

# Convolutional Hough Matching Networks

Juhong Min

Minsu Cho

POSTECH CSE & GSAI

<http://cvlab.postech.ac.kr/research/CHM/>

## Abstract

Despite advances in feature representation, leveraging geometric relations is crucial for establishing reliable visual correspondences under large variations of images. In this work we introduce a Hough transform perspective on convolutional matching and propose an effective geometric matching algorithm, dubbed Convolutional Hough Matching (CHM). The method distributes similarities of candidate matches over a geometric transformation space and evaluate them in a convolutional manner. We cast it into a trainable neural layer with a semi-isotropic high-dimensional kernel, which learns non-rigid matching with a small number of interpretable parameters. To validate the effect, we develop the neural network with CHM layers that perform convolutional matching in the space of translation and scaling. Our method sets a new state of the art on standard benchmarks for semantic visual correspondence, proving its strong robustness to challenging intra-class variations.

## 1. Introduction

Visual correspondence lies at the heart of image understanding, being used as a core component for numerous tasks such as object recognition, image retrieval, motion estimation, object tracking, and reconstruction [16]. With recent advances in deep neural networks [21, 23, 25, 36, 57], there has been substantial progress in learning robust feature representation for establishing correspondences. Despite the effectiveness of deep convolutional features, however, spatial matching with a geometric constraint is still essential to handle image pairs with large variations, *e.g.*, viewpoint and illumination changes, blur, occlusion, lack of texture, etc. In particular, the presence of intra-class variations, *i.e.*, scenes depicting different instances of the same categories, remains a critical challenge for correspondence [18, 20, 26, 32, 38, 42, 45, 47, 51, 52, 54]. The process of geometric matching is the de facto solution of choice, which most recent methods adopt in their models.

Geometric matching commonly relies on exploiting a geometric consensus of candidate matches to verify rela-

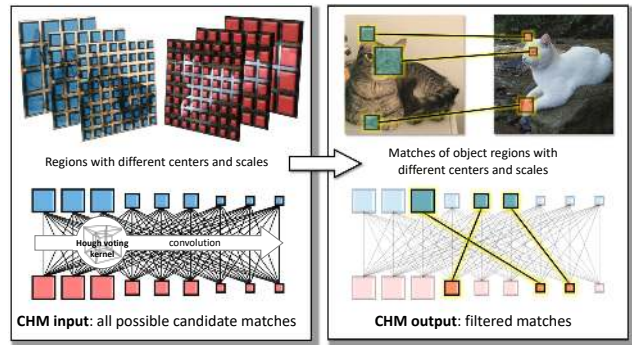


Figure 1: Convolutional Hough matching (CHM) establishes reliable correspondences across images by performing position-aware Hough voting in a high-dimensional geometric transformation space, *e.g.*, translation and scaling.

tive transformations. In computer vision, RANSAC [15] and Hough transform [22] have long been used as geometric verification for wide-baseline correspondence problems with rigid motion models, while graph matching [5, 7, 14, 55] has played a main role in matching deformable objects with non-rigid motion. Recent work [6, 20, 45, 47] has advanced the idea of Hough transform to perform non-rigid image matching, showing that the Hough voting process incorporated in neural networks is effective for challenging correspondence problems with intra-class variations. However, their matching modules are neither fully differentiable nor learnable, and weak to background clutter due to the position-invariant global Hough space.

In this work we introduce *Convolutional Hough Matching* (CHM) that distributes similarities of candidate matches over a geometric transformation space and evaluates them in a convolutional manner. As illustrated in Fig. 1, the convolutional nature makes the output equivariant to translation in the transformation space and also attentive to each position with its surrounding contexts, thus bringing robustness to background clutter. We design CHM as a learnable layer with an semi-isotropic high-dimensional kernel that acts on top of a correlation tensor. The CHM layer is compatible with any neural networks that use correlation computation, allowing flexible non-rigid matching and even multiple

matching surfaces or objects. It naturally generalizes existing 4D convolutions [26, 40, 61, 54] and provides a new perspective of Hough transform on convolutional matching. To demonstrate the effect, we propose the neural network with CHM layers that perform convolutional matching in the high-dimensional space of translation and scaling. Our method clearly outperforms state-of-the-art methods on standard benchmarks for semantic correspondence, proving its strong robustness to challenging intra-class variations.

## 2. Related Work

**Hough transformation.** The Hough transform [22] is a classic method developed to identify primitive shapes in an image via geometric voting in a parameter space. Ballard [1] generalizes the idea to identify positions of arbitrary shapes with R-table. Early approaches [4, 8] in computer vision widely adopt Hough transform for its effectiveness in extracting features of a particular shape in an image. As a representative example, Leibe *et al.* [39] introduce a Hough-based object segmentation and detection method by incorporating information about supporting patterns of parts for the target category. The idea of Hough voting has widely been adopted in diverse tasks including retrieval [24], object discovery [17, 44, 48, 50], shape recovery [59], 3D vision [34, 35], and pose estimation [29] to name a few. In geometric matching, Cho *et al.* [6] first extends it to the Probabilistic Hough Matching (PHM) algorithm for unsupervised object discovery. Recent methods [18, 19, 20, 37, 42, 45, 47, 58] have demonstrated the efficacy of the Hough matching with good empirical performance. They, however, are all limited in the sense that the geometric voting is carried out to discover a *global* offset consensus rather than a *local and individual* consensus for a match, which makes it less accurate and weak to clutter.

**Semantic visual correspondence.** Traditional approaches to the task of semantic correspondence [3, 6, 18, 19, 30, 41, 60, 63] typically use hand-crafted descriptors [2, 11, 43]. Although the classic methods work satisfactorily for some applications, they still suffer apparent disadvantages of such features, *e.g.*, lack of semantic patterns. Recent approaches [20, 27, 28, 42, 45, 47, 51, 52, 54, 56, 61, 62] build upon features from convolutional neural network (CNN) pretrained on classification task [12]. Han *et al.* [20] introduce a CNN-based matching model that learns to compute a correlation tensor. Rocco *et al.* [51] propose to learn a CNN regressor that computes a series of 2D convolutions on a dense correlation matrix to predict global geometric transformation parameters, either affine or TPS [13]. Seo *et al.* [56] improve the framework with offset-aware correlation kernels with attention modules. Jeon *et al.* [27] stack multiple affine transformation networks and compute correspondences in coarse-to-fine manner. Wang *et al.* [62] adopt the CNN architecture to estimate translation and rotation pa-

rameters to learn correspondences from raw video. These methods demonstrate that a series of 2D convolutions acting on correlation tensors is effective in capturing geometric information by exploiting local patterns of similarity.

**4D convolution for visual correspondence.** Rocco *et al.* [54] introduce the neighbourhood consensus network that uses 4D convolution for visual correspondence. They view 4D convolution as an extension of 2D convolution, which learns multiple similarity patterns of local correspondences, and thus use multiple 4D kernels, requiring a large number of parameters to learn. Following the work, recent methods [26, 40, 54, 61] also adopt 4D convolution in a similar manner. They commonly consume a high computational cost with a large number of parameters in the kernels and only consider translation in space. In contrast, we extend the idea of Hough matching [6] for high-dimensional convolution and propose an interpretable and light-weight (semi-isotropic) high-dimensional kernel for visual correspondence. In doing so, it naturally generalizes the existing 4D convolution to higher-dimensional ones and achieves superior performance using only a single kernel per layer with a small number of parameters. The results reveal that the role of high-dimensional convolution on a correlation tensor for matching is to learn a reliable voting strategy rather than to capture diverse patterns in the correlation tensor.

