color transfer. As shown in the top-middle, the Laplace-Beltrami bases of the shapes demonstrate distinctive modes, due to the geometric difference. As a consequence, ZOOMOUT fails to produce accurate and smooth maps (as shown in the top-right).

In contrast, CONSISTENT ZOOMOUT handles the non-isometric deformations more gracefully and produces a better map as shown in the bottom-right. This is thanks to two factors: 1) CONSISTENT ZOOMOUT is capable of refining maps in both directions simultaneously in a coupled manner; 2) the limit shape formulation exploited by CONSISTENT ZOOMOUT makes the corresponding spectral embeddings of the shapes more comparable – as shown in the bottom-middle, facilitating the point-wise map computation.

We demonstrate more examples in Fig. 3 and Fig. 4, where the initial maps, obtained via a set of sparse landmarks, are shown in the respective top row. Due to the large geometric difference between the target shapes and the source shape, ZOOMOUT produces noisy maps with evident discontinuities, while the results of CONSISTENT ZOOMOUT are consistently more accurate and smoother.

6. Results

In this section, we first describe different baseline methods and the datasets that we used for comparison. We then provide multiple quantitative and qualitative results in Section 6.1 and 6.2. In the end, we analyze and discuss the results in Section 6.3.

Baselines. In our experiments, we compare our method to different baselines, which can be categorized into two groups 1) computing maps in a collection of shapes, or 2) computing pairwise maps.

- **Collection:** we mainly compare to the following two methods that synchronize maps in a collection of shapes.
 - *Functional-Map Network (FN)* [HWG14]: is a state-of-the-art map synchronization framework that employs low-rank matrix recovery with the functional map representation.
 - *ICSM* [NBCW*11]: as briefly introduced in Section 5.3, ICSM is an approach based on consistency voting that refines a network of maps by replacing inconsistent maps with compositions of more consistent ones. We implement the version adapted to functional map framework as in [OBSCS*12].
- **Pairwise:** we also show some comparison to the pairwise map refinement techniques including:
 - *Iterative Closet Point (ICP)* [OBSCS*12]: this is the standard ICP registration applied in the spectral domain for map refinement.
 - *Product Manifold Filter (PMF)* [VLR*17] solves for a map based on a series of linear assignment formulations.
 - *Reversible Harmonic Maps (RHM)* [ESBC19] proposes a bi-directional geodesic Dirichlet energy to regularize the maps.

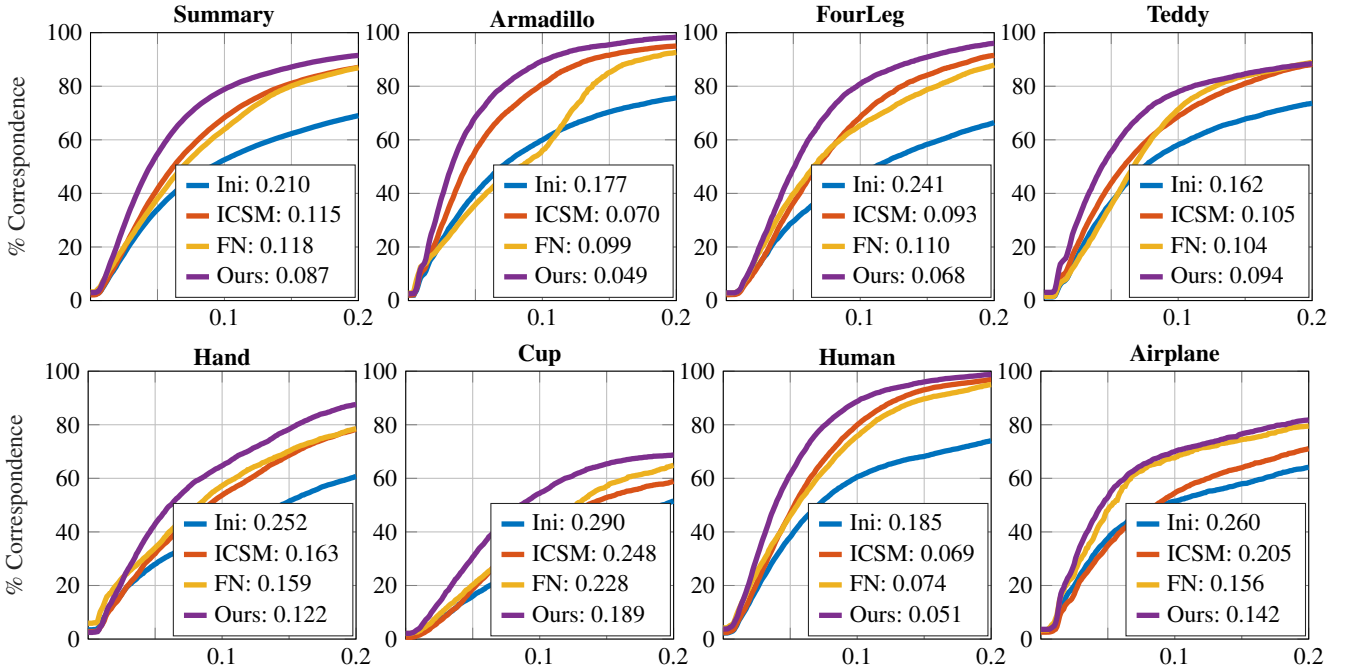


Figure 5: Accuracy evaluation of our method and baselines on 7 categories of the SHREC’07 dataset. We compare our result and that of FN [HWG14] and ICSM [NBCW*11]. The curves read the fraction (y-axis) of computed correspondences that fall within certain (normalized) geodesic distance to the ground-truth ones (x-axis). The numbers in the legends show the average error. Our method achieves the best results in all the categories and overall at least 26.2% improvement over the baselines.

