

Joint Map and Symmetry Synchronization

*Yifan Sun¹, *Zhenxiao Liang^{1,2}, *Xiangru Huang¹, and Qixing Huang¹

¹ The University of Texas at Austin

² Tsinghua University

{yifan,liangzx,xrhuang,huangqx}@cs.utexas.edu

Abstract. Most existing techniques in map computation (e.g., in the form of feature or dense correspondences) assume that the underlying map between an object pair is unique. This assumption, however, easily breaks when visual objects possess self-symmetries. In this paper, we study the problem of jointly optimizing symmetry groups and pair-wise maps among a collection of symmetric objects. We introduce a lifting map representation for encoding both symmetry groups and maps between symmetry groups. Based on this representation, we introduce a computational framework for joint symmetry and map synchronization. Experimental results show that this approach outperforms state-of-the-art approaches for symmetry detection from a single object as well as joint map optimization among an object collection.

Keywords: Correspondences, Symmetry group, Cycle-consistency, Optimization

1 Introduction

Establishing maps across visual objects is a long standing problem in visual computing with rich applications in structure-from-motion [3, 9], joint segmentation [48, 49], and label/attribute transfer [2, 31, 50], among others. While early works focus on computing maps between pairs of objects in isolation (c.f. [21]), a recent trend is to jointly compute consistent maps across a collection of objects [19, 16, 17, 42, 20, 24, 37, 38, 5, 66, 67, 51, 29]. This is motivated from the fact that (i) many applications (e.g., multi-view structure-from-motion [54, 8] and co-segmentation [60, 56]) require consistent correspondences, and (ii) optimizing consistent maps enables us to improve maps computed between pairs of objects in isolation (c.f. [16]). The promise of these approaches is that pairwise maps between dissimilar objects can often be improved by composing maps along paths of similar objects.

However, all existing joint map optimization techniques place the assumption that there is only one underlying map between two visual objects. It turns out this assumption does not hold in many scenarios. In particular, when the underlying objects are symmetric, there are multiple plausible maps between each object pair. This introduces issues both for map computation between pairs of objects and for optimizing consistent maps across multiple objects. To address

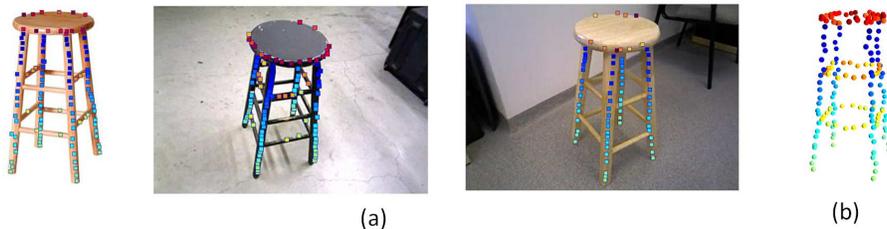


Fig. 1. Category specific reconstruction from internet images. (a) Subset of input images with matched edge feature points [12]. (b) Reconstruction using our approach. Color-coding highlights the recovered symmetry group.

this issue, a straightforward approach is to detect the underlying self-symmetry group of each object first, and then perform matching in factored or quotient spaces (c.f.[39]). This approach, however, is not optimal, since symmetry detection is quite challenging (especially for images where the symmetric parts may be occluded), and errors introduced in the symmetry detection phase cannot be recovered in the joint matching phase.

In this paper, we consider the problem of jointly optimizing the underlying symmetry groups and consistent pairwise maps. We call this problem joint symmetry and map synchronization. The motivation comes from the fact that symmetry detection and map computation are highly correlated. Besides the fact that computing the underlying symmetry groups facilitates map computation, maps can also propagate symmetry groups across objects, allowing us to rectify a noisy symmetry group computed from one object by propagating clean symmetry groups computed from other objects.

To formulate joint symmetry and map synchronization, we propose a representation that encodes symmetry groups and pair-wise maps using the product operator (also known as lifting in the literature [27][18][23, 58]). We prove that for many symmetry groups in our physical world (e.g., reflection symmetry and n -fold rotational symmetry among others), this representation admits exact symmetry group decoding. In addition, we show that the computational cost incurred by the product operator can be effectively addressed via reduced functional basis [38]. Using this representation, we introduce a continuous optimization formulation for joint symmetry and map synchronization. The input consists of symmetry groups and pair-wise maps independently computed using off-the-shelf techniques in each domain. The output consists of consistent symmetry groups and pair-wise maps. Our approach exhibits a few appealing properties. First, the input symmetry groups and pair-wise maps may be noisy and incomplete (e.g., one input self-symmetry for a subset of objects and one object map per pair among a subset of object pairs), and our approach automatically rectifies, propagates and completes symmetry groups and pair-wise maps. Moreover, when we have sufficient data, we do not need to specify the type of the underlying symmetry group. It can be automatically determined by the input data, i.e., by aggregating self-symmetries among the input objects.

We evaluated our approach on a diverse set of object collections ranging from 2D images to 3D models. Experimental results show that our approach outperforms state-of-the-art methods in joint map computation from an object collection and symmetry group detection from each object in isolation.

2 Related Works

Map synchronization. An emerging focus in object matching is to optimize maps jointly among an object collection. The intuition is that the map between a pair of dissimilar objects can be obtained by composing maps along a path of similar object pairs. A way to formalize this intuition is through the cycle-consistency constraint, i.e., composite maps along cycles should be identity maps. This has led to the problem of map synchronization, which takes as input maps computed between pairs of objects in isolation and outputs improved maps that are consistent along cycles. In [16], Huang and Guibas established a connection between the cycle-consistency constraint and the fact that data matrices that store cycle-consistent pair-wise maps in blocks are low-rank and/or positive semi-definite. This work has stimulated several recent works that formulate map synchronization as low-rank matrix recovery using convex optimization [16, 5] and non-convex optimization [20, 67, 51, 7]. Our approach also falls into this category. However, the focus is on establishing a matrix representation for joint map and symmetry synchronization.

Co-symmetry detection and matching. Symmetry detection can be considered a variant of shape matching, which seeks to compute self-maps that preserve extrinsic [36, 44, 22, 52] or intrinsic [40, 28, 65, 46, 45, 25, 64, 63, 62] distances. In this regard, symmetry detection shares a similar limitation with pair-wise matching, namely, existing methods tend to break when the underlying symmetries become more and more approximate. To address this issue, we propose to optimize symmetry groups jointly among an object collection to improve the resulting symmetry group on each object. To the best of our knowledge, existing works on this topic have focused on individual pairs of objects so far. [33] proposed to detect a reflection symmetry axis for boosting the performance of correspondence computation. [39] introduced the first approach for factorizing out the underlying symmetry group when establishing shape maps. However, the approach only considers a pair of objects and does not utilize lifting. In addition, the formulation requires specifying the underlying symmetry group, or in other words, it does not perform joint inference of symmetries and maps. [57] developed an approach for joint structure recovery and matching of man-made 3D shapes. The approach is based on graph matching, and thus requires part-based representations as input.