Our contributions can be summarized as follows:

- We introduce a Hough transform perspective on convolutional matching and propose an effective geometric matching algorithm, CHM, which performs high-dimensional Hough voting in a convolutional manner.
- We develop CHM into a trainable neural layer with a semi-isotropic high-dimensional kernel, which learns non-rigid matching with a small number of interpretable parameters.
- We propose the convolutional Hough matching network (CHMNet) that performs geometric matching in a translation and scaling space using 6D convolution.
- The proposed method sets a new state of the art on standard benchmarks for semantic visual correspondence, proving its robustness to challenging intra-class variations across images to match.

## 3. Convolutional Hough Matching

In this section, we revisit the Hough matching method for visual correspondence and then propose its convolutional version as a high-dimension convolutional layer, which is readily trainable in neural networks.

### 3.1. Hough matching & its convolutional extension

The Hough transform is a powerful detection method for a geometric object, which exploits the duality between parts

and parameters of the object [1, 22]. It performs voting in a parameter space of the target object, called the Hough space, where votes from the object parts are accumulated to form local maxima in the space. The objects are then detected simply by identifying the positions of local maxima. The Hough matching method [6], inspired by the Hough transform, detects reliable correspondences by geometric voting from candidate matches. Given two images, it constructs the Hough space of parameters of geometric transformation between the two images and then accumulates votes of candidate matches for plausible transformation.

Let us assume a local region  $\mathbf{x}$  on an image, that is represented by its geometric attributes, *i.e.*, pose and shape. In principle  $\mathbf{x}$  can be a form of any parameterization, but in this work we simply describe the region  $\mathbf{x}$  by its center and scale. Now let us consider two images,  $I$  and  $I'$ , and two sets of local regions,  $\mathcal{X}$  and  $\mathcal{X}'$ , obtained from the two images, respectively. For any two regions  $(\mathbf{x}, \mathbf{x}') \in \mathcal{X} \times \mathcal{X}'$ , a correlation function  $c$  computes a non-negative similarity  $c(\mathbf{x}, \mathbf{x}')$  using appearance features of the regions. The main idea of Hough matching is to create the Hough space  $\mathcal{H}$ , that is the space of all possible offsets  $\mathbf{h}$  between two regions, *i.e.*, translation and scaling, and accumulate votes from candidate matches onto the Hough space as

$$v(\mathbf{h}) = \sum_{(\mathbf{x}, \mathbf{x}') \in \mathcal{X} \times \mathcal{X}'} c(\mathbf{x}, \mathbf{x}') k_{\text{iso}}(\|\mathbf{x}' - \mathbf{x} - \mathbf{h}\|_{\text{g}}), \quad (1)$$

where  $\|\cdot\|_{\text{g}}$  represents a group-wise distance function that computes the distances separately for two groups, center and scale, *i.e.*,  $\|\mathbf{x}\|_{\text{g}} = [\|\mathbf{x}_{\text{xy}}\|; \|\mathbf{x}_{\text{s}}\|]$  (subscripts xy for center and s for scale) and  $k_{\text{iso}}$  is a kernel function that computes similarity between the observed offset,  $\mathbf{x}' - \mathbf{x}$ , and the given offset  $\mathbf{h}$  in the Hough space.<sup>1</sup> The kernel  $k_{\text{iso}}$  is designed to assign a voting weight for each candidate match according to how close the offset induced by the match  $(\mathbf{x}, \mathbf{x}')$  is to  $\mathbf{h}$ ; we use the group-wise distance to differentiate the effects of center and scale in the kernel. The resultant voting map  $v(\mathbf{h})$  over the Hough space  $\mathcal{H}$  can be used to find reliable matches  $(\mathbf{x}, \mathbf{x}')$  by suppressing spurious ones corresponding to relatively low voting scores  $v(\mathbf{x}' - \mathbf{x})$ , *e.g.*, updating the match score via  $c(\mathbf{x}, \mathbf{x}')v(\mathbf{x}' - \mathbf{x})$  [20]. Despite its good empirical performance [6, 18, 19, 20, 37, 42, 45, 47, 58], the global voting map  $v(\mathbf{h})$ , which is shared for all candidate matches, is limited in the sense that it cannot capture the reliability of a specific candidate match. This global position-invariant Hough space makes the output less accurate and weak to background clutter, *e.g.*, increasing the score of distant outliers that has a similar offset to that of dominant inliers.

As illustrated in Fig. 2, in order to address the issue, we create a local and individual voting space for each candidate

<sup>1</sup>For the kernel function, previous work uses a form of discretized Gaussian [6] or Dirac delta [20] without learning the kernel parameters.

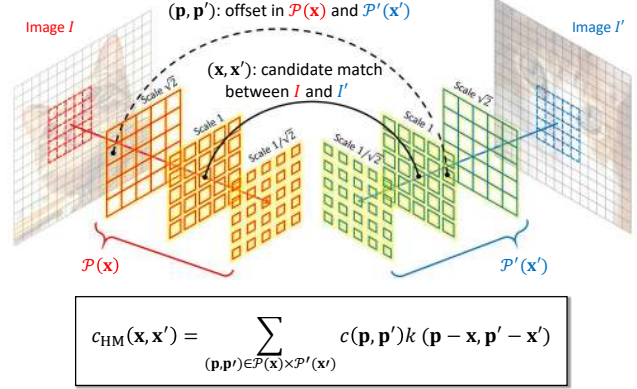


Figure 2: Convolutional Hough matching that carries out geometric voting in 6D space, *e.g.*, translation and scale.

match  $(\mathbf{x}, \mathbf{x}')$  by introducing local windows around the regions,  $\mathbf{x}$  and  $\mathbf{x}'$ :

$$v(\mathbf{x}, \mathbf{x}', \mathbf{h}) = \sum_{(\mathbf{p}, \mathbf{p}') \in \mathcal{P}(\mathbf{x}) \times \mathcal{P}'(\mathbf{x}')} c(\mathbf{p}, \mathbf{p}') k_{\text{iso}}(\|\mathbf{p}' - \mathbf{p} - \mathbf{h}\|_{\text{g}}), \quad (2)$$

where  $\mathcal{P}(\mathbf{x})$  denotes the set of neighbor regions within the local window centered on  $\mathbf{x}$ . Since this local voting space is now dedicated to  $(\mathbf{x}, \mathbf{x}')$ , we can simply assign a match score  $v$  for the candidate match by taking the vote value at the bin with offset zero:

$$v(\mathbf{x}, \mathbf{x}') = \sum_{(\mathbf{p}, \mathbf{p}') \in \mathcal{P}(\mathbf{x}) \times \mathcal{P}'(\mathbf{x}')} c(\mathbf{p}, \mathbf{p}') k_{\text{iso}}(\|\mathbf{p}' - \mathbf{p}\|_{\text{g}}). \quad (3)$$

With a slight abuse of notation, let us use  $k(\mathbf{z}, \mathbf{z}')$  to represent the kernel value corresponding to two positions,  $\mathbf{z}$  and  $\mathbf{z}'$ , each representing a local region in the parameter space of regions, *i.e.*, 3D space of center and scale in our case. The equation above then can be generalized to a form of 6D convolution with an arbitrary kernel  $k$ :

$$\begin{aligned} c_{\text{HM}}(\mathbf{x}, \mathbf{x}') &= \sum_{(\mathbf{p}, \mathbf{p}') \in \mathcal{P}(\mathbf{x}) \times \mathcal{P}'(\mathbf{x}')} c(\mathbf{p}, \mathbf{p}') k(\mathbf{p} - \mathbf{x}, \mathbf{p}' - \mathbf{x}') \\ &= (c * k)(\mathbf{x}, \mathbf{x}'), \end{aligned} \quad (4)$$

which becomes equivalent to Eq. 3 when the group-wise isotropic kernel  $k_{\text{iso}}$  is used.