- BCICP [RPWO18] uses some heuristics to improve the smoothness, bijectivity, and coverage of the maps.
- ZOOMOUT [MRR*19] applies an upsampling strategy in the spectral domain to refine the maps. This is the current state-of-the-art method in near-isometric shape matching.

Datasets. We use the following datasets from various sources to evaluate our method.

- **SHREC’07:** this dataset consists of 20 categories of shapes including humans, animals, and man-made objects. Each category has 20 shapes in different triangulation. Each shape within the same category has a sparse set of landmarks for accuracy evaluation. For each tested category, we evaluated all the pairwise maps.
- **Remeshed FAUST & SCAPE:** The FAUST [BRLB14] and the SCAPE [ASK*05] dataset consist of multiple human shapes that are near-isometric to each other with the same triangulation, which makes it easy to overfit for some refinement methods. Therefore, in our setting, we remeshed those human shapes independently such that different shapes have different triangulation and different number of vertices. For both datasets, we compare different methods on more than 100 shape pairs.
- **SHREC’19:** this dataset collected 44 different human shapes from 11 independent datasets with very different connectivity and mesh resolution. Each shape is aligned to the SMPL model [LMR*15] and thus dense ground-truth correspondences are given for accuracy evaluation. We tested on 430 shape pairs as in [MMR*19].

Metrics. In our tests, we mainly measure the accuracy w.r.t. the given ground-truth, the consistency error, and the runtime of the computed maps from different baselines.

- **Accuracy:** we measure the geodesic distance from the mapped vertex to the given ground-truth as the accuracy evaluation.
- **Bijectivity/Consistency:** with the computed maps T_{ij} from shape S_i to shape S_j , we take the compound map $T_{ji} \circ T_{ij}$ as a map from S_i to itself. We then measure the Euclidean distance between this compound map to the identity map as the consistency or bijectivity error.
- **Runtime:** we also evaluate the average runtime per shape pair for different baselines.

Initialization & Parameters. As our method is one of map synchronization, it requires a set of approximate initial maps as input. We provide the details on the exact procedure for initialization along with all of the parameters used in our implementation in Appendix D. We will release a complete implementation of our method upon acceptance.

6.1. Synchronizing Heterogeneous Shape Collections

We first apply our algorithm, CONSISTENT ZOOMOUT, on heterogeneous shape collections. In particular, we consider the SHREC’07 dataset and mainly compare to the shape collection matching methods including FN and ICSM. We tested on 7 categories of SHREC’07 including Armadillo, FourLeg, Teddy, Hand, Cup, Human, and Airplane. The cumulative matching error per category is reported in Fig. 5 where in the legend we report the average

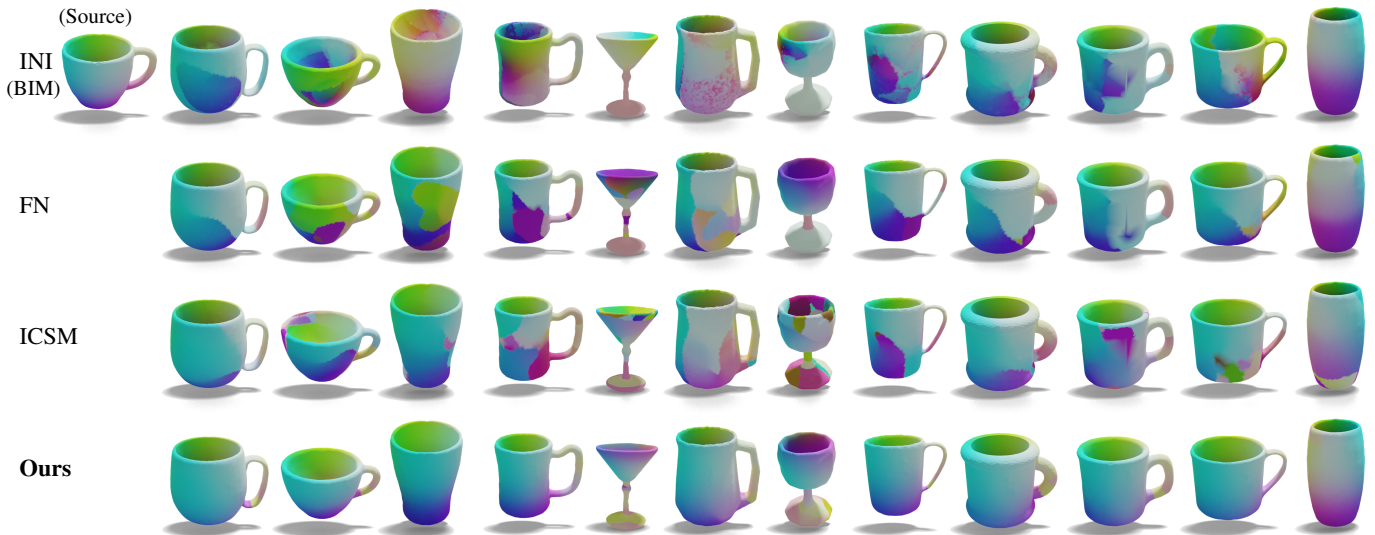


Figure 6: Here we show some qualitative results of the cup category in SHREC'07. The source shape is shown on the left top corner, we then show 12 maps on the target cups via color transfer. We can see that our results are much more smooth and accurate than the other two baselines. This improvement is not fully illustrated in Fig. 5 due to the limited given ground-truth landmarks for accuracy evaluation.

error across 380 shape pairs for each category and each method. We also include a summary of all the tested categories in Fig. 5.