Matching via lifting. Our symmetry group encoding scheme is motivated from the idea of lifting for convex relaxations of second-order MAP inference [27, 18]. When using indicators associated with the first-order potential functions, the second-order potential functions become quadratic in these indicator variables. The idea of lifting is to introduce an additional variable for each quadratic term.

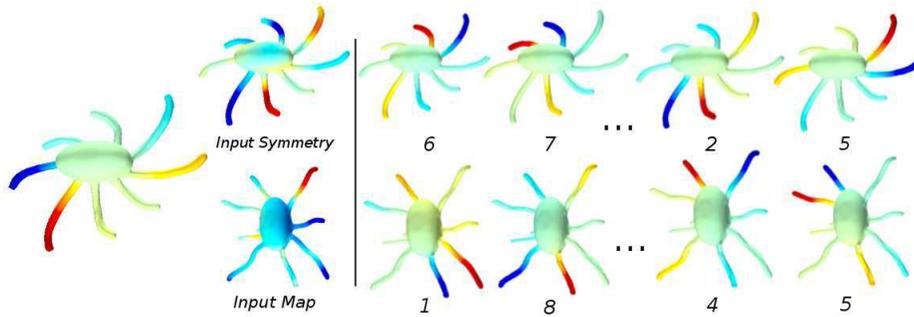


Fig. 2. This figure is better viewed in color. Corresponding points have the same colors. (Left) This input consists of one self-map per-shape and one pair-wise map per-pair. Here we show one shape among 20 shapes and one pair among 380 pairs. (Right) The output of our algorithm on the same shape and shape pair, respectively. Note that our approach not only improves symmetry groups and pair-wise maps, but we also complete the missing ones by propagating self-symmetries and maps.

Using these additional variables (which form the lifted space), the objective function becomes linear and easy to solve. The lifting idea has also been used for object matching [23, 14, 58]. The novelty of our representation includes (1) the decoding scheme, (2) the types of symmetry groups that can be exactly recovered (to be introduced later), and (3) various properties of this encoding scheme, which will be used in formulating joint map and symmetry synchronization.

3 Problem Statement and Approach Overview

This section formally states the joint map and symmetry synchronization problem that is studied in this paper and presents an overview of our approach.

Problem statement. Suppose we are given n relevant objects S_1, \dots, S_n . In this paper, we focus on the case where objects are discrete point sets (e.g., image pixels [32] and feature/sample points extracted from images [34] or shapes [30]), and where a map or a self-map is given by a set of point-wise correspondences.

We assume these objects are generated from an underlying universal object \bar{S}_0 (also a point set), which possesses a symmetry group. The generation process consists of taking a subset and/or deformation. This setting covers many practical scenarios. For example, $\{S_i\}$ could be partial observations of an underlying object (e.g., in multi-view structure from motion). As another example, $\{S_i\}$ could also be objects that fall into the same category (See Figure 2). The input to our approach consists of noisy pair-wise maps $P_{ij}^{in} : S_i \rightarrow S_j, (i, j) \in \mathcal{E}$ computed using off-the-shelf methods along a pre-computed edge set \mathcal{E} . In particular, \mathcal{E} includes self loops, each of which is associated with a pre-computed self-map $P_{ii}^{in} : S_i \rightarrow S_i$. The output of our approach consists of the universal object \bar{S}_0 , its symmetry group \mathcal{G} (we write \mathcal{G} instead of \mathcal{G}_0 to make the notations uncluttered), and an embedding map $P_{i0} : S_i \rightarrow \bar{S}_0$ for each object S_i . As we

will see later, the pairwise maps from S_i to S_j are induced from $\{P_{i0}, 1 \leq i \leq n\}$ and \mathcal{G} .

An important feature of our approach is that it takes as input incomplete observations (e.g., one pair-wise map per object pair despite the fact that the underlying symmetry group suggests multiple plausible pair-wise maps). The promise of our approach comes from the fact that when looking at the input maps among an object collection as a whole, each element of the underlying symmetry group and the corresponding all pair-wise maps are densely sampled, providing sufficient observations for recovery.

Approach overview. The main idea of our approach is to develop a matrix representation of symmetry groups and maps, which allows us to formulate joint map and symmetry synchronization as optimizing matrices. A simple approach is to use a binary correspondence matrix to encode correspondences in symmetry groups and pair-wise maps. As we will discuss later, this representation is insufficient when the size of a symmetry group is bigger than 2. To address this issue, we propose to use the product operator for encoding. We show that our proposed encoding is lossless for many symmetry groups. Using the same idea, we then show how to define and encode embedding maps $P_{i0}, 1 \leq i \leq n$ and pair-wise maps $P_{ij}, 1 \leq i, j \leq n$. Based on this encoding, we propose an optimization framework for recovering the underlying embedding maps $P_{i0}, 1 \leq i \leq n$ and self-symmetry group \mathcal{G} . To address the computational overhead incurred by the tensor operator, we show how to perform joint map and symmetry synchronization using reduced basis (c.f. [38]).

4 Symmetry and Map Representation via Lifting

In this section, we show how to effectively encode symmetry groups and pair-wise maps under the point-based representation. Section 4.1 describes how to encode symmetry groups. Section 4.2 introduces how to encode maps between symmetric objects that are partially similar.

4.1 Symmetry Groups

Consider an object S that consists of m points. A self-map is given by a permutation $P \in [0, 1]^{m \times m}$, where each row and column has exactly one non-zero entry. With \mathcal{P}_m we denote the space of all permutations. Let $\mathcal{G} \subset \mathcal{P}_m$ denote a symmetry group associated with S . For a thorough introduction to symmetry groups, we refer to [11] for more details.

In this paper, we are interested in symmetry groups where the largest orbit size is equal to the group size:

Definition 1. Let $o_{\mathcal{G}}(i) \subset \{1, \dots, m\}$ denote the orbit³ of the i -th element under a symmetry group \mathcal{G} . We call \mathcal{G} a cyclic group if

$$|\mathcal{G}| = \max_{1 \leq i \leq m} |o_{\mathcal{G}}(i)|. \quad (1)$$

³ Intuitively, the orbit collects correspondences induced from a symmetry group.

Note that cyclic groups cover a rich family of symmetry groups in our physical world. For example, standard reflection, n -fold rotational symmetry, and translation symmetries (both in the intrinsic and extrinsic sense) can be described as cyclic groups under suitable placements of samples.

