
SMAC: Simultaneous Mapping and Clustering Using Spectral Decompositions

Chandrajit Bajaj¹ Tingran Gao² Zihang He³ Qixing Huang¹ Zhenxiao Liang³

Abstract

We introduce a principled approach for *simultaneous mapping and clustering* (SMAC) for establishing consistent maps across heterogeneous object collections (e.g., 2D images or 3D shapes). Our approach takes as input a heterogeneous object collection and a set of maps computed between some pairs of objects, and outputs a homogeneous object clustering together with a new set of maps possessing optimal intra- and inter-cluster consistency. Our approach is based on the spectral decomposition of a data matrix storing all pairwise maps in its blocks. We additionally provide tight theoretical guarantees for the accuracy of SMAC under established noise models. We also demonstrate the usefulness of our approach on synthetic and real datasets.

1. Introduction

Establishing maps (e.g. pointwise correspondences) across object collections is a fundamental problem spanning many scientific domains. High-quality maps facilitating information propagation and transformation are key to applications ranging from 3D reconstruction with partial scans (Huber & Hebert, 2001), data-driven geometry completion and reconstruction (Pauly et al., 2005), texture transfer (Schreiner et al., 2004; Kraevoy & Sheffer, 2004), to comparative biology (Boyer et al., 2011; Gao et al., 2017), joint data-analysis (Huang et al., 2011; Kim et al., 2012; Wang et al., 2013; 2014; Huang et al., 2014), and data exploration and organization (Kim et al., 2012; Huang et al., 2014).

High quality object maps are generally difficult to compute. Prior work on map computation focused on optimizing maps between pairs of objects; see (van Kaick et al., 2010) for

¹Department of Computer Science, The University of Texas at Austin ²Department of Statistics, The University of Chicago ³Institute for Interdisciplinary Information Sciences, Tsinghua University. Correspondence to: Chandrajit Bajaj <bajaj@cs.utexas.edu>, Qixing Huang <huangqx@cs.utexas.edu>.

a standard survey and (Kim et al., 2011; Mandad et al., 2017) for some recent advances. Despite the significant progress, state-of-the-art techniques tends to hit a barrier on the quality of maps that are computed in a pairwise manner. Building upon the availability of big-data, a recent line of research (Kim et al., 2012; Nguyen et al., 2011; Ovsjanikov et al., 2012; Huang et al., 2012; Huang & Guibas, 2013; Huang et al., 2014; Chen et al., 2014; Zhou et al., 2015a,b; Shen et al., 2016; Leonardos et al., 2017; Huang et al., 2017) considered computing many pairwise maps jointly among a collection of objects. The promise of these approaches hinges upon the observation that one way to obtain a high quality pairwise map between dissimilar objects is to choose a path connecting these objects but consisting of consecutive similar shapes: maps between similar objects are typically of higher quality, and so is the resulted composed map. From a regularization perspective, joint map computation leverages the generic consistency of a network of maps among multiple objects, in which composition of maps along cycles are expected to be close to the identity map.

However, the performance of these data-driven approaches relies predominantly on the *homogeneity* of the object collection, i.e. the input objects fall into the same category or sub-category (e.g. Chairs, Cars, and Human models). In the presence of *heterogeneous* data, where the input objects fall into multiple underlying categories, applying existing data-driven approaches without the category label information tend to produce unsatisfactory results. In this setting, even though existing methods are able to suppress the noise in intra-cluster maps within a single cluster, jointly computed maps for the entire object collection leads are often significantly worse. One explanation is that high fraction of incorrect inter-cluster maps tends to “contaminate” the regularization effect of intra-cluster maps. A natural resolution is to employ a two-stage cascading strategy that identifies the underlying clusters before computing the intra- and inter-cluster maps. Unfortunately, such clustering requires accurate quantification of the object similarities, which is a difficult problem in its own right. Meanwhile, the error produced in the clustering stage is unlikely remedied by the consistency-based regularization.

In this paper, we propose to solve the mapping and clustering problems *simultaneously*. Instead of explicitly relying on certain pairwise similarity and/or map distortion




Figure 1. The pairwise map distortion scores (Left) and the pairwise cycle-consistency scores (Right) derived from 3-cycles. The scores are normalized by the maximum score in each matrix.

scores, we identify the underlying clusters based on the consistency of intra- and inter-cluster maps, inspired by the observation that maps tend to be more consistent along cycles within a cluster than across clusters. This discrepancy has been observed in many different contexts, and appears to be a consequence of the energy landscape of almost all optimization-based pairwise matching algorithms. The matching energy functional between objects in the same underlying cluster tends to have simple energy landscapes with easily identifiable global optimums, resulting in fairly cycle-consistent intra-cluster maps; in contrast, the highly non-convex energy landscape for dissimilar objects from different clusters leads to more “random” maps due to random initialization and/or multiple strong local minimums, for which cycle-consistency is much less often observed. This map consistency argument is the foundation of our *simultaneous mapping and clustering* (SMAC) algorithm.

1.1. Motivating Example

We validate the map consistency argument through a motivating example (see Figure 1) on a real dataset from SHREC07 Watertight benchmark (Giorgi et al., 2007). This dataset consists of 38 shapes: the first 18 are Human models and the remaining 20 are Fourleg models (e.g., Horses and Dogs). Each shape is represented as a discrete metric space with 1024 sample points generated from farthest-point sampling (Eldar et al., 1997). We compute pairwise *blended intrinsic maps* (Kim et al., 2011) for all objects in this collection and use these maps to compute two similarity scores for each object pair: a *map distortion score* that measures the squared sum of geodesic distortions across all point-pairs (c.f. (Bronstein et al., 2006)), and a *cycle-consistency score* that is the median value (among other options) of the distortion scores of all 3-cycles to which the pair belongs, where the distortion of a 3-cycle is defined as the squared sum of the geodesic distances between each point and its image propagated along the 3-cycle.