Specifically, we achieved 26.2% improvement in matching accuracy w.r.t. the best baseline averaged over 7 categories. In Fig. 6 we show a qualitative comparison of the Cup category. We can see that even though the accuracy curve in Fig. 5 does not reveal a significant improvement of our method over the other two methods, our results are actually more accurate, smoother, and more consistent. Finally, regarding runtime, our method on average takes 0.59s per map in the refinement, while *FN* takes 0.50s and *ICSM* takes 0.79s on average. The statistics are collected over 7 categories, including 2433 maps.

Remark that we also achieved 84.3% improvement in map consistency w.r.t. the best baseline. We refer readers to Appendix E for more detailed results, as well as some qualitative comparison on map consistency.

6.2. Synchronizing Near-isometric Shape Collections

In this part, we test our CONSISTENT ZOOMOUT-iso for synchronizing near-isometric shape collections including **SCAPE**, **FAUST** and **SHREC'19**. Besides the baselines used in Section 6.1, we introduce three more: our CONSISTENT ZOOMOUT and two hybrid baselines, which apply ZOOMOUT refinement on top of the maps synchronized by *FN* and by *ICSM*.

We report the results on the three different datasets in Table 1, and in Fig. 7 we demonstrate a qualitative comparison on map consistency. Overall, we conclude that:

- Our CONSISTENT ZOOMOUT-iso achieves the best accuracy across all tests, and the best consistency in two out of three.
- As shown in the initialization row of Table 1, the initialization quality of **FAUST (sparse)** test is relatively lower than the other two. In this case, we can see that most of the pairwise refinement baselines struggle to produce good results, except for

BCICP. On the other hand, by leveraging the consistency regularization across the *whole* collection, CONSISTENT ZOOMOUT achieves already slightly better result than *BCICP*, while being about 500× faster.

- Moreover, our CONSISTENT ZOOMOUT-iso improves the performance upon CONSISTENT ZOOMOUT by 16%, which suggests the effectiveness of our specified formulation for near-isometric collections.
- We observe that our method outperforms the baselines for collection synchronization, *FN*, *ICSM*, by a large margin across all tests. That is still true even after their outputs are further refined by ZOOMOUT.
- Unlike *BCICP*, *PMF*, our method does not enforce the point-wise map consistency explicitly, while achieving better or comparable performance in the bijectivity evaluation. It is also worth noting that our method indeed promotes *global* consistency among shape collections, instead of merely in map bijectivity.

6.3. Discussion and Analysis

- Our results highlight that conducting map refinement in the context of shape collections is beneficial. This is especially true when the input maps are noisy (e.g., the test on **FAUST** in Section 6.2), or the shapes undergo significant deformations (see, e.g., Fig. 1).
- Similar to our method, the two baselines on map synchronization are both designed to promote the cycle-consistency of functional maps among shape collections. As mentioned in Section 4, the strength of our method comes mainly from the fact that 1) we take a multi-scale approach, while the baselines process the functional maps at a fixed scale all the time; 2) the baselines promote only the consistency of *functional* maps, which not necessarily corresponds to high-quality point-wise maps. On the other hand, our method relates to both map representations, leading to much more accurate and consistent results.
- It is worth noting, however, our method will be affected if the

Table 1: Near-isometric Datasets Summary. We compare with the pairwise refinement techniques ICP, PMF, BCICP, and ZOOMOUT, the collection refinement methods including FN, ICSM, and hybrid baselines combining both types of methods on three different datasets. For PMF and BCICP, we pick the sampling size (1k and 5k respectively) with the best accuracy shown in [MRR*19] as our baselines. The numbers in the table indicate the mean normalized geodesic distance to the ground-truth for the accuracy evaluation. As for the bijectivity evaluation, we use the normalized Euclidean distance instead.