Note that this convolutional extension of Hough matching has a generic form; it reduces to a similar form of 4D convolutions in [26, 40, 54, 61] when the Hough space is restricted to center translation, and generalizes to higher dimensions beyond 6D when additional transformation dimensions is introduced such as rotation, shear, and others.

### 3.2. Convolutional Hough matching layer

We design the convolutional Hough matching (CHM) as a learnable convolution layer:

$$c_{\text{HM}}(\mathbf{x}, \mathbf{x}'; k, b) = b + (c * k)(\mathbf{x}, \mathbf{x}'), \quad (5)$$

where  $b$  is a bias term for the layer and  $k$  represents a kernel with a specific type of weight sharing. The group-wise isotropic kernel  $k_{\text{iso}}$ , which is directly derived from Hough matching, can be implemented by weight sharing among parameters with the same offset  $|\mathbf{z} - \mathbf{z}'|$  in  $k(\mathbf{z}, \mathbf{z}')$ . While it is a reasonable choice, the fully isotropic kernel assigns the same importance to the matches of the same offset regardless of their distances from the kernel position  $(\mathbf{x}, \mathbf{x}')$ . It may be an excessive constraint in the sense that the distance of an object from the center of focus is likely to be relevant to the importance.

We thus relax the isotropy and propose the position-sensitive isotropic kernel  $k_{\text{psi}}(\|\mathbf{p}' - \mathbf{p}\|_{\text{g}}, \|\mathbf{p} - \mathbf{x}\|_{\text{g}}, \|\mathbf{p}' - \mathbf{x}'\|_{\text{g}})$  that differentiates the distances from the kernel position,  $\|\mathbf{p} - \mathbf{x}\|_{\text{g}}$  and  $\|\mathbf{p}' - \mathbf{x}'\|_{\text{g}}$ . The kernel  $k_{\text{psi}}$  is implemented by sharing parameters whose triplets,  $(\|\mathbf{p}' - \mathbf{p}\|_{\text{g}}, \|\mathbf{p} - \mathbf{x}\|_{\text{g}}, \|\mathbf{p}' - \mathbf{x}'\|_{\text{g}})$ , are the same.

The CHM layer is compatible with any neural network layer that computes correlations between images, and can be stacked multiple times to improve the performance. As a result of substantial parameter sharing, the 6D kernels,  $k_{\text{iso}}^{6\text{D}}$  and  $k_{\text{psi}}^{6\text{D}}$  in  $\mathbb{R}^{H_k \times W_k \times S_k \times H_k \times W_k \times S_k}$ , contain only a small number of parameters, thus making CHM resistant to overfitting in training; *e.g.*, the kernels with  $H_k = W_k = 5$  and  $S_k = 3$  contains only 45 and 220 parameters, respectively, while the full kernel has 5,625. More importantly, the perspective of Hough matching on convolution provides the interpretability of the learned kernel: each element in the kernel is a voting weight of the corresponding neighbor match in the local offset space. Based on this perspective, Figure 3 visualizes the kernel  $k_{\text{psi}}^{6\text{D}} \in \mathbb{R}^{H_k \times W_k \times S_k \times H_k \times W_k \times S_k}$  of size  $H_k = W_k = 5$  and  $S_k = 3$  trained in our experiment. For the ease of visualizing 6D tensor, we decompose it into multiple (four in case of  $k_{\text{psi}}^{6\text{D}}$ ) 4D tensors in which each of the map shows parameter values of the kernel with the same offset, where the arrows represent the offset vectors relative to the kernel position  $(\mathbf{x}, \mathbf{x}')$ , and the circles mean zero offset. The maps reveal that weights for matches with smaller offsets and closer distance are learned to be higher (darker), which appears to be a reasonable voting strategy. For more information, refer to our supplementary.

### 4. Convolutional Hough Matching Networks

Based on CHM, we develop a family of image matching models, dubbed *Convolutional Hough Matching Networks (CHMNet)*, which consists of three parts: (1) high-dimensional correlation computation, (2) convolutional

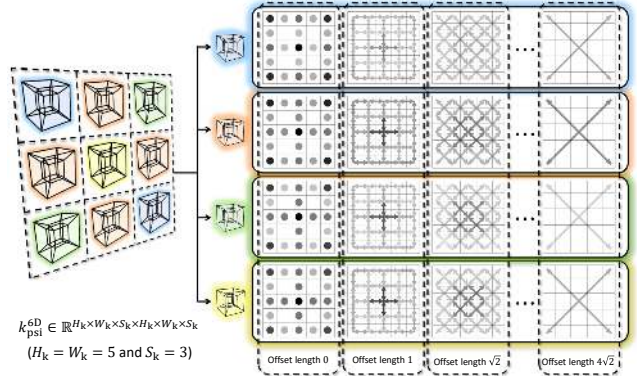


Figure 3: Visualization of learned CHM kernel ( $k_{\text{psi}}^{6\text{D}}$ ). See our supplementary material for the visualization method.

Hough matching, and (3) flow formation (and keypoint transfer). Figure 4 illustrates the overall architecture.

#### 4.1. High-dimensional correlation computation

Following other recent methods [26, 40, 45, 47, 54], we also use as a CNN feature extractor pretrained on ImageNet classification [12]. Given an input image  $I$ , the feature extractor outputs a feature map in  $\mathbb{R}^{C \times H \times W}$ . We construct feature maps of multiple scales  $\{\mathbf{F}_s\}_{s=1}^S$  by resizing the output for  $S - 1$  times by the scaling factor of  $\sqrt{2}$ , followed by  $3 \times 3$  conv layers with parameters  $\{\theta_s\}_{s=1}^S$ , reducing channel dimensions of input feature map by  $1/\rho$ . The  $S$  different conv layers learn to capture effective semantic information of receptive fields with different scales for the subsequent multi-scale (6D) correlation computation. The same is done for  $\{\mathbf{F}'_s\}_{s=1}^S$  given image  $I'$ . We set  $S = 3$ , *i.e.*,  $\{1/\sqrt{2}, 1, \sqrt{2}\}$ , and  $\rho = 4$  in our experiments.