Figure 1 illustrates the distributions of the map distortion scores (Left) and the cycle-consistency scores (Right) on the 38 models. The cycle-consistency scores clearly reveal the

underlying cluster structure and in fact better separates the two clusters of models (Human vs. Fourleg) than the map-distortion scores (intra-cluster blocks in the right figure are darker in blue). The superior cluster separation is verified by comparing the results of spectral clustering (Ng et al., 2002; Lei & Rinaldo, 2015) using the two similarity scores: spectral clustering based on the cycle-consistency scores recovers the two underlying clusters perfectly, whereas the same procedure using the map distortion scores incorrectly puts two Fourleg models in the cluster of Human models. This motivating example illustrates the effectiveness and superiority of the map consistency score as a quantification of object similarity.

1.2. Approach Overview

Motivated by the example above, we propose an algorithm for *simultaneous mapping and clustering* (SMAC). Our SMAC algorithm takes as input (i) an object collection that falls into multiple clusters, and (ii) some noisy maps pre-computed between object pairs, and outputs the underlying clusters together with improved maps between all pairs of objects. Our SMAC algorithm builds upon the equivalence between map consistency and the low-rank property of a data matrix with consistent maps in its blocks (c.f. (Huang & Guibas, 2013)). We show that this low-rank property still holds in the setting of multiple disjoint collections of consistent intra-cluster maps, though the rank is expected to be higher due to multiple clusters. Based on this observation, the first step of our approach simultaneously recovers the underlying clusters and intra-cluster maps by spectral decomposition. We show that properly “rounding off” the leading eigenvectors recovers the ground-truth clusters and intra-cluster maps in a single pass. We then construct inter-cluster maps from the recovered intra-cluster maps. Our theoretical analysis establishes sharp exact recovery conditions for both steps under a fairly general noise model, using some novel tight L^∞ -stability bounds for eigen-decompositions.

1.3. Related Works

Joint object matching, i.e., simultaneously estimating maps among a collection of objects, is an emerging field across many scientific problems. Earlier works use combinatorial optimizations (Nguyen et al., 2011; Kim et al., 2012; Huang et al., 2012). More recent works (Huang & Guibas, 2013; Huang et al., 2014; Chen et al., 2014; Wang & Singer, 2013) rely on convex or non-convex optimization techniques. However, all these methods assume that the underlying object collection is *homogeneous* (all objects belong to a single category). For a *heterogeneous* object collection where the objects fall into multiple distinctive clusters, existing methods usually rarely succeed in producing both high-quality intra- and inter-cluster maps.

Clustering and in particular spectral clustering is another well-studied topic. We refer to (Lei & Rinaldo, 2015; Rohe et al., 2011) for some recent advances and to (Filippone et al., 2008; Luxburg, 2007; Fortunato, 2010) for surveys. Our approach falls into the general category of graph-based clustering, but the pairwise information we utilize is of “functional” rather than “scalar” nature. Instead of the more common approach that derives affinity scores from the pairwise maps for clustering, our SMAC algorithm discovers the cluster structure based purely on the consistency of pairwise maps and demonstrates improved empirical performance. This strategy is reminiscent of *heterogeneous multireference alignment* (Boumal et al., 2017) and *simultaneous alignment and classification* (Lederman & Singer, 2016) for *synchronization problems* in Cryo-Electron Microscopy; in this context, our residue-based clustering strategy is closest in nature to *learning group actions* (Gao et al., 2016).

Our approach relies on tight L^∞ -type bounds on leading eigenvectors of perturbed matrices. Though the stability of eigenvalues and eigenspaces are well-studied, element-wise eigenvector stability appears to be a much harder problem; see recent survey (O’Rourke et al., 2016). We introduce new stability bounds to tackle this technical difficulty.

1.4. Mathematical Notation

We use lower bold letters $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{u}, \mathbf{v}, \mathbf{w}, \dots$ to denote vectors, and upper letters A, B, C, \dots for matrices. For a block matrix $X \in \mathbb{R}^{n_1 m_1 \times n_2 m_2}$, we use $X_{ij} \in \mathbb{R}^{m_1 \times m_2}$ to denote its ij -th block; the ij -th element of a matrix A is denoted as a_{ij} . With \otimes we denote the Kronecker product. For a symmetric matrix $A \in \mathbb{R}^{n \times n}$, we always sort the eigenvalues in non-increasing order, i.e., $\lambda_1(A) \geq \dots \geq \lambda_n(A)$. Matrix norms $\|\cdot\|_{\mathcal{F}}, \|\cdot\|_1, \|\cdot\|_2$, and $\|\cdot\|_\infty$ will be used for a matrix $A \in \mathbb{R}^{n_1 \times n_2}$, of which $\|\cdot\|_2 = \sigma_{\max}(A)$ is the maximum singular value of A , and the spectral norm $\|\cdot\|_2$ is sometimes simplified as $\|\cdot\|$. We denote \mathcal{P}_m for the set of permutation matrices of dimension $m \times m$.

2. Algorithm

For simplicity, we focus on describing and analyzing our algorithm under the setting where pair-wise maps are given by permutations. In Section 4, we show how to modify the algorithm for partially similar objects.

We first describe the problem setup. Consider n objects $\mathcal{S} = \{S_1, \dots, S_n\}$ each represented by m points. With $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ we denote an observation graph among \mathcal{S} . An initial map $X_{ij}^{in} \in \mathcal{P}_m$ is pre-computed on each edge $(i, j) \in \mathcal{E}$ using an off-the-shelf pairwise object matching algorithm from S_i to S_j . We also assume the objects in \mathcal{S} are partitioned into $k \geq 2$ clusters, but k is unknown. Our goal is to identify the underlying clusters, and in the mean-

Algorithm 1 PermSMAC: Simultaneously mapping and clustering

Input: Observation graph $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ and initial pairwise maps $X_{ij}^{in}, (i, j) \in \mathcal{E}$

Output: Underlying clusters $\mathcal{S} = c_1 \cup \dots \cup c_k$ and optimized pairwise maps $X_{ij}, 1 \leq i, j \leq n$

- 1: **{Step 1}** Simultaneously compute the intra-cluster maps and extract the underlying clusters:
- 2: **{Step 1.1}** Form data matrix based on (1).
- 3: **{Step 1.2}** Compute the critical value $r = \operatorname{argmax}_{2 \leq i \leq nm} \frac{\lambda_i - \lambda_{i+1}}{\lambda_i + \lambda_{i+1}}$.
- 4: **{Step 1.3}** Let $U \in \mathbb{R}^{nm \times r}$ store the leading r eigenvectors of X . Compute pair-wise maps X_{ij}^* by solving (2)
- 5: **{Step 1.4}** Use $f_{ij}(X_{ij}^*)$ as the affinity score and apply single-linkage clustering to obtain the underlying clusters
- 6: **{Step 2}** compute the inter-cluster maps by solving (6)

while improve all pairwise maps between objects in \mathcal{S} . As a basis for identifying the underlying clusters, we assume that the intra-cluster maps are more accurate (in terms of cycle-consistency) than inter-cluster maps.