Methods\Metrics	SCAPE (dense)		FAUST (sparse)		SHREC'19		Runtime (s)
	Acc. ($\times 10^{-3}$)	Bij. ($\times 10^{-3}$)	Acc. ($\times 10^{-3}$)	Bij. ($\times 10^{-3}$)	Acc. ($\times 10^{-3}$)	Bij. ($\times 10^{-3}$)	
Initialization	50.0	35.3	99.2	108.8	60.4	31.4	-
ICP	32.6	20.5	58.6	50.0	47.0	18.6	87.3
RHM	33.6	13.7	56.4	25.5	51.8	4.0	118.1
PMF (1k)	98.0	0	101.5	0	30.1	4.3	437.9
BCICP (5k)	29.7	8.8	34.0	8.7	42.6	4.5	2313
ZOOMOUT	27.0	13.8	45.9	36.8	28.8	8.1	1.5
FN	38.1	33.4	114.5	122.3	60.9	24.2	3.8
ICSM	33.4	21.2	71.2	70.3	42.9	18.5	10.0
FN + ZOOMOUT	27.2	18.9	36.5	28.1	28.9	7.1	6.7
ICSM + ZOOMOUT	27.2	19.2	45.9	38.9	28.5	7.7	13.3
Ours	22.6	6.6	33.4	8.7	28.9	3.7	4.5
Ours-iso	22.5	5.3	28.0	6.8	27.0	4.5	4.5




Figure 7: We demonstrate the bijectivity of the resulting maps from various baselines and our method. In the top row, we compare the texture transferred from S_1 to itself via the map composition going through S_2 in the second row, i.e. $T_{21} \circ T_{12}$ (and similarly for S_2 in the bottom row). Under perfect bijectivity, the map composition should be an identity map. It is then evident that our method significantly outperforms the collection map refinement baselines and even achieves comparable results with the frameworks aiming to optimize for bijectivity in the pairwise setting. Compared to them, our method is far more efficient (see Table 1) and without any specific pairwise refinement.

given initial maps are heavily inaccurate and inconsistent, especially in the presence of symmetries. Fig. 8 shows some qualitative results of the Teddy class from SHREC'07 dataset (see the quantitative result in Fig. 5). When the initial maps are partially symmetrical or even wrong, our method fails to correct all maps but instead returned twisted maps. We also refer readers to Appendix. F for the results of the more challenging Ants class.

7. Conclusion and limitations

In this paper, we proposed an efficient map synchronization method, CONSISTENT ZOOMOUT. Our method is based on a multi-scale functional map consistency energy and results in a simple iterative upsampling scheme in practice. Our method results in high-accuracy maps, compared to various often sophisticated baselines in different settings, but also significantly improves the map consistency, without any additional post-processing.

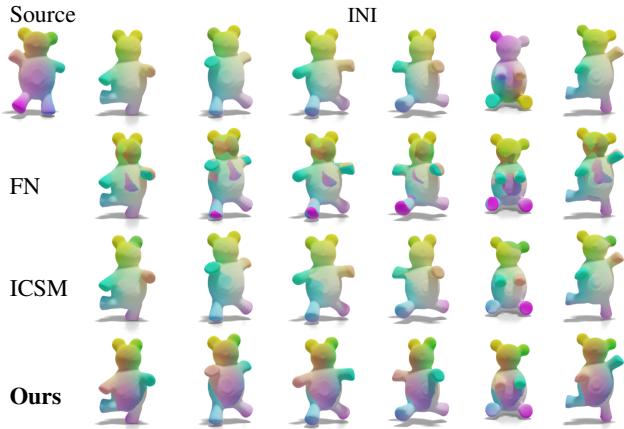


Figure 8: Qualitative results of the Teddy shapes from the SHREC'07 dataset, where the corresponding quantitative evaluations are shown in Fig. 5. In this example, the initial BIM maps are not accurate enough. For example, some shapes are mapped from left to right, back to front, or even foot to head. Starting from this challenging initialization, our method returns twisted maps, where the bodies are correctly mapped but the legs are consistently mapped from left to right. On the other hand, the results of the baselines are less consistent (see, for example, how the ears are mapped in ICSM) than ours, while either being noisy (FN) or suffering from multiple symmetries (ICSM).



Figure 9: CONSISTENT ZOOMOUT can utilize the collection information to refine the initial maps with large outliers. In this example, the initial maps computed on the armadillo pairs involve left-to-right ambiguity, partial shape distortions, and large errors. Our method can avoid those errors by enforcing the circle consistency and produce high-quality refined maps.

Our method still possesses a few limitations. First, as mentioned in Section 6.3, our method's performance depends on the quality of the initial maps. Also, since the canonical consistent latent basis formulation assumes implicitly that bijective maps exist across the shape collection, our method cannot provide theoretical guarantees on processing collections of partial shapes (though, interestingly, as shown in Fig. 9, our method works well with armadillos having missing parts). In the future, we also plan on improving scalability by better exploiting its potential for parallelization.

Acknowledgement The authors thank the anonymous reviewers for their valuable comments. Parts of this work were supported by the KAUST OSR Award No. CRG-2017-3426 and the ERC Starting Grant No. 758800 (EXPROTEA).