When considering matrix representations of \mathcal{G} , a straight-forward approach is to use a correspondence matrix (of dimension $m \times m$) to encode all correspondences induced from the symmetry group. The downside of this encoding, however, is that when $|\mathcal{G}| > 2$, there may exist multiple symmetry groups that correspond to the same encoding. One such example is given in Appendix A. To address this issue, we propose to consider the following encoding using the tensor operator:

$$Q := \sum_{P \in \mathcal{G}} P \otimes P \in \{0, 1\}^{m^2 \times m^2}. \quad (2)$$

A key property of Q is that when \mathcal{G} is a cyclic group, then it can be recovered from the elements of Q :

Proposition 1. *If a symmetry group \mathcal{G} is a cyclic group, then it can be exactly recovered from the elements of Q .*

To prove Prop. 1 (details are deferred to Appendix B.2) and for later usage, we define the following linear operator that shuffles the elements of Q :

Definition 2. *Consider a matrix $A \in \mathbb{R}^{m_1^2 \times m_2^2}$. Let $A_{a,b}$ denote (a, b) -th element of matrix A . Define $\mathcal{F} : \mathbb{R}^{m_1^2 \times m_2^2} \rightarrow \mathbb{R}^{m_1 m_2 \times m_1 m_2}$ as*

$$\mathcal{F}(A)_{i_1 m_2 + i_2, j_1 m_2 + j_2} = A_{i_1 m_1 + j_1, i_2 m_2 + j_2}, \quad 0 \leq i_1, j_1 \leq m_1 - 1, 0 \leq i_2, j_2 \leq m_2 - 1.$$

It is easy to check that $\mathcal{F}(Q)$ is low-rank:

Fact 1. *Let $\text{vec}(P) = (\mathbf{p}_1^T, \dots, \mathbf{p}_m^T)^T \in \mathbb{R}^{m^2}$ be the vector that unfolds the columns of $P = (\mathbf{p}_1, \dots, \mathbf{p}_m) \in \mathbb{R}^{m \times m}$. Then*

$$\mathcal{F}(Q) = \sum_{P \in \mathcal{G}} \text{vec}(P) \cdot \text{vec}(P)^T. \quad (3)$$

Proof. See Appendix B.1. □

4.2 Maps between Symmetry Groups

We proceed to define and encode maps between symmetry objects. To handle the case, where two objects being matched are subsets of an underlying complete object, we first define maps from a partial object to a complete object. Later, we show how to extend this setting to properly defined maps between partially similar objects.

Let us consider a source object S_1 and a target object \bar{S}_0 (which denotes the underlying universal object in this paper). To define the map from S_1 to \bar{S}_0 , we introduce a latent object \bar{S}_1 that is (i) a superset of S_1 , and (ii) a copy of \bar{S}_0 ,

i.e., $|\overline{S}_1| = |\overline{S}_0|$, and the symmetry group \mathcal{G}_1 of \overline{S}_1 is isomorphic to \mathcal{G} of \overline{S}_0 . Our goal is to define the map from S_1 to \overline{S}_0 through the inclusion map $E_1 : S_1 \rightarrow \overline{S}_1$ and a properly defined map from \overline{S}_1 to \overline{S}_0 :

Definition 3. Let $f_1 : \mathcal{G} \rightarrow \mathcal{G}_1$ be the group isomorphism. We say a map $\overline{P}_{10} : \overline{S}_1 \rightarrow \overline{S}_0$ is proper if

$$\overline{P}_{10} \cdot f_1(P) = P \cdot \overline{P}_{10}, \quad \forall P \in \mathcal{G}. \quad (4)$$

In other words, \overline{P}_{10} is proper if it is consistent with the underlying symmetry groups. It is clear that each map \overline{P}_{10} induces an equivalence class $\{P\overline{P}_{10}, P \in \mathcal{G}\}$. Now we are ready to define proper maps from S_1 to \overline{S}_0 :

Definition 4. We say a map $P_{10} : S_1 \rightarrow \overline{S}_0$ is proper if there exist $\overline{S}_1 \supset S_1$ and a proper map $\overline{P}_{10} : \overline{S}_1 \rightarrow \overline{S}_0$ such that

$$P_{10} = \overline{P}_{10} \cdot E_1. \quad (5)$$

Note that P_{10} also induces an equivalent class $\mathcal{M}_{10} := \{PP_{10}, P \in \mathcal{G}\}$. We again use the product operator to encode \mathcal{M}_{10} :

$$Q_{10} = \sum_{P_{10} \in \mathcal{M}_{10}} P_{10} \otimes P_{10}. \quad (6)$$

Similar to symmetry groups, we have the following recovering condition:

Proposition 2. When \mathcal{G} is a cyclic group and S_1 contains one element of a maximal orbit of \overline{S}_1 , then \mathcal{M}_{10} can be exactly recovered from the elements of Q_{10} described in (6).

The proof is similar to that of Proposition 1 and is deferred to Appendix B.3. Now we are ready to define maps between a partially similar object pair (S_1, S_2) :

Definition 5. We say a map $P_{12} : S_1 \rightarrow S_2$ is proper, if there exist proper maps $P_{i0} : S_i \rightarrow \overline{S}_0$, so that

$$P_{12} = P_{20}^T P_{10}. \quad (7)$$

Note that each proper map P_{12} induces an equivalence class

$$\mathcal{M}_{12} := \{P_{20}^T P P_{10}, P \in \mathcal{G}\}. \quad (8)$$

It is easy to check that \mathcal{M}_{12} is independent of the particular choice of P_{i0} , $1 \leq i \leq 2$.

Again, we encode \mathcal{M}_{12} using the tensor product:

$$Q_{12} := \sum_{P_{12} \in \mathcal{M}_{12}} P_{12} \otimes P_{12} \stackrel{(8)}{=} (P_{20} \otimes P_{20})^T Q_{00} (P_{10} \otimes P_{10}). \quad (9)$$

Finally, we introduce a property of $\mathcal{F}(Q_{12})$, which we will use later:

Proposition 3. *When $|o_{\mathcal{G}}(i)| = |\mathcal{G}|, 1 \leq i \leq m$, then each $\text{vec}(P_{12})$ is an eigenvector of $\mathcal{F}(Q_{12})$, i.e.,*

$$\mathcal{F}(Q_{12})\text{vec}(P_{12}) = \|\text{vec}(P_{12})\|^2 \cdot \text{vec}(P_{12}), \quad \forall P_{12} \in \mathcal{M}_{12}. \quad (10)$$

The proof is straight-forward, as two different members of \mathcal{M}_{12} , their non-zero entries do not overlap. In practice, the number of elements where $|o_{\mathcal{G}}(i)| < |\mathcal{G}|$ remains small (e.g., points on 1D reflection axis versus points on the entire image). This means (10) is at least approximately satisfied.

5 Functional-Based Joint Map and Symmetry Synchronization

In this section, we introduce a computational framework for joint map and symmetry synchronization. Section 5.1 describes a generalized representation for symmetry groups and maps using reduced basis. Section 5.2 then introduces an efficient approach using this reduced representation.

5.1 Symmetry and Map Representation Using Reduced Basis

Laplacian reduced basis. We begin with reviewing the reduced basis representation for encoding maps between discrete objects [38, 61, 47]. Consider an object S with m points, we associate S with an orthogonal matrix $B \in \mathbb{R}^{m \times k}$, which projects S onto the feature space spanned by the columns of B . Following [38, 61], we use the first k eigenvectors of Laplacian matrices, e.g., graph Laplacian for images and Laplace-Beltrami for 3D meshes. Note that it is possible to use other basis computation methods such as partial functional map [47]. We leave this as a future work.