Our algorithm proceeds in two major steps. The first step simultaneously extracts the underlying clusters and computes intra-cluster maps. The second step then computes inter-cluster maps. Now we introduce these two steps in details. Algorithm 1 provides the pseudo code.

2.1. Step I: Extract underlying clusters and compute intra-cluster maps.

Motivated from prior works for map synchronization (Pachauri et al., 2013; Shen et al., 2016), we use a block data matrix $X \in \mathbb{R}^{nm \times nm}$ to encode the input maps:

$$X_{ij} = \begin{cases} (X_{ij}^{in} - \frac{1}{m} \mathbf{1}\mathbf{1}^T) & (i, j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Our approach is motivated by the empirical observation that the leading eigen-vectors of X reveal the underlying clusters and simultaneously denoise intra-cluster maps if intra-cluster maps are more accurate than inter-cluster maps. We provide the algorithmic details below; an analysis is provided in Section 3.

Given the data matrix, we first estimates the number of stable eigen-vectors as

$$r = \operatorname{argmax}_{m \leq i \leq nm} \frac{\lambda_i - \lambda_{i+1}}{|\lambda_i| + |\lambda_{i+1}|}.$$

Here we search within the range of $[m, nm]$, as we expect multiple underlying clusters. Let $U \in \mathbb{R}^{nm \times r}$ store the

top r eigen-vectors of X , and divide U into n matrices U_1, \dots, U_n of shape $m \times r$ such that $U = (U_1^T, \dots, U_n^T)^T$. We then compute the estimated map X_{ij}^* along each edge $(i, j) \in \mathcal{E}$ by solving a linear assignment problem, i.e.,

$$X_{ij}^* = \min_{X \in \mathcal{P}_m} f_{ij}(X), \quad f_{ij}(X) := \|X \cdot U_i - U_j\|_{\mathcal{F}}^2 \quad (2)$$

for all $1 \leq i, j \leq n$. Note that (2) admits an exact solution via linear assignment. In fact,

$$X_{ij}^* = \operatorname{argmin}_{X \in \mathcal{P}_m} \|XU_i - U_j\|_{\mathcal{F}}^2 \quad (3)$$

$$= \operatorname{argmin}_{X \in \mathcal{P}_m} (\|U_i\|_{\mathcal{F}}^2 + \|U_j\|_{\mathcal{F}}^2 - 2\langle XU_i, U_j \rangle) \quad (4)$$

$$= \operatorname{argmax}_{X \in \mathcal{P}_m} \langle X, U_j U_i^T \rangle. \quad (5)$$

Intuitively, $U_j U_i^T$ provides an approximation of the underlying map X_{ij} , and the linear assignment procedure projects this approximation onto the space of permutations.

For clustering, we treat the residual score $f_{ij}(X_{ij}^*)$, $1 \leq i, j \leq n$ as the distance measure between S_i and S_j , and apply single-linkage clustering (Gower & Ross, 1969) to obtain the underlying clusters. We set $k = \lfloor \frac{r}{m-1} \rfloor$ as the number of desired clusters.

Empirically, we found that when the input inter-cluster maps are inaccurate, the quality of estimated inter-cluster maps appear to be much more noisy than estimated intra-cluster maps. This motivates us to re-estimate inter-cluster maps as a second step, described in Section 2.2 below.

2.2. Step II: Estimate inter-cluster maps.

We estimate the inter-cluster maps between each pair of clusters independently. Specifically, consider two clusters c_s and c_t . Let $S_{i_s} \in c_s$ and $S_{i_t} \in c_t$ be a pair of objects selected from each cluster, respectively. We optimize the inter-cluster map X_{st}^{inter} , represented as a pairwise object map $X_{i_s i_t}$, by solving the following linear assignment:

$$\begin{aligned} X_{st}^{\text{inter}} &= \operatorname{argmin}_{X \in \mathcal{P}_m} \sum_{i \in c_s, j \in c_t, (i,j) \in \mathcal{E}} \|X - X_{j i_t} X_{ij}^{\text{in}} X_{i i_s}\|_1 \\ &= \operatorname{argmax}_{X \in \mathcal{P}_m} \sum_{i \in c_s, j \in c_t, (i,j) \in \mathcal{E}} \langle X, X_{j i_t} X_{ij}^{\text{in}} X_{i i_s} \rangle. \end{aligned} \quad (6)$$

Note that it is possible to jointly optimize X_{st}^{inter} among all pairs of clusters. However, since the number of clusters is usually significantly smaller than the size of each cluster, we found the gain of doing so is insignificant.

3. Analysis

We first describe our noise model in Section 3.1. We then analyze our method under this model and present the exact

recovery conditions for both underlying clusters and the pairwise maps in Section 3.2. Our analysis is based on a set of new stability bounds of eigen-decompositions. The proofs of all Lemma's and Theorem's with technical details can be referred to in the supplemental material.

3.1. Model for Analysis

We consider two models, one for the pairwise map and the other one for the observation graph and the underlying cluster structure.

