

# Multi-Image Semantic Matching by Mining Consistent Features

Qianqian Wang<sup>†</sup>      Xiaowei Zhou<sup>†\*</sup>      Kostas Daniilidis<sup>‡</sup>

<sup>†</sup> State Key Lab of CAD&CG, Zhejiang University

<sup>‡</sup> GRASP Laboratory, University of Pennsylvania

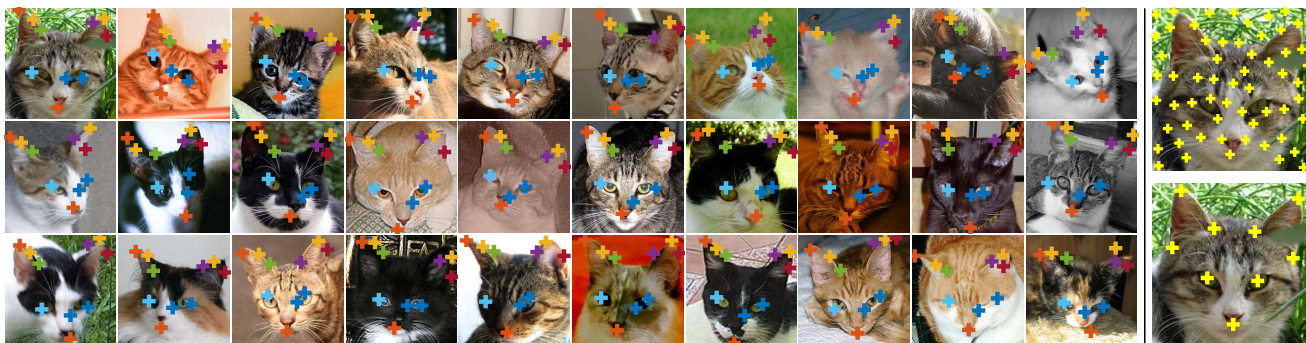


Figure 1: Given initial feature candidates in multiple images and their noisy correspondences between each pair of images, the proposed method identifies a sparse set of reliable feature points and establishes cyclically and geometrically consistent correspondences across all images. The figure gives an example of identified features points (colored crosses) from 1000 cat head images (only 30 images are shown here). The color indicates the correspondence. The last column shows initial feature candidates (top) and manually-annotated landmarks provided by the dataset (bottom), both in the first image. Interestingly, the unsupervisedly-identified feature points by the proposed method roughly coincide with the manually-annotated landmarks.

## Abstract

*This work proposes a multi-image matching method to estimate semantic correspondences across multiple images. In contrast to the previous methods that optimize all pairwise correspondences, the proposed method identifies and matches only a sparse set of reliable features in the image collection. In this way, the proposed method is able to prune nonrepeatable features and also highly scalable to handle thousands of images. We additionally propose a low-rank constraint to ensure the geometric consistency of feature correspondences over the whole image collection. Besides the competitive performance on multi-graph matching and semantic flow benchmarks, we also demonstrate the applicability of the proposed method for reconstructing object-class models and discovering object-class landmarks from images without using any annotation.*

## 1. Introduction

Computing feature correspondences across images is a fundamental problem in computer vision. Low-level geometric features (e.g., SIFT [30]) are successful for matching images of the same scene. Recently, there has been an increasing interest in semantic matching (e.g., [28, 22]), i.e., establishing semantic correspondences across different object instances or scenes. Most research on semantic matching focuses on the pairwise case that considers only image pairs. Finding consistent correspondences across multiple images is important in many situations, e.g., object-class model reconstruction [21] and automatic landmark annotation [38]. The multi-image semantic matching problem is the focus of this work.

Despite remarkable advances in solving semantic matching and multi-image matching problems (see related work section), the following challenges remain. First, repeatable feature point detection for semantic matching is an open problem [25, 38]. Previous work bypassed this issue by either using all pixels (dense flow) [28] or randomly sampled points [41], resulting in numerous nonre-

\*Corresponding author. E-mail: xzhou@cad.zju.edu.cn

repeatable features that have no real correspondences in other images. Second, previous multi-image matching methods (e.g., [35, 54]) mainly optimize the cycle consistency of correspondences and seldom consider the geometric consistency simultaneously. While there have been effective ways to enforce geometric constraints in the pairwise setting (e.g., RANSAC [15] and graph matching [26]), few solutions exist for the multi-image case. Finally, the existing multi-image matching methods are computationally expensive, which could hardly process hundreds of images. Analyzing large datasets requires more scalable algorithms.

In most situations one only needs the correspondences of a sparse set of highly repeatable features which are cyclically and geometrically consistent across images. Dense correspondences could be achieved by interpolation. Therefore, in contrast to the previous multi-image matching methods that optimize all pairwise correspondences, we formulate the problem as a feature selection and labeling problem: Starting from fussy pairwise correspondences, we aim to select a sparse set of feature points from the initial set of candidates in each image, and establish their correspondences across images by assigning labels to them. The selection and labeling are accomplished by optimizing both cycle consistency and geometric consistency of selected features. Formulating the problem in this way allows us to 1) explicitly deal with nonrepeatable feature points in the initial feature sets and 2) dramatically decrease the number of variables, resulting in a scalable algorithm that is able to jointly analyze thousands of images. Finally, inspired by classic results on factorization-based structure from motion [39], we propose a low-rank constraint that enforces the geometric consistency for matching multiple images and is very efficient to solve in optimization. Figure 1 gives an example illustrating our problem and the proposed method.