Given two objects S_1 and S_2 and the associated basis B_1 and B_2 , we represent a map from S_1 to S_2 in the feature space as a matrix $X_{12} \in \mathbb{R}^{k \times k}$ (called functional map in the literature ([38])). Given a point-based map $P_{12} : S_1 \rightarrow S_2$, we can derive the corresponding functional map as

$$X_{12} := \underset{X}{\text{argmin}} \|XB_1^T - B_2^T P_{12}\|_{\mathcal{F}}^2 = B_2^T P_{12} B_1. \quad (11)$$

In the other direction, we convert a functional map X_{12} into a point-map P_{12} via nearest neighbor query with respect to rows of B_1 and B_2 . More precisely, for the i -th point of S_1 , we compute the index of its target point as $j^* = \underset{j}{\text{argmin}} \|B_2^T P_{12} \mathbf{e}_j - X_{12} B_1^T \mathbf{e}_i\|$,⁴ where $A \mathbf{e}_i$ extracts the i -th column of A .

Given the underlying point-based map P_{12} , this functional map encoding scheme is accurate if there exists a linear map X_{12} such that $B_2^T P_{12} \approx X_{12} B_1^T$.

⁴ It is also possible to enforce injectivity by solving a linear assignment.

In our experiments, we found using $k = 25$ Laplacian basis provides a fairly accurate encoding between discrete objects. We will also analyze the effects of varying k in the supplemental material.

Encoding using reduced functional basis. Using reduced basis to encode symmetry groups and maps between symmetry groups is quite similar to that in the point-based representation. Let B_0 be a basis of \overline{S}_0 (As we see later, our algorithm will recover both \overline{S}_0 and the reduced basis B_0). With $X_{00} = B_0^T P B_0$, $X_{10} = B_0^T P_{10} B_1^T$ and $X_{12} = B_2^T P_{12} B_1$ we denote the functional map representations of P , P_{10} and P_{12} , respectively. With this setup, we rewrite the functional representations of (2), (6) and (9) as

$$Y_{00} := \sum_{P \in \mathcal{G}} (X_{00} \otimes X_{00}), \quad Y_{10} := \sum_{P_{10} \in \mathcal{M}_{10}} (X_{10} \otimes X_{10}), \quad (12)$$

$$Y_{12} := (X_{20} \otimes X_{20})^T Y_{00} (X_{10} \otimes X_{10}). \quad (13)$$

Encoding using reduced Laplacian basis shares many properties as using the point-based representation. As discussed in Proposition 3, $\text{vec}(P_{12})$ is approximately an eigen-vector of Q_{12} . We find that under mild conditions (See Appendix B.5 for a discussion), this property also holds under the reduced basis:

$$\mathcal{F}(Y_{12}) \text{vec}(X_{12}) \approx \|\text{vec}(X_{12})\|^2 \text{vec}(X_{12}). \quad (14)$$

Decoding via alternating minimization. Once we have obtained an encoding Y_{i0} under reduced basis, we recover the underlying functional maps and point-based maps via alternating minimization (c.f. [38]). Motivated from Fact 1, the functional maps are forced to lie within the leading eigenspace of $\mathcal{F}(Y_{i0})$. Due to space constraint, the details are left in Appendix C.

5.2 Joint Map and Symmetry Synchronization

In this section, we describe our approach for joint map and symmetry synchronization using the representation developed above. At this stage, we assume we have computed reduced basis $B_i, 1 \leq i \leq n$. We also assume we have computed initial functional maps $X_{ij}^m, (i, j) \in \mathcal{E}$ along an edge set \mathcal{E} . Both the input maps and the edge set are specified by off-the-shelf approaches, and we will discuss them in Section 6. Our goal is to recover the underlying universal object \overline{S}_0 , its symmetry group G_0 , and the maps $P_{i0} : S_i \rightarrow \overline{S}_0$. Our approach consists of a synchronization step and an extraction step.

Synchronization. We introduce a $n \times n$ block matrix $Y \in \mathbb{R}^{nk^2 \times nk^2}$ that encodes each Y_{ij} in the block of Y , i.e., Y_{ij} corresponds to j -th row and i -th column of Y . Similar to the point-based representation, Y is low-rank. To see this, we introduce $Z_i = Y_{00}^{\frac{1}{2}} (X_{i0} \otimes X_{i0}), 1 \leq i \leq n$. Let $Z = (Z_1^T, \dots, Z_n^T)^T$. Then (13) gives rise to

$$Y = Z Z^T. \quad (15)$$

We propose to utilize numerical optimization to obtain Y and Z . Specifically, by combing (15) and (14), we arrive at the following objective function for

recovery:

$$\min_{Y,Z} \sum_{(i,j) \in \mathcal{E}} \|\mathcal{F}(Y_{ij})\text{vec}(X_{ij}^{in}) - \|\text{vec}(X_{ij}^{in})\|^2 \text{vec}(X_{ij}^{in})\| + \lambda \|Y - ZZ^T\|_{\mathcal{F}}^2 \quad (16)$$

where \mathcal{F} is the Frobenius norm. Note that we use the block-wise L1-norm to suppress the noise in the input. We set $\lambda = \frac{1}{\sqrt{n}}$ in our experiments. In this paper, we combine spectral initialization and reweighted non-linear least squares to solve (16). The details are deferred to Appendix D.

Decoding. Since our goal is to recover the underlying object \bar{S}_0 , merely recovering pair-wise maps Y_{ij} is insufficient. In contrast, we seek to recover $X_{i0}, 1 \leq i \leq n$ and Y_{00} . Let Z^* be the optimal solution to (16). If Z^* is an exact recovery, then there exists $V \in \mathbb{R}^{k^2 \times k^2}$, so that

$$V \cdot Z_i^{*T} = X_{i0} \otimes X_{i0}, \quad 1 \leq i \leq n, \quad Y_{00} = V^{-1T} V^{-1}. \quad (17)$$

For decoding, we set up the following objective function to recover V and $X_{i0}, 1 \leq i \leq n$:

$$\min_{V, X_{i0}, 1 \leq i \leq n} \sum_{i=1}^n \|V \cdot Z_i^{*T} - X_{i0} \otimes X_{i0}\|_{\mathcal{F}}^2 + \mu \|V - I_{k^2}\|_{\mathcal{F}}^2 \quad (18)$$

where the second term is introduced to avoid obtaining the trivial solution $V = 0, X_{i0} = 0, 1 \leq i \leq n$. In our experiment, we used $\mu = 10^{-3}$. We again perform alternating minimization to solve (18). Please refer to Appendix E for details.

As X_{i0} aligns the feature space associated with S_i with that of the latent universal object S_0 , we apply clustering on the columns of $X_{i0} B_i^T, 1 \leq i \leq n$ to recover B_0 , which also specifies S_0 . In our experiments, we used single-linkage clustering and set the total number of clusters as m . We discard all clusters whose size are smaller than 3. Each cluster center corresponds to one column of B_0 . Finally, we apply the decoding scheme described in (20) to recover the underlying symmetry group G .

6 Results

In this section, we evaluate our approach for joint map and symmetry synchronization on 2D images (Section 6.1) and 3D shapes (Section 6.2).