Map model. We generalize the map model described in (Shen et al., 2016) to multiple clusters: Suppose there are k underlying clusters. With $c_s, 1 \leq s \leq k$ we denote the vertex indices of the s -th cluster. In other words, we have $\{1, \dots, n\} = c_1 \cup c_2 \cup \dots \cup c_k$. Given an observation graph, the input pairwise maps are independent, and they follow

$$X_{ij}^{\text{in}} = \begin{cases} I_m & \text{with probability } \eta_{ij} \\ U_{\mathcal{P}_m} & \text{with probability } 1 - \eta_{ij} \end{cases}$$

where $U_{\mathcal{P}_m}$ denotes a random permutation matrix satisfying

$$E[U_{\mathcal{P}_m}] = \frac{1}{m} \mathbf{1} \mathbf{1}^T. \quad (7)$$

η_{ij} depends on the edge type: $\eta_{ij} = p$ if (i, j) is an intra-cluster edge, i.e., $(i, j) \in \mathcal{E} \cap (\cup_{1 \leq s \leq k} c_s \times c_s)$, and $\eta_{ij} = q$ if (i, j) is an inter-cluster edge, i.e., $(i, j) \in \mathcal{E} \cap (\cup_{1 \leq s \neq t \leq k} c_s \times c_t)$. We assume $p > q$.

Remark 1. Note that one can also generalize the model by assuming there exist underlying permutations $P_i, 1 \leq i \leq n$, so that the ground-truth map $X_{ij} = P_j P_i^T$. Nevertheless, it turns out that the two models are identical (Shen et al., 2016). Hence we adopt the simpler form for convenience.

Model for the observation graph and clusters. To obtain concise and interpretable exact recovery conditions, we are particularly interested in analyzing our algorithm when the observation graph and underlying clusters are generated from the following established model: assume $n = n_0 k$, and the size of each underlying cluster $|c_s| = n_0, 1 \leq s \leq k$; the observation graph is generated from the Erdős-Rényi $G(n, t)$ with edge selection probability t . However, the stability bounds we introduce in the supplemental material can be used for analyzing more general noisy models, e.g., sizes of the underlying clusters are different.

3.2. Map and Cluster Recovery

We begin by analyzing the leading eigenvalues and eigenvectors of $E[X]$ to gain insights on the necessary conditions for map and cluster recovery. To make our discussion more general, we first assume the underlying cluster and the observation graph are fixed. Consider then the following (p, q) -

reweighted normalized adjacency matrix $\bar{A}(p, q) \in \mathbb{R}^{n \times n}$:

$$(\bar{A}(p, q))_{ij} = \begin{cases} p & (i, j) \in \mathcal{E} \cap \cup_{1 \leq s \leq m} (c_s \times c_s), \\ q & (i, j) \in \mathcal{E} \cap \cup_{1 \leq s \neq t \leq m} (c_s \times c_t). \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

It is clear that

$$E[X] = \bar{A}(p, q) \otimes (I_m - \frac{1}{m} \mathbf{1}\mathbf{1}^T)$$

and thus the non-zero eigenvalues of $E[X]$ are non-zero eigenvalues of $\bar{A}(p, q)$ with multiplicity $m-1$. Furthermore, let $(\frac{1}{\sqrt{m}}\mathbf{1}, H_m)$ be an orthonormal basis for \mathbb{R}^m . Then the leading $k(m-1)$ eigenvectors of $E[X]$ are given by $S_k \otimes H_m$. This leads two conditions on the eigenvalues and eigenvectors of $\bar{A}(p, q)$ for map and cluster recovery:

- **Eigenvalue separation.** Since our method leverages the largest eigengap, we assume that $\lambda_k(\bar{A}(p, q)) - \lambda_{k+1}(\bar{A}(p, q))$ has the largest eigengap. Define

$$\gamma = \frac{\max_{1 \leq i \neq k \leq n-1} \lambda_i(\bar{A}(p, q)) - \lambda_{i+1}(\bar{A}(p, q))}{\lambda_k(\bar{A}(p, q)) - \lambda_{k+1}(\bar{A}(p, q))},$$

Then a necessary condition for map recovery is $\gamma < 1$.

- **Eigenvector separation.** We further assume that the underlying clusters can be recovered by reading off the rows of S_k . Formally, consider rows of S_k as coordinates of the corresponding objects, and define

$$d_{\text{intra}} = \max_{1 \leq s \leq k} \max_{i, j \in c_s} \|(e_i - e_j)^T S_k\|, \quad (9)$$

$$d_{\text{inter}} = \min_{1 \leq s < t \leq k} \min_{i \in c_s, j \in c_t} \|(e_i - e_j)^T S_k\|. \quad (10)$$

d_{intra} and d_{inter} essentially measure the maximum distance within each cluster and the minimum distance between different clusters, respectively. Thus, a necessary condition for cluster recovery is that $d_{\text{intra}} < \mu \cdot d_{\text{inter}}$ for some small constant μ .

Under these two conditions, it is easy to see that when $X \approx E[X]$, we have $U_i \approx (e_i^T S_k) \otimes H_m$ and $U_j U_i^T \approx (e_j^T S_k S_k^T e_i)(I_m - \frac{1}{m} \mathbf{1}\mathbf{1}^T)$. It follows that both the underlying clusters and intra-cluster maps can be exactly recovered.

These two separation conditions are quite general. In fact, they hold for the noisy model described in Section 3.1. To gain some further insight, one can show that (c.f. (Le et al., 2017)) with high probability

$$\lambda_1(\bar{A}(p, q)) = \begin{cases} q + \frac{p-q}{k} + O(\frac{1}{\sqrt{nt}}) & i = 1 \\ \frac{p-q}{k} + O(\frac{1}{\sqrt{nt}}) & 2 \leq i \leq k \\ O(\frac{1}{\sqrt{nt}}) & k+1 \leq i \leq n \end{cases}$$

indicating that $\gamma = o(1)$. Moreover, under this model

$$\bar{A}(p, q) \approx (p - q)I_k \otimes (\mathbf{1}\mathbf{1}^T) + q(\mathbf{1}\mathbf{1}^T).$$

If p and q are well-separated, the top k eigenvectors of $\bar{A}(p, q)$ approximate $I_k \otimes \mathbf{1}$, meaning $d_{\text{intra}} \approx 0$ and $d_{\text{inter}} \approx 1$. Now we formally state the exact recovery conditions:

Theorem 3.1. (Intra-Cluster and Map Recovery) *Assume $t = \Omega(\frac{\log(n)}{n})$. Consider the noise model described in Section 3.1. There exists an absolute constant c_{intra} such that PermSMAC recovers the underlying intra-cluster maps and the underlying clusters with high probability if*

$$p - q \geq c_{\text{intra}} k \sqrt{\frac{\log(n)}{nt}}. \quad (11)$$

Remark 2. Note that the gap between p and q is used to ensure the recovery of the underlying clusters. Moreover, the recovery rate matches the information theoretic lower bound established in (Chen et al., 2016) up to $O(\sqrt{\log(n)})$, indicating the tightness of our condition for PermSMAC.