References

- [AL19] AZENCOT O., LAI R.: Shape analysis via functional map construction and bases pursuit, 2019. 2
- [APL15] AIGERMAN N., PORANNE R., LIPMAN Y.: Seamless surface mappings. *ACM Trans. Graph.* 34, 4 (July 2015), 72:1–72:13. 1
- [ARFF18] ARRIGONI F., ROSSI B., FRAGNETO P., FUSIELLO A.: Robust synchronization in so (3) and se (3) via low-rank and sparse matrix decomposition. *Computer Vision and Image Understanding* 174 (2018), 95–113. 2
- [ASK*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: SCAPE: Shape Completion and Animation of People. In *ACM Transactions on Graphics (TOG)* (2005), vol. 24, ACM, pp. 408–416. 1, 7
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized Multidimensional Scaling: A Framework for Isometry-Invariant Partial Surface Matching. *PNAS* 103, 5 (2006), 1168–1172. 1
- [BCBB16] BIASOTTI S., CERRI A., BRONSTEIN A., BRONSTEIN M.: Recent trends, applications, and perspectives in 3d shape similarity assessment. *Computer Graphics Forum* 35, 6 (2016), 87–119. 2
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and Evaluation for 3d Mesh Registration. In *Proc. CVPR* (2014), pp. 3794–3801. 7
- [BTGT19] BERNARD F., THUNBERG J., GONCALVES J., THEOBALT C.: Synchronisation of partial multi-matchings via non-negative factorisations. *Pattern Recognition* 92 (2019), 146–155. 2
- [CMG13] CHATTERJEE A., MADHAV GOVINDU V.: Efficient and robust large-scale rotation averaging. In *Proceedings of the IEEE International Conference on Computer Vision* (2013), pp. 521–528. 2
- [CRA*17] COSMO L., RODOLÀ E., ALBARELLI A., MÈMOLI F., CREMERS D.: Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum* 36, 1 (2017), 209–221. 1, 2
- [ERGB16] EYNARD D., RODOLA E., GLASHOFF K., BRONSTEIN M. M.: Coupled functional maps. In *Fourth International Conference on 3D Vision (3DV)*, 2016 (2016), IEEE, pp. 399–407. 1
- [ESBC19] EZUZ D., SOLOMON J., BEN-CHEN M.: Reversible harmonic maps between discrete surfaces. *ACM Trans. Graph.* 38, 2 (Mar. 2019), 15:1–15:12. 1, 6
- [GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: Shape retrieval contest 2007: Watertight models track. *SHREC competition* (2007). 2, 6
- [HAGO19] HUANG R., ACHLIOPTAS P., GUIBAS L., OVSJANIKOV M.: Limit shapes – a tool for understanding shape differences and variability in 3d model collections. *Computer Graphics Forum* 38, 5 (2019), 187–202. 2, 3, 5, 12
- [HG13] HUANG Q.-X., GUIBAS L.: Consistent shape maps via semidefinite programming. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 177–186. 1, 2, 5
- [HLBH17] HUANG X., LIANG Z., BAJAJ C., HUANG Q.: Translation synchronization via truncated least squares. In *Advances in neural information processing systems* (2017), pp. 1459–1468. 2
- [HLW*19] HUANG Q., LIANG Z., WANG H., ZUO S., BAJAJ C.: Tensor maps for synchronizing heterogeneous shape collections. *ACM Trans. Graph.* 38, 4 (July 2019), 106:1–106:18. 1
- [HO17] HUANG R., OVSJANIKOV M.: Adjoint map representation for shape analysis and matching. In *Proc. SGP* (2017), vol. 36. 1
- [HSS*09] HASLER N., STOLL C., SUNKEL M., ROSENHAHN B., SEIDEL H.-P.: A Statistical Model of Human Pose and Body Shape. In *Computer Graphics Forum* (2009), vol. 28, pp. 337–346. 1
- [Hub02] HUBER D. F.: *Automatic three-dimensional modeling from reality*. PhD thesis, Citeseer, 2002. 2
- [HW14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 36. 2, 6, 7