6.1 Experimental Evaluations on 2D Images

Experimental Setup. As illustrated in Figure 1 and Figure 3, we consider the application of reconstructing a generic sparse 3D pointcloud from multi-views of similar but different objects [67, 4, 59] (i.e., category-specific reconstruction). A crucial task for the reconstruction is to establish consistent feature correspondences across multiple images. We share this general motivation in our experiments, but focus on symmetric objects. In this regard, we collect two datasets. The first dataset consists of 16 images of similar stool objects that possess a

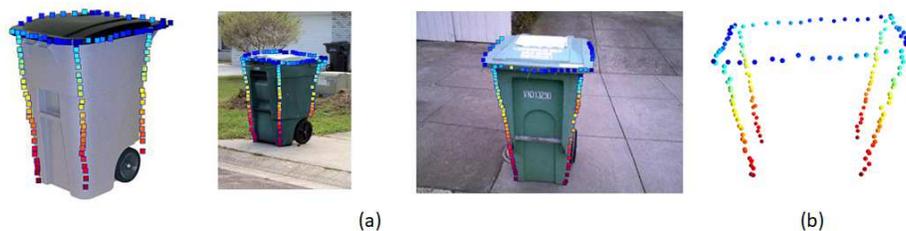


Fig. 3. Category specific reconstruction from internet images. (a) Subset of input images with matched edge feature points [12]. (b) Reconstruction using our approach. Color-coding highlights the recovered symmetry group.

four-way rotational symmetry (See Figure 1). The second dataset consists of 16 images of similar trash container objects (See Figure 3). For both datasets, we annotate keypoints for evaluating feature correspondences. The full datasets and annotated feature correspondences are included in Appendix H. The same as [67, 4, 59], we assume the underlying objects are segmented out. In addition, we also evaluated on Sedan and SUV from [67]. We use these two datasets to compare our approach against standard joint map optimization approaches that do not explicitly consider the underlying symmetry group.

We follow the procedure described in [67], which applies structural forests [12] and graph matching [6] to generate the feature points and perform pair-wise matching, respectively. The reduced basis we utilize are the first $k = 25$ eigenvectors of unnormalized graph Laplacian of the graphs used in pair-wise matching (See [6] for more details). For Stool and TrashContainer, we also apply this procedure to match each image with itself, which gives a self-map. To test whether the improvement of our algorithm on symmetry detection is consistent with respect to different input symmetries, we also evaluated on applying [13] for symmetry detection (we used the same procedure to exclude identity maps). Regarding the edge set \mathcal{E} , we connect each image with 8-closest images in terms of affinity scores for Stool and TrashContainer. The same as [67], we let \mathcal{E} connect all image pairs for Sedan and SUV.

We consider four state-of-the-art synchronization techniques [19, 66, 67, 7] for baseline evaluation. In particular, [19] synchronizes consistent functional maps; [66] optimizes point-based correspondences by enforcing three-cycles; [67] applies a fast low-rank matrix recovery approach to optimize consistent point-wise correspondences. [7] uses non-convex optimization to obtain consistent point-wise correspondences. The same as [67], we report the percentage of feature correspondences whose error with respect to annotated feature correspondences fall within a varying threshold δ . Note that for Stool and TrashContainer, we evaluate with respect to the closest map (in terms of cumulative feature correspondence error) induced by the underlying symmetry group. This is in contrast to evaluating the quality of each individual correspondence with respect to symmetric correspondences, where the closest correspondences may be inconsistent with each other.

	Stool	TC	Sedan	SUV	Octop.	Arma.	Ant	Bird	Fish	4leg	Glass.	Hand	Human	Plane	Mean
Sym-In	71.2	68.7	na	na	52.4	72.9	69.1	59.3	54.5	64.0	78.2	55.1	89.3	61.2	67.0
Sym-Out	84.5	81.2	na	na	76.1	91.2	81.0	67.2	71.3	79.2	90.1	69.2	94.5	73.2	79.7
SymII-In	65.9	63.1	na	na	48.3	68.1	66.4	57.4	51.5	63.2	71.7	49.2	81.3	60.0	63.2
SymII-Out	79.8	77.4	na	na	68.0	81.2	77.3	74.2	70.1	80.7	86.5	57.1	90.3	72.4	76.6
Input	60.7	58.5	na	na	44.9	69.5	64.1	53.0	48.1	58.0	76.3	51.7	83.8	35.7	60.1
Huang14	68.2	60.6	na	na	58.1	69.4	69.0	54.3	42.8	76.9	71.1	66.3	90.9	42.4	64.8
Zhou15a	63.2	59.8	na												
Zhou15b	64.6	62.8	na	na	56.5	70.1	71.1	51.3	43.7	71.9	66.7	68.1	86.1	44.4	63.7
Cosmo17	61.9	59.7	na	na	54.9	68.2	71.6	51.8	42.8	71.9	64.7	67.0	86.4	44.4	63.2
Ours	80.6	75.2	na	na	74.6	81.2	76.3	62.9	54.1	80.1	78.3	71.1	92.1	53.7	72.2
Input	na	na	78.7	77.6	na	65.7	47.8	39.8	45.6	56.6	54.1	47.5	69.4	30.2	50.7
Huang14	na	na	85.9	87.1	na	69.6	59.0	52.2	42.3	76.5	59.3	66.2	90.9	39.4	61.7
Zhou15a	na	na	86.1	85.2	na										
Zhou15b	na	na	87.7	86.2	na	68.2	59.3	47.6	46.1	70.3	59.2	59.7	88.1	38.5	59.7
Cosmo17	na	na	84.5	83.4	na	67.1	61.4	44.9	42.8	71.3	57.9	61.4	85.8	37.9	58.9
Ours	na	na	91.2	90.8	na	74.9	64.1	54.8	51.2	83.2	65.7	68.2	93.2	38.0	65.9

Table 1. Comparisons with respect to human annotated feature points. We report percentage of correspondences whose error are below $0.1d$, where $d = \max(w, h)$ for images and d is the shape diameter for 3D shapes, respectively. (Top) Symmetry detection. Images:Sym-IN [6] and SymII-IN [13]. Shapes:Sym-IN [25] and SymII-IN [40]. (Middle) Evaluation of object maps after factoring out the underlying self-symmetry group. Baseline approaches: Huang14:[19], Zhou15a:[66], Zhou15b:[67], Cosmo17:[7]. (Bottom). Evaluation of object maps in the original space.

Table 1 collects the statistics of each method when $\delta = 0.1 \max(w, h)$, where w and h are width and height of an input image.

Joint map and symmetry detection improves symmetry detection. As shown in Table 1(Top), our approach can drastically improve the input symmetry detection. For the first set of input symmetries (i.e., using [6]), our approach improves from 71.2% and 68.7% on Stool and TrashContainer, respectively, to 84.5% and 81.2%, respectively. For the second set of input symmetries (i.e., using [13]), our approach improves from 65.9% and 63.1% on Stool and TrashContainer, respectively, to 79.8% and 77.4%, respectively. Figure 4(Left) shows that the improvements are consistent when varying the error threshold.