Given a set of feature pairs from multiple scales  $\{(\mathbf{F}_s, \mathbf{F}'_s)\}_{s=1}^S$ , we compute all possible 4D correlation tensors placed on the  $S \times S$  grid:

$$\mathbf{C}_{mn}^{(0)}(\mathbf{x}_m, \mathbf{x}'_n) = \text{ReLU} \left( \frac{\mathbf{F}_m(\mathbf{x}_m) \cdot \mathbf{F}'_n(\mathbf{x}'_n)}{\|\mathbf{F}_m(\mathbf{x}_m)\| \|\mathbf{F}'_n(\mathbf{x}'_n)\|} \right), \quad (6)$$

where  $\mathbf{x}_m \in \mathcal{X}_m$  and  $\mathbf{x}'_n \in \mathcal{X}'_n$  are spatial positions of feature map at scale  $m$  and  $n$ , respectively, and ReLU clamps negative correlation scores to zero. To process it in the subsequent 6D CHM layer, we interpolate each 4D correlation  $\mathbf{C}_{ij}^{(0)}$  to have the same spatial size to build 6D correlation tensor  $\mathbf{C}^{(1)} \in \mathbb{R}^{H \times W \times S \times H \times W \times S}$  such that  $\mathbf{C}_{::i::j}^{(1)} = \zeta_1(\mathbf{C}_{ij}^{(0)})$  where  $\zeta_1(\cdot)$  is a function that interpolates input 4D tensor to the size  $H \times W \times H \times W$ .

#### 4.2. Convolutional Hough Matching

A CHM layer takes the 6D correlation tensor  $\mathbf{C}^{(1)}$  to perform convolutional Hough voting in the space of translation and scaling:  $\mathbf{C}^{(2)} = \text{CHM}(\mathbf{C}^{(1)}; k_{\text{psi}}^{6\text{D}})$ , where  $k_{\text{psi}}^{6\text{D}} \in \mathbb{R}^{H_k \times W_k \times S_k \times H_k \times W_k \times S_k}$  is a 6D position-sensitive

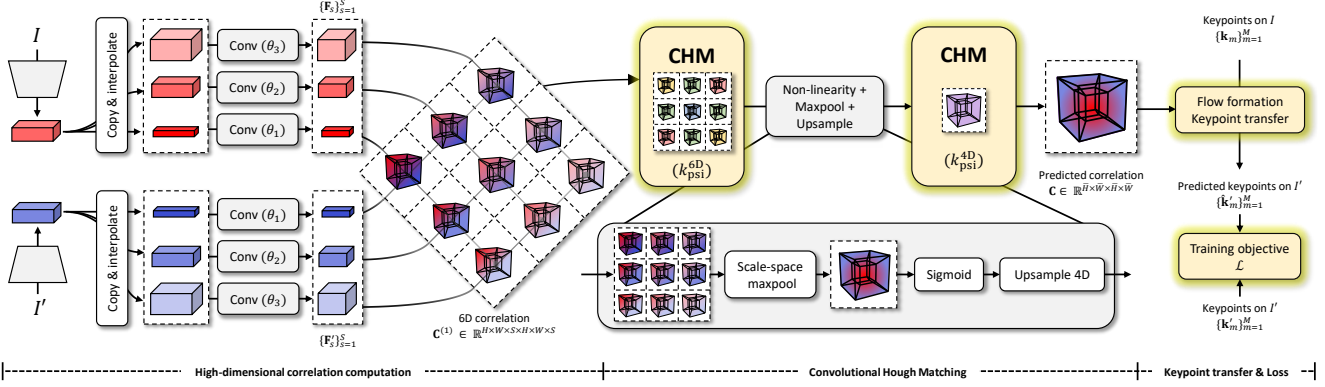


Figure 4: Overall architecture of the proposed method that performs (learnable) geometric voting in high-dimensional spaces.

isotropic kernel. In our experiments, we set  $H_k = W_k = 5$  and  $S_k = 3$  with stride 1 for all dimensions and use zero-padding to the input to retain the same size at the output. We then perform max-pooling on  $C^{(2)}$  to select the most dominant vote among candidate match scores in the scale space, reducing the tensor dimension down to 4D:  $C_{ijkl}^{(3)} = \max_{m,n} C_{ijmkl}^{(2)}$ . We proceed another CHM with a 4D kernel  $k_{\text{psi}}^{4D} \in \mathbb{R}^{H_k \times W_k \times H_k \times W_k}$ :  $C = \text{CHM}(\zeta_2(\sigma(C^{(3)}))); k_{\text{psi}}^{4D}$ , where  $\sigma(\cdot)$  is the sigmoid activation function and  $\zeta_2(\cdot)$  is the upsampling function that resizes input 4D tensor to the size of  $\bar{H} \times \bar{W} \times \bar{H} \times \bar{W}$  for fine-grained localization. We set  $\bar{H} = 2H$  and  $\bar{W} = 2W$  in our experiment.

### 4.3. Flow formation & keypoint transfer

**Flow formation.** The output  $C$  can easily be transformed into a dense flow field by applying kernel soft-argmax [38]. We first normalize the raw correlation scores with softmax:

$$\hat{C} = \frac{\exp(\mathbf{G}_{kl}^{\mathbf{p}} C_{ijkl})}{\sum_{(k',l') \in \bar{H} \times \bar{W}} \exp(\mathbf{G}_{k'l'}^{\mathbf{p}} C_{ij k'l'})}, \quad (7)$$

where and  $\mathbf{G}^{\mathbf{p}} \in \mathbb{R}^{\bar{H} \times \bar{W}}$  is 2-dimensional Gaussian kernel centered on  $\mathbf{p} = \arg \max_{k,l} C_{ijkl}$ . Using the estimated probability map  $\hat{C}$ , we then transfer all the coordinates on dense regular grid  $\mathbf{P} \in \mathbb{R}^{\bar{H} \times \bar{W} \times 2}$  of image  $I$  to obtain their corresponding coordinates  $\hat{\mathbf{P}}' \in \mathbb{R}^{\bar{H} \times \bar{W} \times 2}$  on image  $I'$ :  $\hat{\mathbf{P}}'_{ij} = \sum_{(k,l) \in \bar{H} \times \bar{W}} \hat{C}_{ijkl} \mathbf{P}_{kl}$ . We now can construct a dense flow field at sub-pixel level using the set of estimated matches  $(\mathbf{P}, \hat{\mathbf{P}}')$ .

**Keypoint transfer.** As in [38], one simplest way of assigning a match  $\hat{\mathbf{k}}$  to some keypoint  $\mathbf{k} = (x_k, y_k)$  is to pick a single, discrete sample of a transferred coordinate such that  $\hat{\mathbf{k}} = \hat{\mathbf{P}}'_{y_k x_k}$ . However, this may cause mis-localized keypoints as the discrete sampling under sub-pixel level hinders fine-grained localization. To this end, we define a soft sampler  $\mathbf{W}^{(\mathbf{k})} \in \mathbb{R}^{\bar{H} \times \bar{W}}$  for given keypoint  $\mathbf{k} = (x_k, y_k)$

as follows

$$\mathbf{W}_{ij}^{(\mathbf{k})} = \frac{\max(0, \tau - \sqrt{(x_k - j)^2 + (y_k - i)^2})}{\sum_{i',j'} \max(0, \tau - \sqrt{(x_k - j')^2 + (y_k - i')^2})}, \quad (8)$$

such that  $\sum_{ij} \mathbf{W}_{ij}^{(\mathbf{k})} = 1$  where  $\tau$  is a distance threshold. We assign a match to the keypoint  $\mathbf{k}$  by  $\hat{\mathbf{k}} = \sum_{(i,j) \in \bar{H} \times \bar{W}} \hat{\mathbf{P}}'_{ij} \mathbf{W}_{ij}^{(\mathbf{k})}$ . The soft sampler  $\mathbf{W}^{(\mathbf{k})}$  effectively samples each transferred keypoint  $\hat{\mathbf{P}}'_{ij}$  by giving weights inversely proportional to the distance to  $\mathbf{k}$ .