The main contributions of this work are summarized as follows:

- We propose a novel approach to solving the multi-image semantic matching problem as a feature selection and labeling problem. The proposed algorithm is able to discover consistent features in an image collection and is scalable to handle thousands of images.
- We introduce a novel low-rank constraint for multi-image matching that allows the proposed algorithm to optimize cycle consistency and geometric consistency simultaneously.
- We demonstrate the competitive performance of the proposed method on standard benchmarks. We also show two applications: 1) reconstruct 3D object-class models from images of different instances *without using any manual annotation* and 2) match 1000 cat head images and interestingly find that the automat-

ically selected feature points represent very discriminative landmarks on eyes, ears and mouths, which demonstrates the potential applicability of the proposed method to automatic landmark annotation.

## 2. Related work

**Image matching:** In classical image matching, sparse feature correspondences between images are estimated using low-level geometric feature detectors (e.g., corners and covariant regions [32]) and descriptors (e.g., SIFT [30], SURF [1] and HoG [17]). The geometric consistency is imposed by either using RANSAC [15] as a postprocessing step or solving a graph matching problem that minimizes the geometric distortion between images [26, 8]. Many recent works attempt to find semantic correspondences across different scenes [28, 37]. Hierarchical matching [22] and region-based strategies [16] have been proposed to make use of high-level semantics in images.

**Learning detectors and descriptors:** Recent results (e.g., [14, 29]) show that the deep features extracted from convolutional neural networks (CNNs) are very effective in matching and outperform handcrafted features even if the CNNs are not trained for matching. Supervised learning has been used to explicitly learn descriptors. The supervision is from manually annotated correspondences [9], images transformations [34], and additional cues and data, e.g., silhouettes [20] and CAD models [52]. Meanwhile, there are few efforts towards learning feature detectors that are repeatable and covariant to transformations [49, 25, 38]. The proposed method can be viewed as an unsupervised approach to harvesting reliable features and consistent correspondences from image collections, which may provide training data for detector and descriptor learning.

**Multi-image matching:** The proposed method is technically related to the joint matching methods [23, 19, 35]. Most existing methods aim to make use of the cycle consistency to improve the pairwise correspondences. Various approaches have been proposed such as unclosed cycle elimination [50, 33], constrained local optimization [46, 45, 47, 44, 53], spectral relaxation [23, 18, 35] and convex relaxation [18, 6, 54]. The proposed method differs from them as it aims to identify the most consistent features instead of optimizing all pairwise correspondences, making it more scalable for exploring large datasets. In addition, we introduce a low-rank constraint to optimize the geometric consistency of selected features. The matrix decomposition method for multi-graph matching proposed in [48] imposes the low-rank constraint on graph edges. In this work the constraint is directly imposed on feature locations resulting in more efficient optimization. The recent work [40] proposes an efficient method to discover clusters of discriminative features for matching, but no geometric constraint is considered.

### 3. Preliminaries and notation

#### 3.1. Pairwise matching

Given  $n$  images to match and  $p_i$  feature points in each image  $i$ , the pairwise feature correspondences for each image pair  $(i, j)$  can be represented by a partial permutation matrix  $P_{ij} \in \{0, 1\}^{p_i \times p_j}$ , which satisfies the doubly stochastic constraints:

$$\mathbf{0} \leq P_{ij} \mathbf{1} \leq \mathbf{1}, \mathbf{0} \leq P_{ij}^T \mathbf{1} \leq \mathbf{1} \quad (1)$$

$P_{ij}$  can be estimated by maximizing the inner product between itself and the feature similarities subject to the constraints in (1). This is a linear assignment problem, which is well-studied and can be solved by the Hungarian algorithm. Finding  $P_{ij}$  can also be formulated as a graph matching problem, which can be cast as a quadratic assignment problem (QAP). Specifically, an objective function encoding both local compatibilities (feature similarity) and structural compatibilities (spatial rigidity) is maximized in order to find the assignment. Although QAP is NP-hard, many effective algorithms have been proposed to solve it approximately, e.g., [2, 8, 26]. We will use the output of linear matching or graph matching, denoted by  $W_{ij} \in \mathbb{R}^{p_i \times p_j}$ , as our input.

#### 3.2. Cycle consistency

Recent works [6, 35, 45] propose to use cycle consistency as a constraint to match multiple images. The correspondence between all pairs of images is cyclically consistent if the following equation holds for any triplet of images  $(i, j, z)$ :

$$P_{ij} = P_{iz} P_{zj} \quad (2)$$

The cycle consistency can be described more concisely by introducing a virtual ‘‘universe’’ which is defined as the set of unique features that appear in the image collection [35, 18]. Each feature point in the universe must be observed by at least one image and matched to corresponding image points. Suppose the underlying correspondence between image  $i$  and the universe is denoted by partial permutation matrix  $X_i \in \{0, 1\}^{p_i \times u}$ , where  $u$  is the size of the universe and  $u \geq p_i$  for all  $i$ . The pairwise correspondence  $P_{ij}$  can be represented as  $X_i X_j^T$ .

If the permutation matrices are concatenated as

$$P = \begin{bmatrix} P_{11} & P_{12} & \cdots & P_{1n} \\ P_{21} & P_{22} & \cdots & P_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ P_{n1} & P_{n2} & \cdots & P_{nn} \end{bmatrix}, \quad X = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix}, \quad (3)$$

it has been shown that the set  $\{P_{ij} | \forall i, j\}$  is cyclically consistent if and only if  $P$  can be factorized as  $XX^T$  [26, 18].