- [KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 439–448. 2
- [KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015), pp. 905–914. 2
- [KGB16] KOVNATSKY A., GLASHOFF K., BRONSTEIN M. M.: MADMM: A generic algorithm for non-smooth optimization on manifolds. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (2016), vol. 9909 LNCS, pp. 680–696. 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended Intrinsic Maps. In *ACM Transactions on Graphics (TOG)* (2011), vol. 30, ACM, p. 79. 12, 13
- [LMR*15] LOPER M., MAHMOOD N., ROMERO J., PONS-MOLL G., BLACK M. J.: SMPL: A skinned multi-person linear model. *TOG* 34, 6 (2015), 248:1–248:16. 7
- [LRBB17] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Fully spectral partial shape matching. *Computer Graphics Forum* 36, 2 (2017), 247–258. 2
- [LZD17] LEONARDOS S., ZHOU X., DANIILIDIS K.: Distributed consistent data association via permutation synchronization. In *2017 IEEE International Conference on Robotics and Automation (ICRA)* (2017), IEEE, pp. 2645–2652. 2
- [MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In *Visualization and mathematics III*. Springer, 2003, pp. 35–57. 2
- [MMR*19] MELZI S., MARIN R., RODOLÀ E., CASTELLANI U., REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Matching humans with different connectivity. In *3DOR* (2019), pp. 121–128. 2, 7, 12
- [MRR*19] MELZI S., REN J., RODOLÀ E., WONKA P., OVSJANIKOV M.: Zoomout: Spectral upsampling for efficient shape correspondence. *Proc. SIGGRAPH Asia* (2019). 2, 4, 5, 6, 7, 9, 11, 12
- [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. 1, 2, 5, 6, 7, 12
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional Maps: A Flexible Representation of Maps Between Shapes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 30. 1, 6
- [OCB*17] OVSJANIKOV M., CORMAN E., BRONSTEIN M., RODOLÀ E., BEN-CHEN M., GUIBAS L., CHAZAL F., BRONSTEIN A.: Computing and processing correspondences with functional maps. In *ACM SIGGRAPH 2017 Courses* (2017), ACM, p. 5. 2, 5
- [PKS13] PACHAURI D., KONDOR R., SINGH V.: Solving the multi-way matching problem by permutation synchronization. In *Advances in neural information processing systems* (2013), pp. 1860–1868. 2
- [PP93] PINKALL U., POLTHIER K.: Computing Discrete Minimal Surfaces and their Conjugates. *Experimental mathematics* 2, 1 (1993), 15–36. 2
- [RPPSH17] RABINOVICH M., PORANNE R., PANOZZO D., SORKINE-HORNUNG O.: Scalable locally injective mappings. *ACM Trans. Graph.* 36, 2 (Apr. 2017). 1
- [RPWO18] REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.* 37, 6 (Dec. 2018), 248:1–248:16. 1, 7, 12
- [RSO19] ROUFOSSE J., SHARMA A., OVSJANIKOV M.: Unsupervised deep learning for structured shape matching. *ICCV* (2019). 12
- [SHS16] SHEN Y., HUANG Q., SREBRO N., SANGHAVI S.: Normalized spectral map synchronization. In *Advances in Neural Information Processing Systems* (2016), pp. 4925–4933. 2
- [SVBC19] SHOHAM M., VAXMAN A., BEN-CHEN M.: Hierarchical functional maps between subdivision surfaces. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 55–73. 2
- [TCL*13] TAM G. K., CHENG Z.-Q., LAI Y.-K., LANGBEIN F. C., LIU Y., MARSHALL D., MARTIN R. R., SUN X.-F., ROSIN P. L.: Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. *IEEE TVCG* 19, 7 (2013), 1199–1217. 2
- [VKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 2
- [VLR*17] VESTNER M., LITMAN R., RODOLÀ E., BRONSTEIN A., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proc. CVPR* (2017), pp. 6681–6690. 6
- [WHG13] WANG F., HUANG Q., GUIBAS L. J.: Image co-segmentation via consistent functional maps. In *Proceedings of the IEEE International Conference on Computer Vision* (2013), pp. 849–856. 1, 2, 3, 4, 5, 12
- [WS13] WANG L., SINGER A.: Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA* 2, 2 (2013), 145–193. 2
- [ZKP10] ZACH C., KLOPSCHITZ M., POLLEFEYS M.: Disambiguating visual relations using loop constraints. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2010), IEEE, pp. 1426–1433. 2
- [ZZD15] ZHOU X., ZHU M., DANIILIDIS K.: Multi-image matching via fast alternating minimization. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 4032–4040. 2

Appendix A: Proof of Theorem 4.1

Proof We let Π_{ji}, Π_{ij} be respectively the point-wise maps corresponding to C_{ij}, C_{ji} .

First, if $E_{\text{col-iso}}(\mathcal{G}) = 0$, then $\forall (i, j) \in \mathcal{E}, E_{\text{zm}}(C_{ij}) = 0$ and $E_{\text{zm}}(C_{ji}) = 0$. Following Theorem 4.1 of [MRR*19], Π_{ji}, Π_{ij} are isometries between S_i, S_j . Regarding the consistency argument, in the complete cotangent Laplacian basis, we have

$$I = C_{i_p, i_1} C_{i_{p-1}, i_p} \cdots C_{i_1, i_2} = \Pi_{i_1, i_p} \Pi_{i_p, i_{p-1}} \cdots \Pi_{i_2, i_1},$$

for any cycle $\{i_1, i_2, \dots, i_p\} \in \mathcal{G}$, which proves that $E_{\text{col-iso}}(\mathcal{G}) = 0$ implies cycle consistency as well.

Now we prove the converse, i.e., if Π_{ij}, Π_{ji} are isometries and $\Pi_{i_1, i_p} \Pi_{i_p, i_{p-1}} \cdots \Pi_{i_2, i_1} = I$ holds for any cycle in \mathcal{G} , then $E_{\text{col-iso}}(\mathcal{G}) = 0$.

Again, according to Theorem 4.1 of [MRR*19], since Π_{ij}, Π_{ji} are isometries, then

$$\sum_{(i,j) \in \mathcal{E}} E_{\text{zm}}(C_{ij}) + E_{\text{zm}}(C_{ji}) = 0. \quad (7)$$

Now we prove the consistency terms also vanish. Following Theorem 4.1 of [MRR*19], C_{ij}^k, C_{ji}^k are both diagonal and orthonormal matrices, which implies that the non-zero entries of them are either 1 or -1 . On top of that, by the definition of functional maps (Eq. (1)), we have

$$\Pi_{ij} \Phi_j^k = \Phi_i^k C_{ji}^k. \quad (8)$$

Without loss of generality, we assume that $\{1, 2, 3\}$ forms a cycle

in \mathcal{G} , then according to Eq. (8):

$$\begin{aligned}\Phi_1^k C_{21}^k C_{32}^k C_{13}^k &= \Pi_{12} \Phi_2^k C_{32}^k C_{13}^k \\ &= \Pi_{12} \Pi_{23} \Phi_3^k C_{13}^k \\ &= \Pi_{12} \Pi_{23} \Pi_{31} \Phi_1^k \\ &= \Phi_1^k,\end{aligned}$$

which implies that $C_{21}^k C_{32}^k C_{13}^k = I^k$, and similarly we can prove all the cycle-consistency terms vanish, i.e., $E_{\text{col-iso}}(\mathcal{G}) = 0$. \square

Appendix B: Proof of Theorem 5.1

Proof We prove the follow equality:

$$\begin{aligned}\|\Pi_{ji} \Phi_i^k Y_i^k - \Phi_j^k Y_j^k\|_A^2 &= \|C_{ij}^k C_{ji}^k - I^k\|_F^2 \\ &\quad + \|(I - \Phi_j \Phi_j^\dagger)(\Pi_{ji} \Phi_i Y_i - \Phi_j Y_j)\|_A^2,\end{aligned}\quad (9)$$

where A is the mass matrix of S_1 .