### 4.4. Training objective

We assume that keypoint match annotations are given for each training image pair, as in [9, 20, 40, 45, 47]; each image pair is annotated with a set of coordinate pairs  $\mathcal{M} = \{(\mathbf{k}_m, \mathbf{k}'_m)\}_{m=1}^M$ , where  $M$  is the number of annotations. Following the aforementioned keypoint transfer scheme, we obtain a set of predicted and ground-truth keypoint pairs on image  $I'$ :  $\{(\hat{\mathbf{k}}'_m, \mathbf{k}'_m)\}_{m=1}^M$  by assigning a match  $\hat{\mathbf{k}}'_m$  to each  $\mathbf{k}_m$ . Our objective in training is formulated as  $\mathcal{L} = \frac{1}{M} \sum_{m=1}^M \|\hat{\mathbf{k}}'_m - \mathbf{k}'_m\|$ , which minimizes the average Euclidean distance between the predicted keypoints and the ground-truth ones.

## 5. Experimental Evaluation

In this section we evaluate the proposed method, compare it with recent state of the arts, and discuss the results.

**Implementation detail.** For the feature extractor network, we employ ResNet-101 [21], truncated after the conv4\_23 layer, pre-trained on ImageNet [12]. Both input and output channel sizes of all the CHM layers are set to 1. We set spatial size of the input image to  $240 \times 240$ , thus having  $H = W = 15$  and  $\bar{H} = \bar{W} = 30$ . Due to parameter sharing structure of  $k_{\text{psi}}^*$  and  $k_{\text{iso}}^*$ , magnitudes of the loss gradient with respect to the shared weights are unevenly distributed during training time. To resolve the numerical instability, the shared weights are normalized before the con-

Sup.	Methods	SPair-71k		PF-PASCAL		PF-WILLOW		uses nD conv?	FLOPs (G)	time (ms)	memory (GB)
		PCK @ 0.1 (F)	$\alpha_{\text{bbox}}$ 0.1 (T)	PCK @ 0.05	$\alpha_{\text{img}}$ 0.1	PCK @ 0.05	$\alpha_{\text{bbox}}$ 0.1				
I	NC-Net <sub>res101</sub> [54]	20.1	26.4	54.3	78.9	33.8	67.0	4D	44.9	222	<u>1.2</u>
	DCC-Net <sub>res101</sub> [26]	-	26.7	55.6	82.3	43.6	73.8	4D	47.1	567	2.7
	DHPF <sub>res101</sub> [47]	27.7	28.5	56.1	82.1	<u>50.2</u>	<b>80.2</b>	$\times$	<b>2.0</b>	<u>58</u>	1.6
K	UCN <sub>res101</sub> [9]	-	17.7	-	75.1	-	-	$\times$	-	-	-
	HPF <sub>res101</sub> [45]	28.2	-	60.1	84.8	45.9	74.4	$\times$	-	63	-
	SCOT <sub>res101</sub> [42]	35.6	-	63.1	85.4	47.8	76.0	$\times$	<u>6.2</u>	151	4.6
	DHPF <sub>res101</sub> [47]	<u>37.3</u>	27.4	<u>75.7</u>	<u>90.7</u>	49.5	77.6	$\times$	<b>2.0</b>	<u>58</u>	1.6
	NC-Net <sub>res101</sub> * [54]	-	-	-	81.9	-	-	4D	44.9	222	<u>1.2</u>
	DCC-Net <sub>res101</sub> * [26]	-	-	-	83.7	-	-	4D	47.1	567	2.7
	ANC-Net <sub>res101</sub> * [40]	-	<u>28.7</u>	-	86.1	-	-	4D	44.9	216	<b>0.9</b>
	CHMNet <sub>res101</sub> (ours)	<b>46.3</b>	<b>30.1</b>	<b>80.1</b>	<b>91.6</b>	<b>52.7</b>	<u>79.4</u>	6D	19.6	<b>54</b> <sup>†</sup> (248)	1.6

Table 1: Performance on standard benchmarks in accuracy, FLOPs, per-pair inference time, and memory footprint. Subscripts denote backbone networks. Some results are from [27, 31, 40, 42, 45, 47]. Numbers in bold indicate the best performance and underlined ones are the second best. Models with an asterisk (\*) are retrained using keypoint annotations (strong supervision) from [40]. The first column shows supervisory signals used in training: image-level labels (I), and keypoint matches (K). Superscript <sup>†</sup> denotes inference time using our implementation of nD conv.

volution by dividing by the number of times being shared. The network is implemented in PyTorch [49] and optimized using Adam [33] with a learning rate of 1e-3. We finetune the backbone network by setting its learning rate 100 times smaller than CHM layers, *e.g.*, 1e-5.

**Datasets.** We evaluate the proposed network on three standard benchmark datasets of semantic correspondence: SPair-71k [46], PF-PASCAL [19], and PF-WILLOW [18]. SPair-71k [46] is a highly challenging, large-scale dataset, which contains 70,958 pairs from 18 categories with large variations in view-point and scale. PF-PASCAL [19] and PF-WILLOW [18] respectively contain 1,351 pairs from 20 categories and 900 pairs from 4 categories with small variations in view-point and scale. Each pair in the datasets consists of keypoint match annotations for semantic parts.

**Evaluation metric.** We adopt the standard evaluation metric, percentage of correct keypoints (PCK), for the evaluation. Given a set of predicted and ground-truth keypoint pairs  $\mathcal{K} = \{(\mathbf{k}'_m, \mathbf{k}''_m)\}_{m=1}^M$ , PCK is measured by  $\text{PCK}(\mathcal{K}) = \frac{1}{M} \sum_{m=1}^M \mathbb{1}[\|\mathbf{k}'_m - \mathbf{k}''_m\| \leq \alpha_\tau \cdot \max(w_\tau, h_\tau)]$  where  $w_\tau$  and  $h_\tau$  are the width and height of either an entire image or an object bounding box, *e.g.*,  $\tau \in \{\text{img}, \text{bbox}\}$ , and  $\alpha_\tau$  is a tolerance factor.

## 5.1. Results and analysis

On the SPair-71k dataset, following [45, 47], we evaluate two versions for each model: a finetuned model (F), which is trained on SPair-71k, and a transferred model (T), which is trained on PF-PASCAL. On PF-PASCAL and PF-WILLOW, following the common evaluation protocol [9, 20, 26, 31, 40, 45, 47, 52, 54], our network is trained on the training split of PF-PASCAL [19] and evaluated on the test splits of PF-PASCAL and PF-WILLOW. We use the same training, validation, and test splits of PF-

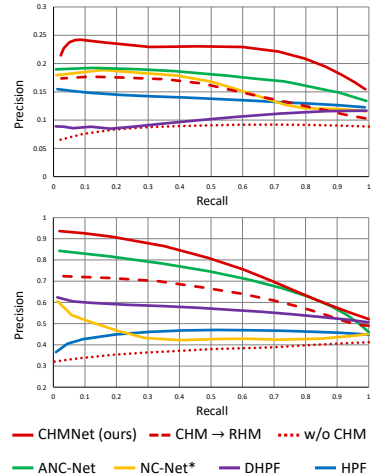


Figure 5: PR curves on SPair-71k (top) and PF-PASCAL (bottom).