The following theorem provides an exact recovery condition for the inter-cluster maps. Compared with the previous lower bound, the lower bound on q for inter-cluster map recovery is significantly cruder. This shows the advantage of recovering inter-cluster maps as a separate step.

Theorem 3.2. *There exists an absolute constant $c_{\text{inter}} > 0$, so that when*

$$q \geq c_{\text{inter}} k \sqrt{\frac{\log(n)}{n^2 t}},$$

PermSMAC recovers the underlying inter-cluster maps with high probability.

4. Partial Matching

In this section we extend the algorithm to handle partially (as opposed to fully) similar objects. Each input object S_i ($1 \leq i \leq n$) in this setting has $m_i \in \mathbb{N}_+$ points, where the m_i 's vary across the collection. Consequently, the pairwise maps $X_{ij}^{in} \in \{0, 1\}^{m_j \times m_i}$ are no longer permutation matrices since some rows and columns may only contain zero elements. We propose the following modified algorithm for SMAC in this partial matching setting.

Step I: Extract underlying clusters and compute intra-cluster maps. Forming the data matrix and leading eigenvector computation stay the same as in the full matching setting, except we replace $X_{ij}^{in} - \frac{1}{m} \mathbf{1}\mathbf{1}^T$ by X_{ij}^{in} in (1). The first difference occurring in the partial matching setting is that we cannot apply (2) to obtain pair-wise maps and affinity scores for clustering. Our strategy is to apply single linkage clustering on the rows of U . Specifically, the distance measure between two points p, p' is given by

$\|\mathbf{u}_p - \mathbf{u}_{p'}\|^2$, where \mathbf{u}_p is the corresponding row of p in U . The number of output clusters is set as r . Each output cluster of this single-linkage procedure collects a set of matched points among the input objects. We merge two clusters if the objects they belong to overlap. Suppose we finally obtain k clusters c_1, \dots, c_k . For each cluster c_i , we introduce a binary matrix $Y_i \in \{0, 1\}^{n_i \times r_i}$, whose columns encode the enclosed point clusters. Then it is easy to see that the blocks of $Y_i Y_i^T$ describe intra-cluster maps. Note that $r_i = m$ in the full setting, but in the partial setting r_i is usually larger than $\max_{1 \leq i \leq n} m_i$, due to partial similarity.

Step II: Compute inter-cluster maps. In the partial setting, we encode the inter-cluster map from cluster c_s to cluster c_t as a matrix $X_{st} \in \{0, 1\}^{r_s \times r_t}$. Consider a object pair $(i, j) \in \mathcal{E}$, where $i \in c_s$ and $j \in c_t$. With $E_{i,s}$ and $E_{j,t}$ we denote the index matrices that extract the corresponding blocks in Y_s and Y_t . It is easy to see that the entries $Y_t^T E_{j,t} X_{ij}^{in} E_{i,s}^T Y_s$ provide cues for the inter-cluster map X_{st} . Similar to the full map setting, we compute

$$C_{st} = \sum_{i \in c_s, j \in c_t, (i,j) \in \mathcal{E}} Y_t^T E_{j,t} X_{ij}^{in} E_{i,s}^T Y_s.$$

Since the inter-cluster maps may not be permutation matrices either, we apply a simple thresholding to obtain the inter-cluster maps:

$$X_{st} = C_{st} > \beta_{st},$$

where β_{st} is set as 0.9 times the maximum element of C_{st} in our experiments.

5. Experimental Results

In this section, we evaluate our approach on both synthetic (Section 5.1) and real datasets (Section 5.2 and Section 5.3). For baseline comparison, we consider state-of-the-art approaches for clustering and joint mapping in each domain.

5.1. Experimental Results on Synthetic Instances

We apply the model described in Section 3 to generate the synthetic data sets for our experiments. Below we summarize how the procedure depends on the model parameters:

- **The observation graph \mathcal{G} .** We employ a standard two-community stochastic block model (Abbe et al., 2016) which enables us to fully control the vertex degree and the spectral gap. We use this model to generate three observation graphs $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$. All of them have $n = 300$ vertices, but the vertex degrees and spectral gaps vary. Specifically, \mathcal{G}_1 is the clique graph. \mathcal{G}_2 is a sparse graph, whose average vertex degree is 50 and the spectral gap is 0.1. The average vertex degree and spectral gap of \mathcal{G}_3 are 50 and 0.5, respectively. In

our experiment, we treat \mathcal{G}_1 as the default observation graph. We also study the influence of the observation graphs on the performance of our algorithm.

- **Number of clusters k .** Without loss of generality, we allocate each object into a underlying cluster with probability $\frac{1}{k}$. For each observation graph, we generate and fix one underlying cluster throughout our experiments.
- **Other parameters m, γ, p, q .** We fix the number of points on each object as $m = 30$. In the partial matching setting, we follow the protocol (Chen et al., 2014) to generate the input objects so that the expected size of each object is $m\gamma$. We sample the ratio of correct inter-cluster maps $q = \exp(-\frac{i}{10})$, $15 \leq i \leq 50$. Since $p > q$, we sample the ratio of correct intra-cluster maps so that $p - q = \frac{i}{100}$, $1 \leq i \leq 30$.

We now study the empirical phase transition curves when varying p and q under different k, γ and observation graphs.

Varying k . The first two rows of Figure 2 show the phase transition curves of map and cluster recovery for $k = 2, 4, 6$. Across all configurations, our approach tolerates a significant portion of noise in the input maps. The fraction of noise that our approach can handle reduces as k increases, which is consistent with the exact recovery conditions in Section 3. In addition, phase transitions with respect to mapping recovery and cluster recovery roughly align. The subtle differences are two-fold: when p and q are close, cluster recovery breaks as there is no cue for clustering; likewise, map recovery breaks when q approaches 0.