### 4. Proposed methods

#### 4.1. Matching by labeling

Recall that  $X \in \{0, 1\}^{m \times u}$  is the map from image features to the universe, where  $m$  and  $u$  denote the total number of local features in the image collection and the size of universe, respectively. Another interpretation of  $X$  is that each row vector of  $X$  is the label of the corresponding feature. The image features with identical labels match each other. To accommodate all image features, previous work [6, 54] usually defines a sufficiently large  $u$ .

However, not all features that appear in the image collection are desirable for matching. Particularly, in semantic matching most of the randomly or uniformly sampled features are nonrepeatable across images and should be excluded during matching. Inspired by this, we select the most repeatable image features and map them to a more compact feature space containing only  $k$  elements, where  $k$  is a predefined small value meaning the number of selected features in each image.

Suppose the correspondences between the feature points in image  $i$  and the selected feature space is represented by  $X_i \in \{0, 1\}^{p_i \times k}$ . Each  $X_i$  is a partial permutation matrix with a small number of columns which satisfies

$$\mathbf{0} \leq X_i \mathbf{1} \leq \mathbf{1}, X_i^T \mathbf{1} = \mathbf{1} \quad (4)$$

The sum of each column in  $X_i$  equals to 1, meaning that every element in the selected feature space should correspond to a feature point in each image. On the contrary, the sum of a row could be zero, meaning that the corresponding feature point is not selected.

The set  $\{X_i | 1 \leq i \leq n\}$  is what we need to estimate.  $X_i X_j^T$  gives the pairwise correspondences between selected features in image  $i$  and  $j$ , which must be cyclically consistent by construction as explained in Section 3.2. As we attempt to identify the discriminative features that are supposed to produce more cyclically consistent correspondences in the initial pairwise matching, we minimize the discrepancy between the initial pairwise matching results and the constructed ones to estimate  $X$ :

$$\begin{aligned} \min_X \quad & \frac{1}{4} \|W - XX^T\|_F^2 \\ \text{s.t.} \quad & X_i \in \mathbb{P}^{p_i \times k}, 1 \leq i \leq n \end{aligned} \quad (5)$$

where  $\mathbb{P}$  denotes the set of partial permutation matrices and  $W \in \mathbb{R}^{m \times m}$  is the collection of  $W_{ij}$ . By solving (5), the most repeatable features in the image collection will be selected and matched in a cyclically consistent way.

#### 4.2. Geometric constraint

Suppose we have tracked  $k$  features over  $n$  frames in a scene. We use  $M_i \in \mathbb{R}^{2 \times k}$  to denote the coordinates of

the  $k$  ordered features of frame  $i$ , and concatenate all  $\tilde{M}_i$  as rows in a matrix  $\tilde{M} \in \mathbb{R}^{2n \times k}$  with each column per feature.  $\tilde{M}$  is known as the *measurement matrix* in structure from motion [39]. It can be shown that under orthographic projection,  $\tilde{M}$  is of rank 4.

In our problem, we represent the coordinates of all feature candidates in image  $i$  by  $C_i \in \mathbb{R}^{2 \times p_i}$ . Then, the coordinates of selected feature points in image  $i$  are given by

$$\tilde{M}_i = C_i X_i \quad (6)$$

where  $\tilde{M}_i \in \mathbb{R}^{2 \times k}$  stores the coordinates of selected feature points that are aligned in the same order as the selected feature space. Similarly, we could concatenate all  $\tilde{M}_i$  as rows in a matrix  $\tilde{M} \in \mathbb{R}^{2n \times k}$ :

$$\tilde{M} = \begin{bmatrix} C_1 X_1 \\ \vdots \\ C_n X_n \end{bmatrix} \quad (7)$$

If feature points are correctly selected and labeled,  $\tilde{M}$  will become a measurement matrix of rank 4 under orthographic projection. Even if the scene is non-rigid,  $\tilde{M}$  can still be approximated by a low-rank matrix [3]. This conclusion can be effectively utilized to better estimate  $X$ . Suppose the groundtruth rank of  $\tilde{M}$  is no larger than  $r$ . Minimizing the following term allows us to impose geometric consistency on selected feature points:

$$f_{\text{geo}} = \frac{1}{2} \|\tilde{M} - Z\|_F^2 = \frac{1}{2} \sum_{i=1}^n \|C_i X_i - Z_i\|_F^2 \quad (8)$$

where  $Z \in \mathbb{R}^{2n \times k}$  is an auxiliary variable whose rank is no larger than  $r$ , and  $Z_i \in \mathbb{R}^{2 \times k}$  denotes the  $(2i - 1)$ -th and  $2i$ -th rows of  $Z$ .

### 4.3. Formulation

Combining the cycle consistency and geometric consistency terms discussed in Section 4.1 and 4.2, we obtain the final optimization problem:

$$\begin{aligned} \min_{X, Z} \quad & \frac{1}{4} \|\mathbf{W} - \mathbf{X} \mathbf{X}^T\|_F^2 + \frac{\lambda}{2} \sum_{i=1}^n \|C_i X_i - Z_i\|_F^2 \\ \text{s.t.} \quad & X_i \in \mathbb{P}^{p_i \times k}, 1 \leq i \leq n \\ & \text{rank}(Z) \leq r \end{aligned} \quad (9)$$

where  $\lambda$  controls the weight of the geometric constraint.

### 4.4. Optimization

We propose to solve the optimization problem in (9) by block coordinate descent, i.e., alternately updating one variable while fixing the others.