PASCAL used in [20]. The quantitative results are summarized in Tab. 1; we note different levels of supervision for each method in the first column to ensure fair comparison. The proposed model finetuned on SPair-71k (F) clearly surpass current state of the art by a significant margin, outperforming [47] by 9%p of PCK ( $\alpha_{\text{bbox}} = 0.1$ ), *i.e.*, 24.1% relative improvement. On PF-PASCAL, our model achieves 4.4%p and 0.9%p improvement with  $\alpha_{\text{img}} \in \{0.05, 0.1\}$ . Robust performance on SPair-71k (T) and PF-WILLOW verifies reliable transferability of our model. Figure 6 visualizes example qualitative results on SPair-71k.

**FLOPs, running time, and memory.** We collect publicly available codes of some recent methods [26, 40, 42, 47, 54] to measure their FLOPs, inference time<sup>2</sup>, and memory footprint and compare them with ours in Tab. 1. Although the proposed method demands larger memory than some 4D conv based models [40, 54], smaller channel sizes of CHM (6D) layers ( $\{1,1\}$  vs.  $\{16,16,1\}$ ) provide noticeable efficiency in terms of GFLOPs (19.6 vs. 44.9). To achieve faster inference time, we further improve the original implementation of 4D conv [54] and develop an efficient nD conv which enables real-time inference (54ms) without increasing FLOPs and memory. See the supplementary material for details on our implementation of nD convolution.

**Robustness to background clutter.** Recent methods for semantic correspondence [18, 20, 26, 31, 32, 38, 45, 47, 51, 52, 54] predict matching scores for all candidate matches but rarely evaluate their robustness to background clutter. Here, we compare some recent methods [40, 45, 47, 54] and ours in terms of robustness to background clutter based

<sup>2</sup>Some inference time results are retrieved from [47], which is measured on a machine with an Intel i7-7820X and an NVIDIA Titan-XP. For fair comparison, inference time and memory footprint of all the methods are measured on a machine with the same CPU and GPU and includes all the pipelines of a model: from feature extraction to keypoint prediction.



Figure 6: Qualitative results on SPair-71k dataset. Our model predicts reliable matches under deformations, and large changes in view-point and scale.

on the predicted matching scores. Each method, however, exploits its correlation tensor differently from others with its own flow formation (keypoint transfer) scheme. Therefore, given all possible candidate matches in correlation tensor, simply defining matches with top- $k$  scores as positive matches may yield biased estimates. To ensure fair comparison, for each model, we define a set of coordinates on a regular grid on the input pair of images and assign their best matches using its own keypoint transfer method, thus providing the same number of (fairly collected) candidate matches to every model that we compare. For each candidate match, we define its match score as a score nearest to spatial position in the correlation tensor. Given top- $k$  matches according to their matching scores, we define true positives (TPs) as matches falling inside object segmentation masks (bounding box)<sup>3</sup> and false positives (FPs) as those lying outside object masks (boxes). Precision and recall are measured by  $\frac{N_{TP}}{N_{TP}+N_{FP}}$  and  $\frac{N_{TP}}{N_{mask}}$ , respectively, where  $N_{TP}$  and  $N_{FP}$  are respectively the number of TPs and FPs while  $N_{mask}$  is the number of all candidate matches that fall in the object segmentation masks. In defining TPs and FPs, we use masks and boxes only due to the absence of dense flow annotation in SPair-71k and PF-PASCAL, but we find that they are good approximation enough to distinguish inliers from outliers in our experimental setup.

Figure 5 plots precision-recall curves for the recent methods [40, 45, 47, 54] and ours. The proposed method clearly outperforms other methods, indicating our model effectively discriminates between semantic parts and background clutters as seen in the last row of Fig. 8 which visualizes sample pairs with top 300 confident matches. When CHM is either removed (w/o CHM) or replaced with global matching module (CHM  $\rightarrow$  RHM), predicted matches become unreliable, being mostly scattered on the background and even hardly regularized. For our model evaluated on SPair-71k, precision and recall have inverse relationship in most cases. Although initial growth in our PR curves on

<sup>3</sup>We use object seg. masks and bounding boxes for SPair-71k and PF-PASCAL respectively due to absence of mask annotation in PF-PASCAL.

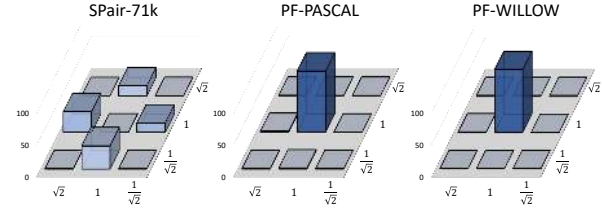


Figure 7: Frequencies over the maxpooled positions in scale-space on SPair-71k, PF-PASCAL, and PF-WILLOW.

Kernel type	SPair-71k PCK ( $\alpha_{bbox}$ )	0.05	0.1	PF-PAS. PCK ( $\alpha_{img}$ )	0.05	0.1	# params. in CHM	FLOPs (G)	time (ms)
$k_{psi}^{6D-4D}$	<b>27.4</b> $\pm$ 0.16	<b>46.4</b> $\pm$ 0.34	<b>80.4</b> $\pm$ 0.28	<b>91.6</b> $\pm$ 0.23	275	<u>19.6</u>	275	19.6	54
$k_{full}^{6D-4D}$	25.9 $\pm$ 0.74	44.8 $\pm$ 0.65	79.8 $\pm$ 0.67	90.7 $\pm$ 0.19	6,250	<u>19.6</u>	6,250	19.6	43
$k_{iso}^{6D-4D}$	24.5 $\pm$ 0.28	44.9 $\pm$ 0.16	76.5 $\pm$ 0.29	90.2 $\pm$ 0.40	<u>60</u>	<u>19.6</u>	60	19.6	46
$k_{psi}^{4D-4D}$	<u>26.4</u> $\pm$ 0.25	44.5 $\pm$ 0.34	79.3 $\pm$ 0.25	<u>91.1</u> $\pm$ 0.32	110	<b>15.9</b>	110	<b>15.9</b>	32
$k_{full}^{4D-4D}$	26.1 $\pm$ 0.33	43.9 $\pm$ 0.53	78.4 $\pm$ 0.82	90.3 $\pm$ 0.43	1,250	<b>15.9</b>	1,250	<b>15.9</b>	26
$k_{iso}^{4D-4D}$	21.0 $\pm$ 0.54	39.7 $\pm$ 0.73	71.8 $\pm$ 0.99	88.0 $\pm$ 0.49	<b>30</b>	<b>15.9</b>	<b>30</b>	<b>15.9</b>	27
$k_{psisparse}^{6D-4D}$	26.3 $\pm$ 0.18	45.2 $\pm$ 0.41	80.3 $\pm$ 0.86	91.1 $\pm$ 0.05	275	-	275	-	55

Table 2: Ablation study of CHM kernels over multiple runs.

SPair-71k indicates that some true matches have in fact low match scores, it still surpasses the other models, revealing the reliability of our approach under large variations.