Versus mapping only. Figure 2 compares our approach to state-of-the-art map recovery technique (Huang & Guibas, 2013). SMAC clearly exhibits a clear advantage in the regime, where q is small and p is significantly larger than q . This is expected through our exact recovery condition.

Varying observation graph. Figure 3 shows phase transition curves of map and cluster recovery when varying the observation graphs ($k = 2, \gamma = 1$). Our approach tolerates a larger fraction of incorrect maps for larger vertex degrees. Moreover, when vertex degrees are comparable, a small spectral gap means higher recovery rate.

Varying γ . Figure 4 shows the phase transition curves when varying the overlapping ratio γ . We again show three configurations, i.e., $\gamma = 1, \gamma = 0.8$ and $\gamma = 0.6$. Still, our approach can tolerate a large rate of noise in the input maps. Moreover, the rate reduces as γ becomes smaller. This is expected, as low overlapping ratio introduces weak signal for mapping and cluster recovery.




Figure 2. Phase transition curves when varying q and $p - q$ under different configurations of k . The first and second rows show mapping and cluster recovery of PermsMAC, respectively. The third row shows mapping recovery of (Huang & Guibas, 2013). (a) $k = 2$. (b) $k = 4$. (c) $k = 6$. Red means more success.

5.2. Experimental Results on 3D Shapes

We proceed to evaluate SMAC on 3D shapes. We consider two datasets for this task. The first dataset collects four categories of 3D shapes from SHREC07-Watertight (Giorgi et al., 2007), namely, Human, Fourleg, Armadillo and Teddy. These categories appear to have both similar global structures and local geometric details. However, the inter-category variability is salient. The second dataset is more fine-grained. It has 10 underlying shape collections from FAUST training dataset (Bogo et al., 2014), where each collection consists of different poses of the same human subject (10 poses per collection). For evaluating shape maps, we follow the protocol of (Kim et al., 2011) by collecting statistics on the geodesic distortion of predicted correspondences with respect to human annotated feature correspondences.

Mapping performance. Figure 5(c) plots the accuracy of predicted correspondences of our approach versus the input. For a detailed assessment, we separate the statistics of intra-cluster maps and inter-cluster maps. We consider two approaches for baseline comparison: the first applies (Huang & Guibas, 2013) to the entire dataset, and the second applies (Huang & Guibas, 2013) to each category in isolation then applies the third step of our approach to compute inter-cluster maps. The second baseline may be considered as a performance upper bound. Our proposed SMAC is significantly better than mapping without clustering (which is seriously affected by the noise in inter-cluster maps) and competitive against the second baseline.

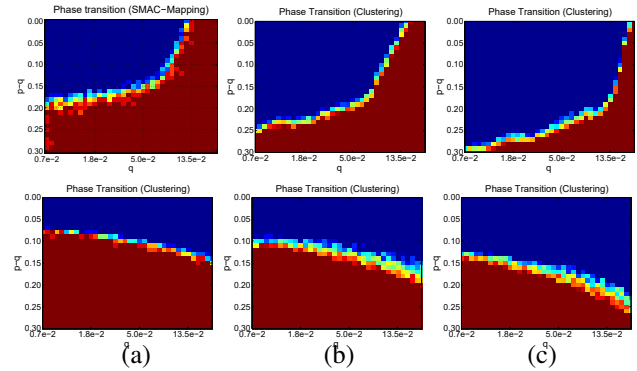


Figure 3. Phase transition curves of map recovery and cluster recovery when varying the observation graph. The top row shows mapping recovery, and the bottom row shows cluster recovery. (a) \mathcal{G}_1 . (b) \mathcal{G}_2 . (c) \mathcal{G}_3 .




Figure 4. Phase transition curves of map recovery and cluster recovery when varying the overlapping ratio. The top row shows mapping recovery, and the bottom row shows cluster recovery. (a) $\gamma = 1$. (b) $\gamma = 0.8$. (c) $\gamma = 0.6$.

Clustering performance. As shown in Table 1, our approach correctly identifies all underlying clusters in SHREC07 and FAUST. We also applied two baseline clustering approaches on the same dataset. The first approach performs k -means on the spectral shape descriptors (Rustamov, 2007). This approach only yields 84.6% and 72.0%, respectively. The second approach utilizes the mapping distortion as an affinity measure and applies spectral clustering. This approach yields 94.9% and 74.0%, respectively, which are better than the first baseline. However, our approach is still better, which shows the advantage of using the cycle-consistency constraint for clustering.

5.3. Experimental Results on 2D Images

Finally, we evaluate our approach on two datasets of 2D images. The first dataset (Figure 6(Left)) consists of 600 internet images of Notre Dame. These images naturally fall into 3 categories, each of which collects images from a similar camera pose (Snaveely et al., 2006). The second dataset (Figure 6(Right)) collects 600 internet images of 4 landmark churches in Europe (Amiens Cathedral (200 im-

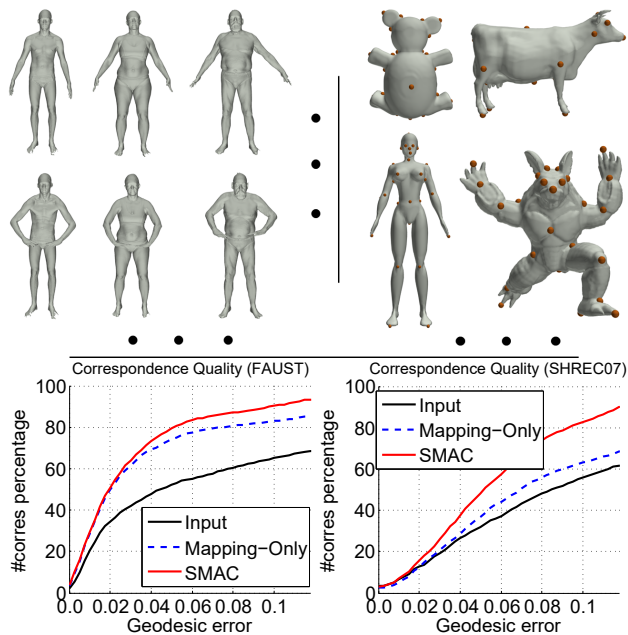


Figure 5. Experimental results on 3D shapes. We consider two datasets: (left) Human/Fourleg/Armadillo/Teddy, and (right) Different poses of xxx human subjects. For each dataset, we show (top) snapshots of each dataset, and (bottom) error distribution of predicted intra-cluster and inter-cluster shape maps.