Analysis of correspondence quality. Our approach outperforms baseline approaches considerably. Specifically, on Stool and TrashContainer, our approach achieved 80.6% and 75.2%, respectively. In contrast, the top-performing state-of-the-art methods only had 68.2% (from [67]) and 62.8% (from [19]), respectively. Figure 4(Middle) shows that the improvements are also consistent when varying the error threshold. One explanation is that due to the underlying symmetry, there are multiple plausible maps between each object pair. However, all existing methods are forced to output a single consistent map. In this case, it turns out their output tends to average multiple plausible maps, and the resulting correspondences may not be consistent across different feature points from the same image. It follows that their output may be far from the underlying map even after factoring out the underlying symmetry group.

To see the usefulness of the output of our approach, we perform SFM on each dataset of Stool and TrashContainer in isolation. Specifically, we run rigid reconstruction from optimized feature correspondences with an orthographic camera model and annotated viewpoints. As illustrated in Figure 1 and Figure 3, our approach can recover 3D point clouds that reflect the overall shapes of the Stool

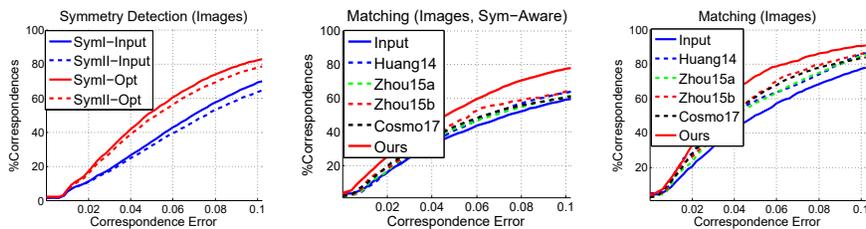


Fig. 4. Quantitative evaluations of joint map and symmetry synchronization on 2D images. (Left) Symmetry detection on Stool and TrashContainer. (Middle) Object matching evaluation in the quotient space on Stool and TrashContainer. (Right) Object matching evaluation in the original space on Sedan and SUV.

and TrashContainer. In contrast, the method of [67] was not successful on these two datasets.

We proceed to quantitatively compare our approach with baseline approaches on Sedan and SUV. In this case, we simply set the diagonal blocks as identity matrices in our approach. As shown in Table 1 and Figure 4(Right), our approach outperforms existing approaches. The benefits come from the product operator, which enables the joint matching procedure to detect consistent correspondence pairs obtained in the graph matching phase. Similar to using pair-wise consistency for boosting the performance of graph matching, these consistent correspondence pairs enhance optimized correspondences.

Timing. Our approach is efficient. In average, the running time of our approach on each dataset is 137.0s using a Matlab implementation on a single-core 3.2G CPU with 32G memory. This includes 12.5s for converting point-based maps into functional maps, 95.6s for joint map and symmetry optimization, and 28.9s for converting functional maps back to the original point-based maps.

6.2 Experimental Evaluations on 3D Shapes

Experimental Setup. We perform experimental evaluation on SHREC07-Watertight [15], which is a challenging dataset for evaluating shape maps. Specifically, SHREC07-Watertight contains 400 shapes across 20 categories. Among them, we choose 10 categories (i.e., Ant, Armadillo, Bird, Fish, Fourleg, Glasses, Hand, Human, Plane, and Octopus) that are suitable for inter-shape mapping. In particular, Octopus contains a non-trivial rotational symmetry group. Each other category contains a reflection symmetry. All categories have human-annotated correspondences for experimental evaluation. As for symmetry detection, we employ [25]. We also tested on [40] to see if the improvements are consistent with respect to different input symmetries. For pair-wise maps, we employ blended intrinsic maps [26]. We consider the same set of baseline approaches described in Section 6.1 except [66], which is specifically designed for image matching. In terms of evaluation protocol, we report the percentage of feature correspondences whose geodesic errors fall into a varying threshold in $[0, 0.1d]$, where d is the diameter of each shape in geodesic distance.

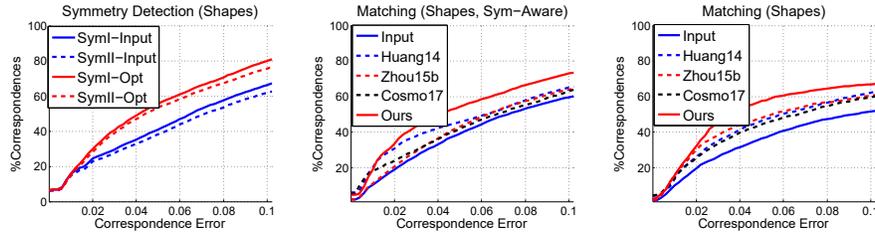


Fig. 5. Quantitative evaluations of joint map and symmetry synchronization on SHREC07-Watertight. (Left) Symmetry detection. (Middle) Object matching evaluation in the quotient space. (Right) Object matching evaluation in the original space.

Analysis of results. We first evaluate our approach on the performance of symmetry detection. As shown in Table 1, our approach improves the quality of input symmetries obtained from [25], by 23.7% and 12.5% on Octopus and remaining models, respectively. Replacing the input symmetries by the ones computed from [40], the improvements become 19.7% and 13.3% on Octopus and remaining categories, respectively. As shown in Figure 5(Left), the improvements are also consistent when varying the cutoff threshold. This again shows the huge potential of detecting symmetries jointly among an object collection.

Table 1 and Figure 5 compare our approach with baseline approaches for joint map optimization. Our approach outperforms baseline approaches both in the quotient space (i.e., after factoring out the self-symmetry) and in the original space. In particular, the relative improvements between our approach and baseline approaches in the quotient space are higher than those in the original space. This justifies the promise of joint map and symmetry optimization.

Timing. We employ a full observation graph on each category. The average running time on the same Desktop for each category is 560.1s, where functional map conversion, joint map and symmetry recovery and point-based map extraction take 38.2s, 313.1s and 208.8s, respectively.

7 Conclusions

In this paper, we have described an approach for joint map and symmetry synchronization. Our approach builds on a novel symmetry and map representation using the tensor operator. Based on this representation, we introduce a non-convex optimization scheme for recovering consistent symmetry groups and pair-wise maps from noisy input. Experimental results demonstrate that our approach is better than state-of-the-art methods for joint map synchronization without symmetry detection and symmetry detection from each object in isolation.

Acknowledgement. Qixing Huang would like to acknowledge support of this research from NSF DMS-1700234, a Gift from Snap Research, and a hardware Donation from NVIDIA.