We attempted to relax the integer constraint of (9) and treated  $X$  as a real matrix  $X \in [0, 1]^{m \times k}$ , which is a common practice to solve quadratic assignment problems. However, we observed that, if the integer constraint on  $X$  is relaxed, the effect of the geometric constraint is negligible as the system  $C_i X_i = Z_i$  is ill-posed for arbitrary  $Z_i$ . Therefore, we keep the integer constraint on  $X$ . To make the optimization tractable, we decouple the two terms in (9) by replacing  $X$  in the first term with an auxiliary variable  $Y \in \mathbb{R}^{m \times k}$  and rewrite the optimization as:

$$\begin{aligned} \min_{X, Y, Z} \quad & \frac{1}{4} \|\mathbf{W} - \mathbf{Y} \mathbf{Y}^T\|_F^2 + \frac{\lambda}{2} \sum_{i=1}^n \|C_i X_i - Z_i\|_F^2 \\ & + \frac{\rho}{2} \|X - Y\|_F^2 \\ \text{s.t.} \quad & X_i \in \mathbb{P}^{p_i \times k}, 1 \leq i \leq n \\ & Y \in \mathcal{C} \\ & \text{rank}(Z) \leq r \end{aligned} \quad (10)$$

where  $\mathcal{C}$  denotes the set of matrices satisfying the following constraints:

$$0 \leq Y \leq 1, 0 \leq Y_i \mathbf{1} \leq 1, Y_i^T \mathbf{1} = 1, 1 \leq i \leq n \quad (11)$$

and  $\rho$  is a parameter controlling the degree of similarity between  $X$  and  $Y$ . When  $\rho$  approaches infinity, the problem in (10) is equivalent to the original one (9).

The motivation for rewriting the problem as (10) is that each subproblem in the block coordinate descent will be much easier to solve. We alternately update  $Y$ ,  $X$ ,  $Z$  in the following manner.

$Y$  is updated via projected gradient descent [36, 31]:

$$Y \leftarrow \Pi_{\mathcal{C}}[Y - \eta(Y Y^T Y - W Y + \rho(Y - X))] \quad (12)$$

where  $\Pi_{\mathcal{C}}$  denotes projection onto  $\mathcal{C}$  and  $\eta > 0$  is the step-size. We update  $Y$  until convergence before updating  $X$  and  $Z$ .

Each  $X_i$  is updated via the Hungarian algorithm, whose cost matrix is constructed as

$$H_i = \lambda D(C_i, Z_i) - 2\rho Y_i, \quad (13)$$

where  $D(C_i, Z_i) \in \mathbb{R}^{p_i \times k}$  denotes the squared Euclidean distance between each pair of observations in  $C_i$  and  $Z_i$ .

$Z$  is updated via singular value decomposition (SVD):

$$Z = U \tilde{\Sigma} V^T, \quad (14)$$

where the columns of  $U$  and the columns of  $V$  are the left and right singular vectors of  $\tilde{M}$  respectively, and  $\tilde{\Sigma}$  is a diagonal matrix with the diagonal elements corresponding to the  $r$  largest singular values of  $\tilde{M}$ .

For better convergence, we use an increasing sequence of  $\rho$  to enforce the geometric constraint gradually. For each



value of  $\rho$ , we update the variables alternately until the objective in (10) does not decrease. As each update never increases the objective, the local convergence is guaranteed. In our experiment, we set the sequence of  $\rho$  as (1, 10, 100),  $\lambda$  as 1 and  $r$  as 4 empirically.

As the optimization is nonconvex and involves both continuous and discrete variables, a reliable initialization is necessary. We first initialize  $\mathbf{Y}$  by ignoring the geometric constraint and solving

$$\begin{aligned} \min_{\mathbf{Y}} \quad & \frac{1}{4} \|\mathbf{W} - \mathbf{Y}\mathbf{Y}^T\|_F^2 \\ \text{s.t.} \quad & \mathbf{Y} \in \mathcal{C} \end{aligned} \tag{15}$$

using the projected gradient descent as (12) with  $\rho = 0$ .  $\mathbf{X}$  can be initialized by discretizing  $\mathbf{Y}$ .

## 5. Experiments

### 5.1. Multi-graph matching

We first validate the effectiveness of the proposed optimization algorithm in the setting of multi-graph matching, where the feature locations are annotated but their correspondences need to be estimated. The matching accuracy is evaluated by the recall, which is defined as the number of true correspondences found by the algorithm divided by the number of groundtruth correspondences.

We use the CMU datasets and the WILLOW Object Class dataset for evaluation. The CMU datasets contain the hotel sequence (111 frames) and the house sequence (101 frames). SIFT descriptors [30] are extracted at 30 feature point annotations provided by [4] in each frame. The WILLOW Object Class dataset [7] provides images of five object classes (Car, Duck, Motorbike, Face, Winebottle) and 10 annotated points corresponding to several discriminative parts of each class. Each class contains at least 40 images with different instances. As the object appearance in each class varies greatly, geometric descriptors like SIFT can hardly work. Instead, we adopt the deep features extracted from pretrained convolutional neural networks, which have proven to be effective in previous work [29]. Specifically, each image is fed through the AlexNet [24] (pretrained on ImageNet [11]) and the feature map responses of Conv4 and Conv5 corresponding to each landmark are extracted and concatenated as its descriptor. For both datasets, the initial pairwise correspondences are obtained from the linear matching solver Hungarian algorithm and then fed into the proposed algorithm. Three alternative methods with publicly available code are used as baselines: the spectral method [35], MatchLift [6] and MatchALS [54]. For all methods, the size of universe is set as the number of annotations in each image.