	SHREC07	FAUST	Notre Dame	Church
SMAC	100%	94.0%	99.3%	96.1%
Distort.	94.9%	74.0%	94.3%	91.9%
Descrip.	84.6%	72.0%	84.3%	86.7%

Table 1. Classification accuracy of our method and baseline approaches on the four datasets show in Figure 5 and Figure 6. Distort. and Descrip. stand for using the pairwise mapping distortion score and global object descriptor for clustering, respectively.

ages), York Minster (200 images), Duomo (100 images) and Westminster Abbey (100 images)). As inter-cluster maps do not make much sense here, we only evaluate clustering results and intra-cluster maps in this experiment. We sample 400 SIFT features (Lowe, 2004) for each image and apply SIFT flow (Liu et al., 2011) to establish pairwise correspondences between the features. We manually mark feature correspondences for evaluation.

Mapping performance. Figure 6 compares our approach with the two baseline approaches introduced in Section 5.2. The relative performance is consistent. Specifically, due to the small-overlapping region across different clusters, inter-cluster maps are rather noisy, so applying joint-mapping directly leads to sub-optimal results. In addition, our approach is competitive against the approach of computing intra-cluster and inter-cluster maps in a sequential manner.

Clustering performance. Finally, we evaluate our approach in terms of clustering accuracy. We choose two base-

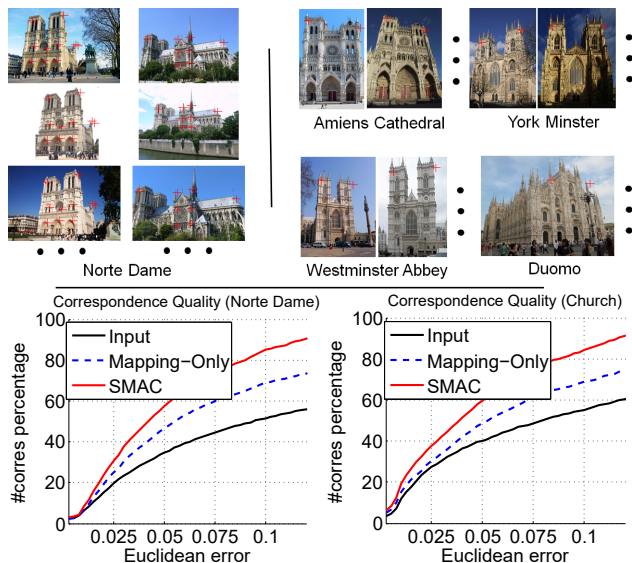


Figure 6. Experimental results on 2D images. We consider two datasets: (Left) Internet images of Norte Dame from three diverse camera poses, and (Right) Internet images of five churches from similar view points. Visualization is the same as Figure 5.

line approaches, where the first baseline approach performs k -means on image descriptors. In this case, we employ GIST (Oliva & Torralba, 2001). The second baseline uses the residual of SIFT flow as the affinity score for clustering. As shown in Table 1, our approach leads to a clustering accuracy of 99.3% and 96.1% on Notre Dame and Church, respectively. They are higher than those of the top performing baselines, i.e., 94.5% and 92.1%, respectively.

6. Conclusions

We have introduced SMAC for simultaneously computing consistent maps across a heterogeneous data collection and identifying the underlying clusters. The key idea is to leverage the higher self-consistency within intra-cluster maps than inter-cluster maps. Enforcing this variation of consistency allows us to denoise the input maps in a sequential manner while simultaneously identifying the underlying cluster structures. The approach is based on spectral decomposition, for which we provided tight exact recovery conditions for both the input maps and the underlying clusters. Experimental results on synthetic data sets justify our exact recovery conditions, and experimental results on real data sets demonstrate the efficacy of our approach.

Acknowledgement. This research by Chandrajit Bajaj was supported in part by a grant from the NIH, R01 GM117594. Tingran Gao acknowledges support from Simons Math+X Investigators Award 400837, DARPA D15AP00109, and NSF IIS 1546413. Qixing Huang would like to acknowledge support for this research from NSF DMS-1700234, NSF CI-P-1729486, and NSF IIS-1618648.