According to the constraints and Eq. (1), we have:

$$\|C_{ij}^k C_{ji}^k - I^k\|_F^2 = \|(\Phi_j^k)^\dagger \Pi_{ji} \Phi_i^k Y_i^k (Y_j^k)^\dagger - I^k\|_F^2.$$

(1) We first prove that

$$\|(\Phi_j^k)^\dagger \Pi_{ji} \Phi_i^k Y_i^k (Y_j^k)^\dagger - I^k\|_F^2 = \|(\Phi_j^k)^\dagger \Pi_{ji} \Phi_i^k Y_i^k - Y_j^k\|_F^2\quad (10)$$

According to the constraint, $C_{ij}^k = (C_{ji}^k)^T$, then according to the construction of the CLB, we have

$$\begin{aligned}Y_j Y_i^\dagger &= (Y_i Y_j^\dagger)^T \\ \Rightarrow Y_j Y_i^\dagger &= (Y_j^\dagger)^T Y_i^T \\ \Rightarrow (Y_j^T Y_i)^\dagger &= (Y_i^T Y_j) = I\end{aligned}\quad (11)$$

Note that Eq. (11) holds for any i, j , and that $\sum_i Y_i^T Y_i = nI$ by the construction of CCLB, one can conclude that $Y_i^T Y_i = I, i = 1, \dots, n$, which implies Eq. (10).

(2) Once Eq. (10) is proven, we borrow the following equation proven in [MRR*19]:

$$\|X\|_A^2 = \|B^\dagger X\|_F^2 + \|(I - BB^\dagger)X\|_A^2,$$

where $B^T AB = I$, i.e., $B^\dagger = B^T A$. We then let $X = \Pi_{ji} \Phi_i^k Y_i^k - \Phi_j^k Y_j^k$ and $B = \Phi_j^k$, finishing the proof. \square

Appendix C: Computing CCLB

We give a brief overview of the details for computing CCLB, which is used in step 3 of Alg. 1. Given an FMN \mathcal{G} and the corresponding weight matrix ω computed via the consistency voting [NBCW*11], we follow the formulation of [WHG13] and construct the following block-wise matrix W of size $nk \times nk$:

$$W_{i,j} = \begin{cases} \sum_{(i,j') \in \mathcal{G}} \omega_{ij'} (I + C_{ij'}^T C_{ij'}) & \text{if } i = j, \\ -(\omega_{ji} C_{ji} + \omega_{ij} C_{ij}^T) & \text{otherwise.} \end{cases}\quad (12)$$

We then perform an eigen-decomposition on W to compute the

smallest k eigenvectors and stack them into a $nk \times k$ matrix Y . Letting $Y = [Y_1; Y_2; \dots; Y_n]$, as proven in [WHG13], this is the optimal solution for Eq. (6).

Now, in order to obtain the CCLB formulated in [HAGO19], we let $E = \sum_i Y_i^T \Lambda_i Y_i$ and do eigen-decomposition on E , yielding $EU = U\Lambda$. Finally, the CCLB is obtained by letting $\tilde{Y}_i = Y_i U$. Note that, for simplicity, we denote by Y_i the CCLB in our main paper directly.

Appendix D: Initialization & Parameters

Initialization.

In our tests, we used the same initialization for different baseline methods.

- **SHREC'07**: for the tests involving this dataset, we always use the maps computed using BIM [KLF11] as initialization.
- **SCAPE (dense)**: we pick 18 shapes from this dataset, and use a learning-based framework, SURFMNet [RSO19], to compute the densely pairwise maps as initialization.
- **FAUST (sparse)**: we use the provided maps in [RPWO18], which contains 300 pairs of maps among the whole 100 shapes.
- **SHREC'19**: we use the provided maps in [MMR*19] as initialization, which are computed from 5 given landmarks.

Parameters. Throughout all the experiments, we use the same set of parameters except for the way of augmenting dimension of functional maps: For computing the canonical consistent latent basis, we consistently use the first $\lfloor 0.9k \rfloor$ canonical consistent latent basis, where k is the dimension of functional maps at current iteration. We use Euclidean farthest point sampling starting from a random seed point, and always sample 1000 vertices from each shape in the collection through all experiments. The intermediate refinement iterations are maintained only among the sampled points, and dense point-wise maps are only computed at the output stage.

In Section 6.1, we augment the dimension of functional maps from 30×30 to 80×80 with step size of 2, while in Section 6.2, we augment from 20×20 to 100×100 with step size of 5.

Appendix E: Consistency Evaluation of Results on SHREC'07

We plot in Fig. 10 the accumulative consistency error curves for each of the 7 categories we evaluated. It is evident that the bijectivity of our results outperform that of the baselines by a large margin.