The recall rates are reported in Table 1, which shows that the proposed method outperforms other methods on

Dataset	Input	[35]	[6]	[54]	Ours <sup>-</sup>	Ours	Input <sup>+</sup>	Ours <sup>+</sup>
Hotel	0.57	0.53	0.64	0.58	0.63	0.90	0.85	1
House	0.74	0.74	0.79	0.75	0.79	0.93	0.95	1
Car	0.48	0.55	0.66	0.65	0.72	0.75	0.83	1
Duck	0.43	0.59	0.56	0.56	0.63	0.77	0.65	0.88
Face	0.86	0.92	0.93	0.94	0.95	0.95	0.99	1
Motorbike	0.30	0.25	0.28	0.27	0.40	0.61	0.85	1
Winebottle	0.52	0.64	0.71	0.72	0.73	0.82	0.92	1

Table 1: The recall rates on the CMU datasets and the WILLOW Object Class dataset. The proposed method is compared with spectral method [35], MatchLift [6] and MatchALS [54]. Ours<sup>-</sup> represents the recall rates without the geometric constraint. Input<sup>+</sup> and Ours<sup>+</sup> represent the initial and optimized recall rates respectively when graph matching is applied to obtain the initial pairwise correspondences.




Figure 2: The matching results with and without the geometric constraint are shown in bottom and top rows, respectively. The true matches and false matches are shown in blue and red, respectively.

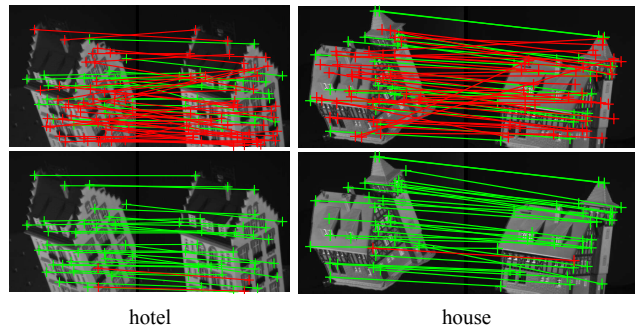


Figure 3: The ability to remove outliers. The true matches and false matches are shown in green and red, respectively. The top and bottom rows correspond to the results of pairwise matching and our joint matching method, respectively. Besides 30 annotated feature points, we introduce 30 randomly located points as outliers in each frame. We set  $k = 30$  in our method.

all datasets. Table 1 also demonstrates another two cases: 1) If the graph matching solver RRWM [8] is leveraged to improve the initial pairwise correspondences, the matching accuracy of the proposed method can achieve 100% on all datasets except the duck. 2) If the geometric consistency is ignored, the matching accuracy reduces significantly. Two

Methods	car(S)	car(G)	car(M)	duc(S)	mot(S)	mot(G)	mot(M)	win(w/o C)	win(w/C)	win(M)	Avg.
LOM [16] + Ours	0.89	0.62	0.56	0.70	0.49	0.31	0.28	0.91	0.52	0.72	0.60
LOM [16]	0.86	0.58	0.52	0.65	0.48	0.28	0.28	0.91	0.37	0.65	0.56
DeepFlow [37]	0.33	0.13	0.22	0.20	0.20	0.08	0.13	0.46	0.08	0.18	0.20
GMK [13]	0.48	0.25	0.34	0.27	0.31	0.12	0.15	0.41	0.17	0.18	0.27
SIFT Flow [28]	0.54	0.37	0.36	0.32	0.41	0.20	0.23	0.83	0.16	0.33	0.38
DSP [22]	0.46	0.30	0.32	0.25	0.31	0.15	0.14	0.85	0.25	0.64	0.37
Zhou <i>et al.</i> [52]	0.77	0.34	0.52	0.42	0.34	0.19	0.20	0.78	0.19	0.38	0.41

Table 2: PCK ( $\alpha = 0.1$ ) for dense flow on the PF-WILLOW dataset (SS w/HOG).