References

1. Bajaj, C., Gao, T., He, Z., Huang, Q., Liang, Z.: SMAC: simultaneous mapping and clustering using spectral decompositions. In: Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018. pp. 334–343 (2018), <http://proceedings.mlr.press/v80/bajaj18a.html>
2. Berg, A.C., Berg, T.L., Malik, J.: Shape matching and object recognition using low distortion correspondences. In: CVPR (1). pp. 26–33. IEEE Computer Society (2005)
3. Bregler, C., Hertzmann, A., Biermann, H.: Recovering non-rigid 3d shape from image streams. In: CVPR. pp. 2690–2696. IEEE Computer Society (2000)
4. Carreira, J., Kar, A., Tulsiani, S., Malik, J.: Virtual view networks for object reconstruction. In: Computer Vision and Pattern Recognition (CVPR) (2015)
5. Chen, Y., Guibas, L.J., Huang, Q.: Near-optimal joint object matching via convex relaxation. In: Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014. pp. 100–108. Proceedings of Machine Learning Research, NY, USA (2014)
6. Cho, M., Lee, J., Lee, K.M.: Reweighted random walks for graph matching. In: ECCV (5). Lecture Notes in Computer Science, vol. 6315, pp. 492–505. Springer (2010)
7. Cosmo, L., Rodolà, E., Albarelli, A., Mémoli, F., Cremers, D.: Consistent partial matching of shape collections via sparse modeling. *Comput. Graph. Forum* **36**(1), 209–221 (2017)
8. Crandall, D., Owens, A., Snavely, N., Huttenlocher, D.: SfM with MRFs: Discrete-continuous optimization for large-scale structure from motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* **35**(12), 2841–2853 (December 2013)
9. Dai, Y., Li, H., He, M.: A simple prior-free method for non-rigid structure-from-motion factorization. *International Journal of Computer Vision* **107**(2), 101–122 (2014)
10. Daubechies, I., DeVore, R.A., Fornasier, M., Güntürk, C.S.: Iteratively reweighted least squares minimization: Proof of faster than linear rate for sparse recovery. In: 42nd Annual Conference on Information Sciences and Systems, CISS 2008, Princeton, NJ, USA, 19-21 March 2008. pp. 26–29 (2008). <https://doi.org/10.1109/CISS.2008.4558489>
11. Diaconis, P.: Group representations in probability and statistics. Institute of Mathematical Statistics Lecture Notes—Monograph Series, 11, Institute of Mathematical Statistics, Hayward, CA (1988)
12. Dollár, P., Zitnick, C.L.: Structured forests for fast edge detection. In: Proceedings of the 2013 IEEE International Conference on Computer Vision. pp. 1841–1848. ICCV '13, IEEE Computer Society, Washington, DC, USA (2013). <https://doi.org/10.1109/ICCV.2013.231>, <http://dx.doi.org/10.1109/ICCV.2013.231>
13. Duchenne, O., Bach, F.R., Kweon, I., Ponce, J.: A tensor-based algorithm for high-order graph matching. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(12), 2383–2395 (2011)
14. Dym, N., Maron, H., Lipman, Y.: Ds++: A flexible, scalable and provably tight relaxation for matching problems. *ACM Trans. Graph.* **36**(6), 184:1–184:14 (Nov 2017). <https://doi.org/10.1145/3130800.3130826>, <http://doi.acm.org/10.1145/3130800.3130826>

15. Giorgi, D., Biasotti, S., Paraboschi, L.: Shape retrieval contest 2007: Watertight models track (2007)
16. Huang, Q.X., Guibas, L.: Consistent shape maps via semidefinite programming. *Computer Graphics Forum* **32**(5), 177–186 (2013)
17. Huang, Q., Zhang, G., Gao, L., Hu, S., Butscher, A., Guibas, L.J.: An optimization approach for extracting and encoding consistent maps in a shape collection. *ACM Trans. Graph.* **31**(6), 167:1–167:11 (2012)
18. Huang, Q., Chen, Y., Guibas, L.: Scalable semidefinite relaxation for maximum a posterior estimation. In: *Proceedings of the 31st International Conference on Machine Learning*, PMLR 32(2). pp. 64–72 (2014)
19. Huang, Q., Wang, F., Guibas, L.: Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.* **33**(4), 36:1–36:11 (Jul 2014). <https://doi.org/10.1145/2601097.2601111>, <http://doi.acm.org/10.1145/2601097.2601111>
20. Huang, X., Liang, Z., Bajaj, C., Huang, Q.: Translation synchronization via truncated least squares. In: *NIPS* (2017)
21. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. *Computer Graphics Forum* **30**(6), 1681–1707 (2011)
22. Kerber, J., Bokeloh, M., Wand, M., Seidel, H.P.: Scalable symmetry detection for urban scenes. *Computer Graphics Forum* **32**(1), 3–15 (2013)
23. Kezurer, I., Kovalsky, S.Z., Basri, R., Lipman, Y.: Tight relaxation of quadratic matching. *Comput. Graph. Forum* **34**(5), 115–128 (2015)
24. Kim, V.G., Li, W., Mitra, N.J., DiVerdi, S., Funkhouser, T.: Exploring collections of 3d models using fuzzy correspondences. *ACM Trans. Graph.* **31**(4), 54:1–54:11 (Jul 2012). <https://doi.org/10.1145/2185520.2185550>
25. Kim, V.G., Lipman, Y., Chen, X., Funkhouser, T.A.: Möbius transformations for global intrinsic symmetry analysis. *Comput. Graph. Forum* **29**(5), 1689–1700 (2010)
26. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *ACM Trans. Graph.* **30**(4), 79:1–79:12 (Jul 2011)
27. Kumar, M.P., Kolmogorov, V., Torr, P.H.S.: An analysis of convex relaxations for MAP estimation of discrete mrfs. *Journal of Machine Learning Research* **10**, 71–106 (2009). <https://doi.org/10.1145/1577069.1577072>, <http://doi.acm.org/10.1145/1577069.1577072>
28. Lasowski, R., Tevs, A., Seidel, H., Wand, M.: A probabilistic framework for partial intrinsic symmetries in geometric data. In: *ICCV*. pp. 963–970. IEEE Computer Society, Kyoto, Japan (2009)
29. Leonardos, S., Zhou, X., Daniilidis, K.: Distributed consistent data association via permutation synchronization. In: *ICRA*. pp. 2645–2652. IEEE (2017)
30. Lipman, Y., Funkhouser, T.: Möbius voting for surface correspondence. *ACM Trans. Graph.* **28**(3), 72:1–72:12 (Jul 2009). <https://doi.org/10.1145/1531326.1531378>, <http://doi.acm.org/10.1145/1531326.1531378>
31. Liu, C., Yuen, J., Torralba, A.: Nonparametric scene parsing via label transfer. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(12), 2368–2382 (2011)
32. Liu, C., Yuen, J., Torralba, A.: Sift flow: Dense correspondence across scenes and its applications. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(5), 978–994 (May 2011)
33. Liu, T., Kim, V.G., Funkhouser, T.A.: Finding surface correspondences using symmetry axis curves. *Comput. Graph. Forum* **31**(5), 1607–1616 (August 2012)

34. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* **60**(2), 91–110 (Nov 2004). <https://doi.org/10.1023/B:VISI.0000029664.99615.94>, <https://doi.org/10.1023/B:VISI.0000029664.99615.94>
35. Markovskiy, I.: *Low Rank Approximation: Algorithms, Implementation, Applications*. Springer Publishing Company, Incorporated (2011)
36. Mitra, N.J., Guibas, L.J., Pauly, M.: Partial and approximate symmetry detection for 3d geometry. *ACM Trans. Graph.* **25**(3), 560–568 (July 2006)
37. Nguyen, A., Ben-Chen, M., Welnicka, K., Ye, Y., Guibas, L.J.: An optimization approach to improving collections of shape maps. *Comput. Graph. Forum* **30**(5), 1481–1491 (2011). <https://doi.org/10.1111/j.1467-8659.2011.02022.x>
38. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.J.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* **31**(4), 30:1–30:11 (2012). <https://doi.org/10.1145/2185520.2185526>
39. Ovsjanikov, M., M erigot, Q., Patraucean, V., Guibas, L.J.: Shape matching via quotient spaces. *Comput. Graph. Forum* **32**(5), 1–11 (2013). <https://doi.org/10.1111/cgf.12167>
40. Ovsjanikov, M., Sun, J., Guibas, L.J.: Global intrinsic symmetries of shapes. *Comput. Graph. Forum* **27**(5), 1341–1348 (2008)
41. Pachauri, D., Kondor, R., Sargur, G., Singh, V.: Permutation diffusion maps (pdm) with application to the image association problem in computer vision. In: *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*. pp. 541–549. NIPS’14, MIT Press, Cambridge, MA, USA (2014), <http://dl.acm.org/citation.cfm?id=2968826.2968887>
42. Pachauri, D., Kondor, R., Singh, V.: Solving the multi-way matching problem by permutation synchronization. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 26*, pp. 1860–1868. Curran Associates, Inc. (2013)
43. Pachauri, D., Kondor, R., Singh, V.: Solving the multi-way matching problem by permutation synchronization. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*. pp. 1860–1868. NIPS’13, Curran Associates Inc., USA (2013), <http://dl.acm.org/citation.cfm?id=2999792.2999820>
44. Podolak, J., Shilane, P., Golovinskiy, A., Rusinkiewicz, S., Funkhouser, T.: A planar-reflective symmetry transform for 3d shapes. *ACM Trans. Graph.* **25**(3), 549–559 (July 2006)
45. Raviv, D., Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Full and partial symmetries of non-rigid shapes. *International Journal of Computer Vision* **89**(1), 18–39 (2010)
46. Raviv, D., Bronstein, M.M., Sapiro, G., Bronstein, A.M., Kimmel, R.: Diffusion symmetries of non-rigid shapes. In: *In Proc. 3DPVT (2010)*
47. Rodol , E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. *CoRR* **abs/1506.05274** (2015), <http://arxiv.org/abs/1506.05274>
48. Rubinstein, M., Joulun, A., Kopf, J., Liu, C.: Unsupervised joint object discovery and segmentation in internet images. In: *CVPR*. pp. 1939–1946. IEEE Computer Society (2013), <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2013.html#RubinsteinJKL13>
49. Rubio, J.C., Serrat, J., L pez, A.M., Paragios, N.: Unsupervised co-segmentation through region matching. In: *CVPR*. pp. 749–756. IEEE Computer Society (2012)

50. Russell, B., Torralba, A., Liu, C., Fergus, R., Freeman, W.T.: Object recognition by scene alignment. In: Platt, J.C., Koller, D., Singer, Y., Roweis, S.T. (eds.) *Advances in Neural Information Processing Systems 20*. pp. 1241–1248. Curran Associates, Inc. (2007)
51. Shen, Y., Huang, Q., Srebro, N., Sanghavi, S.: Normalized spectral map synchronization. In: NIPS. pp. 4925–4933. Curran Associates, Barcelona, Spain (2016)
52. Shi, Z., Alliez, P., Desbrun, M., Bao, H., Huang, J.: Symmetry and orbit detection via lie-algebra voting. *Comput. Graph. Forum* **35**(5), 217–227 (2016)
53. Singer, A., Wu, H.T.: Spectral convergence of the connection laplacian from random samples. *Information and Inference: A Journal of the IMA* **35**(6), 58–123 (2017)
54. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: Exploring photo collections in 3d. *ACM Trans. Graph.* **25**(3), 835–846 (Jul 2006). <https://doi.org/10.1145/1141911.1141964>, <http://doi.acm.org/10.1145/1141911.1141964>
55. Tam, G.K.L., Martin, R.R., Rosin, P.L., Lai, Y.K.: Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. *ACM Trans. Graph.* **33**(1), 4:1–4:17 (Feb 2014)
56. Tani, T., Sinha, S.N., Sato, Y.: Joint recovery of dense correspondence and cosegmentation in two images. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016. pp. 4246–4255 (2016). <https://doi.org/10.1109/CVPR.2016.460>, <https://doi.org/10.1109/CVPR.2016.460>
57. Tevs, A., Huang, Q., Wand, M., Seidel, H.P., Guibas, L.: Relating shapes via geometric symmetries and regularities. *ACM Trans. Graph.* **33**(4), 119:1–119:12 (Jul 2014). <https://doi.org/10.1145/2601097.2601220>, <http://doi.acm.org/10.1145/2601097.2601220>
58. 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: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
59. Vicente, S., Carreira, J., Agapito, L., Batista, J.: Reconstructing PASCAL VOC. In: CVPR. pp. 41–48. IEEE Computer Society (2014)
60. Wang, F., Huang, Q., Guibas, L.: Image co-segmentation via consistent functional maps. In: In Proceedings of the 14th International Conference on Computer Vision (ICCV) (2013)
61. Wang, F., Huang, Q., Guibas, L.J.: Image co-segmentation via consistent functional maps. In: IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1–8, 2013. pp. 849–856. Computer Vision Foundation., Sydney, Australia (2013)
62. Wang, H., Huang, H.: Group Representation of Global Intrinsic Symmetries. *Computer Graphics Forum* **36**, 51–61 (October 2017)
63. Wang, H., Simari, P., Su, Z., Zhang, H.: Spectral global intrinsic symmetry invariant functions. In: Proceedings of Graphics Interface 2014. pp. 209–215. GI '14, Canadian Information Processing Society, Toronto, Ont., Canada, Canada (2014), <http://dl.acm.org/citation.cfm?id=2619648.2619683>
64. Xu, K., Zhang, H., Jiang, W., Dyer, R., Cheng, Z., Liu, L., Chen, B.: Multi-scale partial intrinsic symmetry detection. *ACM Trans. Graph.* **31**(6), 181:1–181:11 (Nov 2012)

65. Xu, K., Zhang, H., Tagliasacchi, A., Liu, L., Li, G., Meng, M., Xiong, Y.: Partial intrinsic reflectional symmetry of 3d shapes. *ACM Trans. Graph.* **28**(5), 138:1–138:10 (Dec 2009)
66. Zhou, T., Lee, Y.J., Yu, S.X., Efros, A.A.: Flowweb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*. pp. 1191–1200 (2015)
67. Zhou, X., Zhu, M., Daniilidis, K.: Multi-image matching via fast alternating minimization. In: *ICCV*. pp. 4032–4040. IEEE Computer Society, Santiago, Chile (2015)