Beyond the quantitative evaluation above, we also select the category of Hands and compare the cycle-consistency of the maps from different methods qualitatively in Fig. 12. Specifically, for a shape collection $\{S_i\}$ with the computed maps $T_{k-1,k}$ from S_{k-1} to S_k , we compute the compound map $T_{11}^{(k)}$ in a circle of k shapes by $T_{11}^{(k)} = T_{k1} \circ T_{k-1,k} \circ \dots \circ T_{23} \circ T_{12}$, which is a map from S_1 to itself (see the top-left for an example). We then visualize the maps $T_{11}^{(k)}$ with different k , ranging from 2 shapes to 20 shapes (the whole shape collection). we compare the cycle consistency of the our maps to those obtained from other baselines. We can see that,

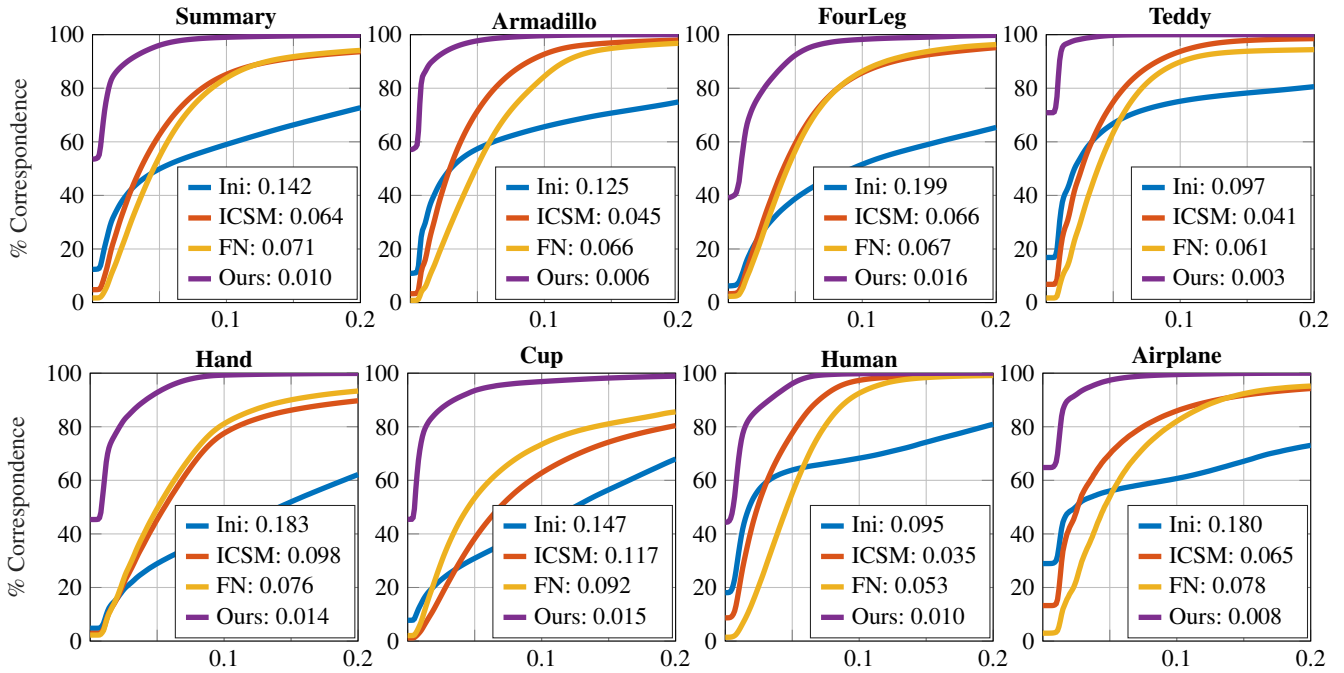


Figure 10: Bijection evaluation on 7 categories of SHREC'07. We accumulate the bijection error as the deviation from $T_{ji} \circ T_{ij}$ to T_{ii} – the identity map on S_i – for all pairs of (i, j) . Note that the curves read in the same way as Fig. 5 and the errors are computed in Euclidean space. Our method significantly outperforms the competing baselines by a large margin: overall 84.3% percent improvement over the best baseline.

even though the input maps are highly inconsistent, our method can achieve consistent maps in a shape collection. As a comparison, even though the baseline methods *FN* and *ICSM* use circle consistency to regularize the map quality, the final maps are not in a consistent manner. Moreover, when compare to the pair-wise refinement baseline *ZOOMOUT*, we can see that when the initial maps are inaccurate and inconsistent, the pair-wise refinement method can barely work, while the collection-based refinement method can still work.

Appendix F: Failure Case due to Symmetry

Fig. 11 shows some example ant pairs, where we can see that the initial maps have mixed symmetries and wrong matches between the legs and the antennas. In this case, our method can produce more accurate maps. However, the symmetry is not fully corrected.




Figure 11: Here we show an example of our failure case on the Ant category of SHREC'07 dataset. The initial maps are computed from BIM [KLF11], where there are a lot of symmetry ambiguities and mismatching between the legs and the antennas. In this case, our method can fail to output correct maps directly.



Figure 12: Qualitative comparison of cycle-consistency of various map syntonization/refinement results (see text for details).