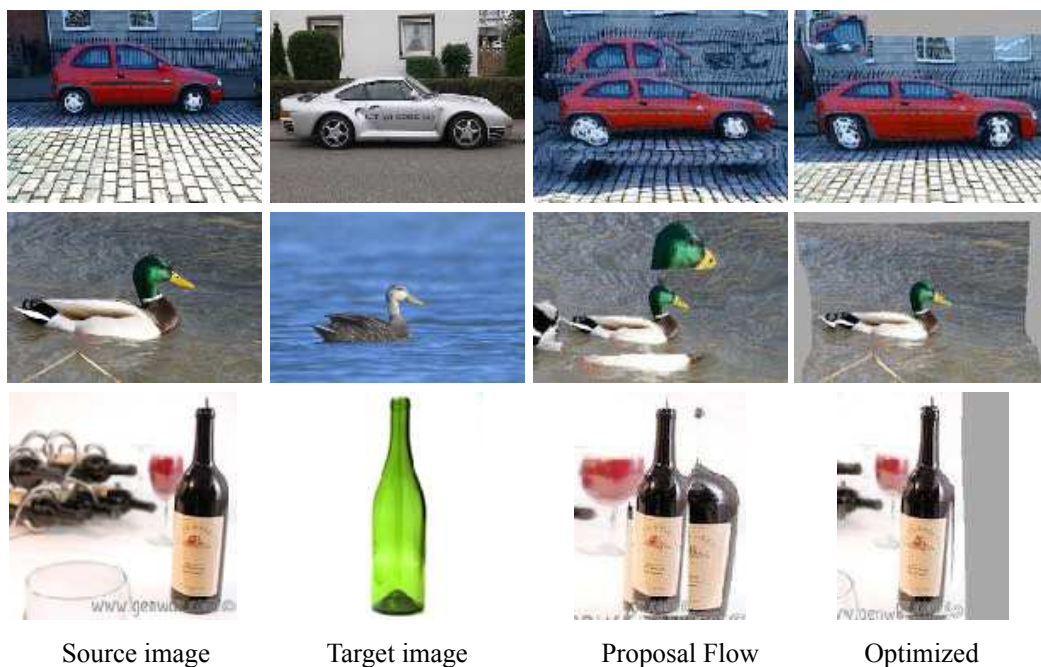


Figure 4: Examples of dense flow. The source images are warped to the target images using the dense correspondences estimated by proposal flow [16] and optimized by the proposed method.

sample image pairs are visualized in Figure 2 which shows that geometrically distorted matches can be corrected after enforcing the geometric consistency.

The proposed method can automatically select reliable features for matching. Figure 3 gives an example where randomly located feature points are added to images in the CMU datasets as outliers. It is shown that the outliers can effectively be pruned and moreover, the correspondences between the selected feature points are improved.

## 5.2. Dense semantic matching

In this section, we show the application of the proposed method in dense matching by combining it with region based semantic flow methods, e.g., proposal flow [16]. In the proposal flow method, the correspondences of region

proposals between images are estimated and then transformed into a dense flow field. For a collection of images, we apply the proposed method on top of proposal flow to improve the estimated pairwise correspondences of proposals, thus improving the dense flow.

We experiment with a benchmark for evaluating semantic flow techniques named the PF-WILLOW dataset [16], which splits the WILLOW Object Class dataset into 10 sub-classes. They are car (S), (G), (M), duck (S), motorbike (S), (G), (M), winebottle (w/oC), (w/C), (M), where (S) and (G) represent side and general viewpoints, respectively, (C) denotes background clutter, and (M) denotes mixed viewpoints. Each sub-class includes 10 images of different object instances.

The percentage of correct keypoints (PCK) is used as

the evaluation metric [16]. It evaluates the percentage of correctly located keypoints when transferring the annotated keypoints from an image to another image using the estimated flow. A predicted feature point is deemed to be correctly located if it lies within  $\alpha \max(h, w)$  pixels from the groundtruth point for  $\alpha$  in  $[0, 1]$ , where  $h$  and  $w$  are the height and width of the object bounding box, respectively. For proposal flow, the selective search (SS) [42] is used as proposal generator, HOG [10] as feature descriptors, and local offset matching (LOM) [16] as geometric matching strategy. 500 proposals are extracted in each image and used for matching and generating dense flow. In our algorithm, each proposal is treated as a feature point, and the center of each proposal is regarded as its coordinates in our geometric constraint. The number of selected features in all sub-classes is set as 10.

The result in Table 2 shows that our method improves the results of original proposal flow on most of the classes. A qualitative example is given in Figure 4.

### 5.3. Object-class model reconstruction

Matching images of different object instances is a main challenge in reconstructing object-class models from images. Some previous works [43, 5, 21] rely on annotated keypoints in images. The recent work [54] requires no keypoint annotation but object masks to remove background. We show that the proposed method can produce consistent correspondences for reconstruction without using any manual annotation. We demonstrate with the FG3DCar dataset [27], match all left-view sedan images (37 in total), and reconstruct a 3D model. In addition, we collect another dataset containing 30 images of different motorbikes with similar views.

Similar to [54], we uniformly sample feature candidates on image edges detected by the structured forests [12]. Unlike [54], we don't need object masks thanks to the capability of the proposed method to prune nonrepeatable features in the background. On average,  $\sim 550$  feature candidates are obtained for each image. The deep features described in Section 5.1 are used as descriptors and the graph matching solver RRWM [8] is adopted for initial pairwise matching.

To illustrate the effect of selection, we use the precision as our metric, which is defined as the number of true correspondences divided by the total number of correspondences found by the algorithm. The definition of true correspondence is similar to that in PCK. We vary the number of selected features and compare the precisions of recovered correspondences in Figure 5. It is shown that the proposed method achieves obvious improvements compared to the original pairwise matching, whose precision is low in the presence of background clutter. Moreover, the fewer the features we select, the higher the precision will be. This justifies the use of selection, which prunes nonrepeatable




Figure 5: The precision for different numbers of selected features on the FG3DCar dataset.




Figure 6: The matches between two sedan images. The true matches and false matches are shown in blue and red, respectively. The top and bottom rows correspond to the results of pairwise matching and the proposed method, respectively. Note that most of the initial feature points in the background are pruned by the proposed method.

features such as the ones in the background. Only the reliable feature points near the objects are selected and matched consistently. An example is visualized in Figure 6.

For reconstruction, we simply run affine reconstruction through the factorization method [39]. The selected feature points, correspondences and reconstructions are visualized in Figure 7. Clearly, most of the selected feature points are located on objects and correctly matched despite the large variety in object appearances and viewpoints. In spite of some noises and missing points, we can see the structures of the sedan and the motorbike from the reconstructions. It is believed that more sophisticated reconstruction techniques can be adopted to obtain better reconstructions. Quantitatively, we evaluate the estimated relative rotations between all pairs of images with the ground truth provided by the FG3DCar dataset. The mean error in geodesic distance is  $18.5^\circ$ . To our knowledge, no previous result exists for unsupervised relative pose estimation between different instances.