References

- Abbe, E., Bandeira, A. S., and Hall, G. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory*, 62(1):471–487, 2016.
- Bogo, F., Romero, J., Loper, M., and Black, M. J. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Piscataway, NJ, USA, June 2014. IEEE.
- Boumal, N., Bendory, T., Lederman, R. R., and Singer, A. Heterogeneous multireference alignment: a single pass approach. *arXiv preprint arXiv:1710.02590*, 2017.
- Boyer, D. M., Lipman, Y., St. Clair, E., Puente, J., Patel, B. A., Funkhouser, T., Jernvall, J., and Daubechies, I. Algorithms to Automatically Quantify the Geometric Similarity of Anatomical Surfaces. *Proceedings of the National Academy of Sciences*, 108(45):18221–18226, 2011. doi: 10.1073/pnas.1112822108.
- Bronstein, A. M., Bronstein, M. M., and Kimmel, R. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. In *PNAS*, volume 103, pp. 1168–1172, 2006.
- Chen, Y., Guibas, L. J., and Huang, Q. Near-optimal joint object matching via convex relaxation. In *ICML*, pp. 100–108, 2014.
- Chen, Y., Suh, C., and Goldsmith, A. J. Information recovery from pairwise measurements. *IEEE Trans. Information Theory*, 62(10):5881–5905, 2016.
- Eldar, Y., Lindenbaum, M., Porat, M., and Zeevi, Y. Y. The farthest point strategy for progressive image sampling. *IEEE Trans. Image Processing*, 6(9):1305–1315, 1997.
- Filippone, M., Camastra, F., Masulli, F., and Rovetta, S. A survey of kernel and spectral methods for clustering. *Pattern Recogn.*, 41(1):176–190, January 2008.
- Fortunato, S. Community detection in graphs. *Physics Reports*, 486(3-5):75 – 174, 2010.
- Gao, T., Brodzki, J., and Mukherjee, S. The Geometry of Synchronization Problems and Learning Group Actions. *arXiv preprint arXiv:1610.09051*, 2016.
- Gao, T., Yapuncich, G. S., Daubechies, I., Mukherjee, S., and Boyer, D. M. Development and Assessment of Fully Automated and Globally Transitive Geometric Morphometric Methods, With Application to a Biological Comparative Dataset With High Interspecific Variation. *The Anatomical Record*, pp. n/a–n/a, 2017. ISSN 1932-8494. Submitted.
- Giorgi, D., Biasotti, S., and Paraboschi, L. Shape retrieval contest 2007: Watertight models track, 2007.
- Gower, J. C. and Ross, G. J. S. Minimum spanning trees and single linkage cluster analysis (ref: 69v18 p106-110). *Applied Statistics*, 18:54–64, 1969.
- Huang, Q. and Guibas, L. J. Consistent shape maps via semidefinite programming. *Comput. Graph. Forum*, 32(5):177–186, 2013.
- Huang, Q., Koltun, V., and Guibas, L. Joint shape segmentation with linear programming. *ACM Trans. Graph.*, 30(6):125:1–125:12, December 2011.
- Huang, Q., Zhang, G., Gao, L., Hu, S., Butscher, A., and 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.
- Huang, Q., Wang, F., and Guibas, L. J. Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.*, 33(4):36:1–36:11, 2014.
- Huang, X., Liang, Z., Bajaj, C., and Huang, Q. Translation synchronization via truncated least squares. In *NIPS*, 2017.
- Huber, D. F. and Hebert, M. Fully automatic registration of multiple 3d data sets. *Image and Vision Computing*, 21: 637–650, 2001.
- Kim, V., Li, W., Mitra, N., DiVerdi, S., and Funkhouser, T. Exploring collections of 3d models using fuzzy correspondences. *ACM Trans. Graph.*, 31(4):54:1–54:11, July 2012.
- Kim, V. G., Lipman, Y., and Funkhouser, T. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4):79:1–79:12, July 2011.
- Kraevoy, V. and Sheffer, A. Cross-parameterization and compatible remeshing of 3d models. In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, pp. 861–869, 2004.
- Le, C. M., Levina, E., and Vershynin, R. Concentration and regularization of random graphs. *Random Struct. Algorithms*, 51(3):538–561, 2017.
- Lederman, R. R. and Singer, A. A representation theory perspective on simultaneous alignment and classification. *arXiv preprint arXiv:1607.03464*, 2016.
- Lei, J. and Rinaldo, A. Consistency of spectral clustering in stochastic block models. *Annals of Statistics*, 43(1): 215–237, February 2015.
- Leonardos, S., Zhou, X., and Daniilidis, K. Distributed consistent data association via permutation synchronization. In *ICRA*, pp. 2645–2652. IEEE, 2017.

- Liu, C., Yuen, J., and Torralba, A. Sift flow: Dense correspondence across scenes and its applications. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):978–994, 2011.
- Lowe, D. G. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, nov 2004.
- Luxburg, U. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, December 2007.
- Mandad, M., Cohen-Steiner, D., Kobbelt, L., Alliez, P., and Desbrun, M. Variance-minimizing transport plans for inter-surface mapping. *ACM Transactions on Graphics*, 36(4), 2017.
- Ng, A. Y., Jordan, M. I., and Weiss, Y. On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems*, pp. 849–856, 2002.
- Nguyen, A., Ben-Chen, M., Welnicka, K., Ye, Y., and Guibas, L. J. An optimization approach to improving collections of shape maps. *Comput. Graph. Forum*, 30(5):1481–1491, 2011.
- Oliva, A. and Torralba, A. Modeling the shape of the scene: A holistic representation of the spatial envelope. *Int. J. Comput. Vision*, 42(3):145–175, May 2001.
- O’Rourke, S., Vu, V., and Wang, K. Eigenvectors of random matrices: A survey. *J. Comb. Theory, Ser. A*, 144:361–442, 2016.
- Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., and Guibas, L. J. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4): 30:1–30:11, 2012.
- Pachauri, D., Kondor, R., and Singh, V. Solving the multi-way matching problem by permutation synchronization. In *NIPS*, pp. 1860–1868, 2013.
- Pauly, M., Mitra, N. J., Giesen, J., Gross, M., and Guibas, L. J. Example-based 3d scan completion. In *Proceedings of the Third Eurographics Symposium on Geometry Processing*, SGP ’05, 2005.
- Rohe, K., Chatterjee, S., and Yu, B. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, August 2011.
- Rustamov, R. M. Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, SGP ’07, pp. 225–233, 2007.
- Schreiner, J., Asirvatham, A., Praun, E., and Hoppe, H. Inter-surface mapping. In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH ’04, pp. 870–877, 2004.
- Shen, Y., Huang, Q., Srebro, N., and Sanghavi, S. Normalized spectral map synchronization. In *NIPS*, pp. 4925–4933, 2016.
- Snively, N., Seitz, S. M., and Szeliski, R. Photo tourism: Exploring photo collections in 3d. *ACM Trans. Graph.*, 25(3):835–846, July 2006.
- van Kaick, O., Zhang, H. R., Hamarneh, G., and Cohen-Or, D. A survey on shape correspondence. In *EuroGraphics: State-of-the-Art Report*, pp. 61–82, 2010.
- Wang, F., Huang, Q., and Guibas, L. J. Image co-segmentation via consistent functional maps. In *ICCV*, pp. 849–856, 2013.
- Wang, F., Huang, Q., Ovsjanikov, M., and Guibas, L. J. Unsupervised multi-class joint image segmentation. In *CVPR*, pp. 3142–3149. IEEE Computer Society, 2014.
- Wang, L. and Singer, A. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2:145193, December 2013.
- Zhou, T., Lee, Y. J., Yu, S. X., and Efros, A. A. Flowweb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In *CVPR*, pp. 1191–1200. IEEE Computer Society, 2015a.
- Zhou, X., Zhu, M., and Daniilidis, K. Multi-image matching via fast alternating minimization. *CoRR*, abs/1505.04845, 2015b.