## 5.2. Ablation study and analysis

**Analyses on CHM kernel.** We conduct ablation study on CHM kernel by replacing position-sensitive isotropic kernels with  $k_{full}^{nD}$ <sup>4</sup> and full isotropic ones  $k_{iso}^{nD}$ . For the ease of notation, we denote by  $k_{psi}^{6D-4D}$  a model with two CHM layers whose kernels are  $k_{psi}^{6D}$  and  $k_{psi}^{4D}$ . Table 2 shows average PCK, its standard deviations, parameter sizes, FLOPs, and average inference time of our model with different kernels over five runs. Despite a huge difference in the number of parameters (110 vs. 1,250), the proposed semi-isotropic kernel  $k_{psi}^{4D-4D}$  outperforms  $k_{full}^{4D-4D}$  on Spair-71k (44.5 vs. 43.9) and extending its voting space to 6D, *e.g.*,  $k_{psi}^{6D-4D}$ , further improves PCK to 46.4 on SPair-71k, which clearly shows efficacy of 6D convolution in scale-space<sup>5</sup>. The comparable performance of  $k_{iso}^{6D-4D}$  to  $k_{full}^{6D-4D}$  reveals that full isotropic parameter sharing can also be a reasonable choice for reducing the large capacity of  $k_{full}^{6D-4D}$ .

In Figure 7, we also plot frequencies over the maxpooled positions in scale-space after 6D CHM layer ( $k_{psi}^{6D-4D}$ ). The maximum votes on both PF-PASCAL and PF-WILLOW are mostly concentrated on the center scale whereas they are distributed over different scales on SPair-71k; this is a

<sup>4</sup>Note  $k_{full}^{nD}$  is a n-dimensional kernel without any parameter sharing. The number of parameters in  $k_{full}^{nD}$  is proportional to  $k^n$ .

<sup>5</sup>To verify the efficacy of the proposed kernel even with sparse match information, we further limit the set of potential matches in  $\mathbf{C}^{(0)}$  using  $K$  nearest neighbors without using MinkowskiEngine [10] as it does not provide high-dim. kernel customization. As seen in shaded row in Tab. 2, our model with the sparse correlation is comparably effective to  $k_{psi}^{6D-4D}$ , which is consistent to the results of [53]. We set  $K = 10$  in our experiment.



Figure 8: Ablation study on matching modules.

Method	SPair-71k		PF-PASCAL	
	PCK ( $\alpha_{\text{bbox}}$ )	PCK ( $\alpha_{\text{img}}$ )	PCK ( $\alpha_{\text{bbox}}$ )	PCK ( $\alpha_{\text{img}}$ )
	0.05	0.1	0.05	0.1
CHMNet <sub>res101</sub>	27.2	46.3	80.1	91.6
CHM $\rightarrow$ RHM	21.8	38.2	77.1	89.6
w/o last CHM layer ( $k_{\text{psi}}^{4D}$ )	24.9	43.1	79.5	89.7
w/o CHM	10.1	21.6	61.6	78.5
w/o kernel $\mathbf{G}$	26.6	45.5	79.5	91.3
w/o soft sampler $\mathbf{A}^{(k)}$	23.1	43.8	78.9	89.6

Table 3: Ablation study of core modules in our model.

reasonable voting strategy as objects in PF-PASCAL and PF-WILLOW hardly vary in scale while those in SPair-71k show large variations in both scale and view-point.

**Ablation study on matching modules.** We analyze the effect of CHM, by either removing or replacing them with the matching module of [45]. Figure 8 and Table 3 summarize qualitative and quantitative results, respectively. The output of global offset voting (CHM  $\rightarrow$  RHM) includes many outliers from the background, showing its weakness to the background clutter. Without the last CHM layer (w/o last CHM), the model fails to effectively refine upsampled correlation scores. The model prediction is severely damaged without any matching modules (w/o CHM) as seen in second row of Fig. 8. For keypoint transfer, kernel  $\mathbf{G}$  and soft sampler  $\mathbf{A}^{(k)}$  help our model find reliable matches by suppressing noisy match scores in  $\mathbf{C}$  and effectively aggregating neighborhood transfers, respectively.

**Effect of channel size.** To study the effect of channel size, we train our model<sup>6</sup> using three different kinds of kernels ( $k_{\text{psi}}^{4D-4D}$ ,  $k_{\text{iso}}^{4D-4D}$ , and  $k_{\text{full}}^{4D-4D}$ ) with different channel sizes, *i.e.*, different number of kernels. Table 9 summarizes the results, showing that increasing the channel size rarely brings performance gain and typically harms the quality of prediction for kernels  $k_{\text{psi}}^{4D-4D}$  and  $k_{\text{full}}^{4D-4D}$ . We train the models on the training split of PF-PASCAL and evaluate on test splits of PF-PASCAL and SPair-71k. For  $k_{\text{iso}}^{4D-4D}$ , although increasing channel size improves performance up to certain amount due to its small capacity, it eventually exhibits sim-

<sup>6</sup>We use the models in the middle section of Tab. 2, *e.g.*,  $k_{*}^{4D-4D}$ .

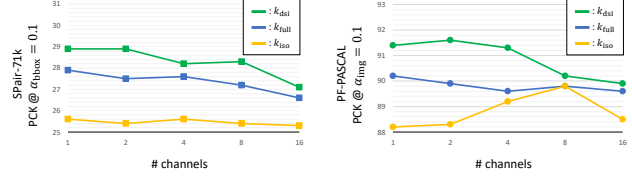


Figure 9: PCK performance on SPair-71k and PF-PASCAL with different channel sizes of 1, 2, 4, 8, and 16.

ilar patterns to other kernels after all.

These experiments imply that the high-dimensional convolution on a correlation tensor may play a different role from 2D convolution on an image feature tensor; the role of convolutional matching is to learn a reliable voting strategy rather than to capture diverse patterns in the correlation tensor. This is consistent with the Hough matching perspective, but previous 4D convolution methods [26, 40, 54, 61] with a different perspective commonly use multiple full kernels ( $k_{\text{full}}^{4D}$ ) for layers. To verify our result, we have conducted a similar experiment using the model of [54] and obtained the consistent result; the original model, which uses channel sizes of  $\{16, 16, 1\}$  for three layers of 4D convolution, achieves 76.2% PCK on our machine while the model with reduced channels of  $\{1, 1, 1\}$  achieves 76.4% PCK. Note that in terms of the number of parameters in a layer, our CHM layers ( $k_{\text{psi}}^{6D-4D}$ ) have 247  $\sim$  654 times smaller number of parameters than the 4D convolution layers used in previous methods [26, 40, 54, 61]. This light-weight layer design is particularly important in practice, since the use of multiple channels, *i.e.* kernels, for high-dim convolution quickly increases the cost both in computation and memory.

For additional results and analyses, we refer the readers to the supplementary material.

## 6. Conclusion

We have introduced the convolutional Hough matching (CHM) and proposed the powerful matching model, CHMNet, that leverages CHM in a high-dimensional geometric transformation space for establishing reliable visual correspondence. The extensive experiments on several standard benchmarks for semantic visual correspondence demonstrate the benefits of our approach. In particular, our method generalizes existing 4D convolutions and also provides the perspective of Hough transform for geometric matching with interpretable high-dimension kernels. We believe further research on this direction can benefit a wide range of other problems related to correspondence.

**Acknowledgements.** This work was supported by Samsung Advanced Institute of Technology (SAIT), the NRF grants (NRF-2017R1E1A1A01077999, NRF-2021R1A2C3012728), and the IITP grant (No.2019-0-01906, AI Graduate School Program - POSTECH) funded by Ministry of Science and ICT, Korea.