Figure 7: Matching sedans and motorbikes. Only four images are selected and shown for each image set. Note that the instances are different. All the feature points are automatically sampled and selected by the proposed method. The markers with the same color indicate the matched points. The 3D reconstruction is rendered with the colors in the first image and visualized in two viewpoints.

#### 5.4. Automatic landmark annotation

We apply the proposed algorithm to the first 1000 images from the cat head dataset [51]. Similar to the previous experiment, the feature candidates are sampled from detected edges in images, yielding  $\sim 43$  candidates per image on average. We set the number of selected features to be 10. The results are shown in Figure 1. As the figure shows, while initial candidates distribute randomly over the whole image including the background, the selected features are all on the objects with correct correspondences established across very different instances with a variety of appearances and poses. More interestingly, the automatically selected features roughly coincide with human annotations provided by the dataset, representing the discriminative parts of cat such as ears, eyes and mouth. This demonstrates the potential of the proposed method for automatic landmark annotation, which imitates humans’ annotation process: we compare a collection of images and find a set of parts that are invariant in appearance and geometry across images.

#### 5.5. Computational complexity

The bottleneck restricting the scalability of the proposed method is the matrix multiplication when updating  $\mathbf{Y}$ , while the other update steps only involve much smaller matrices. In (12), the dominant part is  $\mathbf{Y}\mathbf{Y}^T\mathbf{Y}$  which takes  $O(mk^2)$  flops for each update, where  $m$  and  $k$  are the total number of features in all images and the number of selected features in each image, respectively. As a comparison, the complexity of MatchALS [54] is  $O(m^2k)$ . As  $k$  is much smaller than

$m$ , the proposed method is much more scalable. For example,  $m \approx 43,000$  and  $k = 10$  in the cat head experiment in Section 5.4. We implement the proposed algorithm in Matlab on a PC with an Intel i7 3.4GHz CPU and 16G RAM. The CPU time for the cat head experiment is  $\sim 650s$  regardless of pairwise matching, which can hardly be solved by the previous multi-image matching algorithms.

## 6. Conclusion

We presented a novel method that solved the problem of semantic matching across multiple images as a feature selection and labeling problem. The proposed method could establish reliable feature correspondences among a collection of images which satisfy both cycle consistency and geometric consistency. Experiments showed that the proposed method outperformed the previous multi-image matching methods while being highly scalable to match thousands of images. Several applications were demonstrated: improving dense flow estimation on top of the proposal flow method, reconstructing object-class models without using any manual annotation, and automatically annotating image landmarks in 1000 cat head images.

**Acknowledgements:** The authors are grateful for support through the following grants: ARL MAST-CTA W911NF-08-2-0004, ARL RCTA W911NF-10-2-0016, ONR N00014-17-1-2093, and Fundamental Research Funds for the Central Universities (No. 2018FZA5011).



## References

- [1] H. Bay, T. Tuytelaars, and L. Van Gool. Surf: speeded up robust features. In *ECCV*, 2006. 2
- [2] A. C. Berg, T. L. Berg, and J. Malik. Shape matching and object recognition using low distortion correspondences. In *CVPR*, 2005. 3
- [3] C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3d shape from image streams. In *CVPR*, 2000. 4
- [4] T. S. Caetano, J. J. McAuley, L. Cheng, Q. V. Le, and A. J. Smola. Learning graph matching. *T-PAMI*, 31(6):1048–1058, 2009. 5
- [5] J. Carreira, A. Kar, S. Tulsiani, and J. Malik. Virtual view networks for object reconstruction. In *CVPR*, 2015. 7
- [6] Y. Chen, L. Guibas, and Q. Huang. Near-optimal joint object matching via convex relaxation. In *ICML*, 2014. 2, 3, 5
- [7] M. Cho, K. Alahari, and J. Ponce. Learning graphs to match. In *ICCV*, 2013. 5
- [8] M. Cho, J. Lee, and K. M. Lee. Reweighted random walks for graph matching. In *ECCV*, 2010. 2, 3, 5, 7
- [9] C. B. Choy, J. Gwak, S. Savarese, and M. Chandraker. Universal correspondence network. In *NIPS*, 2016. 2
- [10] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005. 7
- [11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [12] P. Dollár and C. L. Zitnick. Structured forests for fast edge detection. In *ICCV*, 2013. 7
- [13] O. Duchenne, A. Joulin, and J. Ponce. A graph-matching kernel for object categorization. In *ICCV*, 2011. 6
- [14] P. Fischer, A. Dosovitskiy, and T. Brox. Descriptor matching with convolutional neural networks: a comparison to sift. *arXiv preprint arXiv:1405.5769*, 2014. 2
- [15] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 2
- [16] B. Ham, M. Cho, C. Schmid, and J. Ponce. Proposal flow: Semantic correspondences from object proposals. *T-PAMI*, 2017. 2, 6, 7
- [17] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 2
- [18] Q.-X. Huang and L. Guibas. Consistent shape maps via semidefinite programming. *Computer Graphics Forum*, 32(5):177–186, 2013. 2, 3
- [19] Q.-X. Huang, G.-X. Zhang, L. Gao, S.-M. Hu, A. Butscher, and L. Guibas. An optimization approach for extracting and encoding consistent maps in a shape collection. *ACM Transactions on Graphics*, 31(6):167, 2012. 2
- [20] A. Kanazawa, D. W. Jacobs, and M. Chandraker. WarpNet: Weakly supervised matching for single-view reconstruction. In *CVPR*, 2016. 2
- [21] A. Kar, S. Tulsiani, J. Carreira, and J. Malik. Category-specific object reconstruction from a single image. In *CVPR*, 2015. 1, 7
- [22] J. Kim, C. Liu, F. Sha, and K. Grauman. Deformable spatial pyramid matching for fast dense correspondences. In *CVPR*, 2013. 1, 2, 6
- [23] V. G. Kim, W. Li, N. J. Mitra, S. DiVerdi, and T. A. Funkhouser. Exploring collections of 3d models using fuzzy correspondences. *ACM Transactions on Graphics*, 31(4):54, 2012. 2
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 5
- [25] K. Lenc and A. Vedaldi. Learning covariant feature detectors. In *ECCV Workshops*, 2016. 1, 2
- [26] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *ICCV*, 2005. 2, 3
- [27] Y.-L. Lin, V. I. Morariu, W. Hsu, and L. S. Davis. Jointly optimizing 3d model fitting and fine-grained classification. In *ECCV*, 2014. 7
- [28] C. Liu, J. Yuen, and A. Torralba. Sift flow: Dense correspondence across scenes and its applications. *T-PAMI*, 33(5):978–994, 2011. 1, 2, 6
- [29] J. L. Long, N. Zhang, and T. Darrell. Do convnets learn correspondence? In *NIPS*, 2014. 2, 5
- [30] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004. 1, 2, 5
- [31] Y. Lu, K. Huang, and C.-L. Liu. A fast projected fixed-point algorithm for large graph matching. *Pattern Recognition*, 60:971–982, 2016. 4
- [32] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *IJCV*, 65(1-2):43–72, 2005. 2
- [33] A. Nguyen, M. Ben-Chen, K. Welnicka, Y. Ye, and L. Guibas. An optimization approach to improving collections of shape maps. *Computer Graphics Forum*, 30(5):1481–1491, 2011. 2
- [34] D. Novotny, D. Larlus, and A. Vedaldi. AnchorNet: A weakly supervised network to learn geometry-sensitive features for semantic matching. In *CVPR*, 2017. 2
- [35] D. Pachauri, R. Kondor, and V. Singh. Solving the multi-way matching problem by permutation synchronization. In *NIPS*, 2013. 2, 3, 5
- [36] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013. 4
- [37] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Deepmatching: Hierarchical deformable dense matching. *IJCV*, 120(3):300–323, 2016. 2, 6
- [38] J. Thewlis, H. Bilen, and A. Vedaldi. Unsupervised learning of object landmarks by factorized spatial embeddings. In *ICCV*, 2017. 1, 2
- [39] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. *IJCV*, 9(2):137–154, 1992. 2, 4, 7
- [40] R. Tron, X. Zhou, C. Esteves, and K. Daniilidis. Fast multi-image matching via density-based clustering. In *CVPR*, 2017. 2

- [41] N. Ufer and B. Ommer. Deep semantic feature matching. In *CVPR*, 2017. 1
- [42] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders. Selective search for object recognition. *IJCV*, 104(2):154–171, 2013. 7
- [43] S. Vicente, J. Carreira, L. Agapito, and J. Batista. Reconstructing pascal voc. In *CVPR*, 2014. 7
- [44] J. Yan, M. Cho, H. Zha, X. Yang, and S. Chu. Multi-graph matching via affinity optimization with graduated consistency regularization. *T-PAMI*, 2015. 2
- [45] J. Yan, Y. Li, W. Liu, H. Zha, X. Yang, and S. M. Chu. Graduated consistency-regularized optimization for multi-graph matching. In *ECCV*, 2014. 2, 3
- [46] J. Yan, Y. Tian, H. Zha, X. Yang, Y. Zhang, and S. Chu. Joint optimization for consistent multiple graph matching. In *ICCV*, 2013. 2
- [47] J. Yan, J. Wang, H. Zha, X. Yang, and S. Chu. Consistency-driven alternating optimization for multigraph matching: A unified approach. *IEEE Transactions on Image Processing*, 24(3):994–1009, 2015. 2
- [48] J. Yan, H. Xu, H. Zha, X. Yang, H. Liu, and S. Chu. A matrix decomposition perspective to multiple graph matching. In *ICCV*, 2015. 2
- [49] K. M. Yi, E. Trulls, V. Lepetit, and P. Fua. Lift: Learned invariant feature transform. In *ECCV*, 2016. 2
- [50] C. Zach, M. Klopschitz, and M. Pollefeys. Disambiguating visual relations using loop constraints. In *CVPR*, 2010. 2
- [51] W. Zhang, J. Sun, and X. Tang. Cat head detection-how to effectively exploit shape and texture features. In *ECCV*, 2008. 8
- [52] T. Zhou, P. Krahenbuhl, M. Aubry, Q. Huang, and A. A. Efros. Learning dense correspondence via 3d-guided cycle consistency. In *CVPR*, 2016. 2, 6
- [53] T. Zhou, Y. J. Lee, S. X. Yu, and A. A. Efros. Flowweb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In *CVPR*, 2015. 2
- [54] X. Zhou, M. Zhu, and K. Daniilidis. Multi-image matching via fast alternating minimization. In *ICCV*, 2015. 2, 3, 5, 7, 8