## References

- [1] Dana H. Ballard. Generalizing the hough transform to detect arbitrary shapes. *Pattern Recognition*, 13, 1981.
- [2] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *Proc. European Conference on Computer Vision (ECCV)*, 2006.
- [3] Hilton Bristow, Jack Valmadre, and Simon Lucey. Dense semantic correspondence where every pixel is a classifier. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [4] Hsin-Yi Chen, Yen-Yu Lin, and Bing-Yu Chen. Robust feature matching with alternate hough and inverted hough transforms. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [5] Minsu Cho, Karteek Alahari, and Jean Ponce. Learning graphs to match. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [6] Minsu Cho, Suha Kwak, Cordelia Schmid, and Jean Ponce. Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [7] Minsu Cho, Jungmin Lee, and Kyoung Mu Lee. Reweighted random walks for graph matching. In *Proc. European Conference on Computer Vision (ECCV)*, 2010.
- [8] Minsu Cho and Kyoung Mu Lee. Progressive graph matching: Making a move of graphs via probabilistic voting. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [9] Christopher Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [10] Christopher Choy, Jaesik Park, and Vladlen Koltun. Fully convolutional geometric features. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [11] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005.
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [13] Gianluca Donato and Serge Belongie. Approximate thin plate spline mappings. In *Proc. European Conference on Computer Vision (ECCV)*, 2002.
- [14] Matthias Fey, Jan E. Lenssen, Christopher Morris, Jonathan Masci, and Nils M. Kriege. Deep graph matching consensus. In *International Conference on Learning Representations (ICLR)*, 2020.
- [15] Martin Fischler and Robert Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981.
- [16] David Forsyth and Jean Ponce. *Computer Vision: A Modern Approach. (Second edition)*. Prentice Hall, Nov. 2011.
- [17] Juergen Gall and Victor Lempitsky. Class-specific hough forests for object detection. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [18] Bumsu Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [19] Bumsu Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow: Semantic correspondences from object proposals. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018.
- [20] Kai Han, Rafael S Rezende, Bumsu Ham, Kwan-Yee K Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. Snet: Learning semantic correspondence. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [22] Paul V.C. Hough. Method and means for recognizing complex patterns. *U.S. Patent, 3069654*, 1962.
- [23] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [24] Li Huan, Qin Yujian, and Wang Li. Vehicle logo retrieval based on hough transform and deep learning. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [25] Gao Huang\*, Zhuang Liu\*, Laurens van der Maaten, and Kilian Weinberger. Densely connected convolutional networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [26] Shuaiyi Huang, Qiuyue Wang, Songyang Zhang, Shipeng Yan, and Xuming He. Dynamic context correspondence network for semantic alignment. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [27] Sangryul Jeon, Seungryong Kim, Dongbo Min, and Kwanghoon Sohn. Pam: Pyramidal affine regression networks for dense semantic correspondence. In *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [28] Sangryul Jeon, Dongbo Min, Seungryong Kim, Jihwan Choe, and Kwanghoon Sohn. Guided semantic flow. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [29] Wadim Kehl, Fausto Milletari, Federico Tombari, Slobodan Ilic, and Nassir Navab. Deep learning of local rgb-d patches for 3d object detection and 6d pose estimation. In *Proc. European Conference on Computer Vision (ECCV)*, 2016.
- [30] Jaechul Kim, Ce Liu, Fei Sha, and Kristen Grauman. Deformable spatial pyramid matching for fast dense correspondences. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [31] Seungryong Kim, Stephen Lin, Sangryul Jeon, Dongbo Min, and Kwanghoon Sohn. Recurrent transformer networks for semantic correspondence. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [32] Seungryong Kim, Dongbo Min, Stephen Lin, and Kwanghoon Sohn. Dctm: Discrete-continuous transformation matching for semantic flow. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017.

- [33] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- [34] Jan Knopp, Mukta Prasad, and Luc Van Gool. Scene cut: Class-specific object detection and segmentation in 3d scenes. *International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*, 2011.
- [35] Jan Knopp, Mukta Prasad, and Luc Van Gool. Orientation invariant 3d object classification using hough transform based methods. In *Proceedings of the ACM Workshop on 3D Object Retrieval*, 2010.
- [36] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.
- [37] Suha Kwak, Minsu Cho, Ivan Laptev, Jean Ponce, and Cordelia Schmid. Unsupervised object discovery and tracking in video collections. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [38] Junghyup Lee, Dohyung Kim, Jean Ponce, and Bumsuh Ham. Sfnets: Learning object-aware semantic correspondence. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [39] Bastian Leibe and Bernt Schiele. Interleaved object categorization and segmentation. In *Proc. British Machine Vision Conference (BMVC)*, 2003.
- [40] Shuda Li, Kai Han, Theo W. Costain, Henry Howard-Jenkins, and Victor Prisacariu. Correspondence networks with adaptive neighbourhood consensus. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [41] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2011.
- [42] Yanbin Liu, Linchao Zhu, Makoto Yamada, and Yi Yang. Semantic correspondence as an optimal transport problem. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [43] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)*, 2004.
- [44] Fausto Milletari, Seyed-Ahmad Ahmadi, Christine Kroll, Annika Plate, Verena Rozanski, Juliana Maiostre, Johannes Levin, Olaf Dietrich, Birgit Ertl-Wagner, Kai Bötzel, and Nassir Navab. Hough-cnn: Deep learning for segmentation of deep brain regions in mri and ultrasound. *Computer Vision and Image Understanding*, 2017.
- [45] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Hyperpixel flow: Semantic correspondence with multi-layer neural features. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [46] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. SPair-71k: A large-scale benchmark for semantic correspondence. *arXiv preprint arXiv:1908.10543*, 2019.
- [47] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Learning to compose hypercolumns for visual correspondence. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [48] David Novotny, Samuel Albanie, Diane Larlus, and Andrea Vedaldi. Semi-convolutional operators for instance segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [49] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*. 2019.
- [50] Charles R. Qi, Or Litany, Kaiming He, and Leonidas J. Guibas. Deep hough voting for 3d object detection in point clouds. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [51] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [52] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. End-to-end weakly-supervised semantic alignment. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [53] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Efficient neighbourhood consensus networks via submanifold sparse convolutions. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [54] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [55] Michal Rolínek, Paul Swoboda, Dominik Zietlow, Anselm Paulus, Vít Musil, and Georg Martius. Deep graph matching via blackbox differentiation of combinatorial solvers. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [56] Paul Hongsuck Seo, Jongmin Lee, Deunsol Jung, Bohyung Han, and Minsu Cho. Attentive semantic alignment with offset-aware correlation kernels. In *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [57] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [58] Waqas Sultani and Mubarak Shah. What if we do not have multiple videos of the same action? — video action localization using web images. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [59] Min Sun, Gary Bradski, Bing-Xin Xu, and Silvio Savarese. Depth-encoded hough voting for joint object detection and shape recovery. In *Proc. European Conference on Computer Vision (ECCV)*, 2010.
- [60] Tatsunori Taniai, Sudipta N Sinha, and Yoichi Sato. Joint recovery of dense correspondence and cosegmentation in two

- images. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [61] Prune Truong, Martin Danelljan, and Radu Timofte. GLU-Net: Global-local universal network for dense flow and correspondences. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [62] Xiaolong Wang, Allan Jabri, and Alexei A. Efros. Learning correspondence from the cycle-consistency of time. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [63] Fan Yang, Xin Li, Hong Cheng, Jianping Li, and Leiting Chen. Object-aware dense semantic correspondence. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.