

Semantic-Aware Fine-Grained Correspondence

Yingdong Hu¹, Renhao Wang¹, Kaifeng Zhang¹, and Yang Gao^{1,2*}

¹ Tsinghua University

² Shanghai Qi Zhi Institute

{huyd21,wangrh21,zhangkf19}@mails.tsinghua.edu.cn

{gaoyangiii}@tsinghua.edu.cn

Abstract. Establishing visual correspondence across images is a challenging and essential task. Recently, an influx of self-supervised methods have been proposed to better learn representations for visual correspondence. However, we find that these methods often fail to leverage semantic information and over-rely on the matching of low-level features. In contrast, human vision is capable of distinguishing between distinct objects as a pretext to tracking. Inspired by this paradigm, we propose to learn semantic-aware fine-grained correspondence. Firstly, we demonstrate that semantic correspondence is implicitly available through a rich set of image-level self-supervised methods. We further design a pixel-level self-supervised learning objective which specifically targets fine-grained correspondence. For downstream tasks, we fuse these two kinds of complementary correspondence representations together, demonstrating that they boost performance synergistically. Our method surpasses previous state-of-the-art self-supervised methods using convolutional networks on a variety of visual correspondence tasks, including video object segmentation, human pose tracking, and human part tracking.

Keywords: self-supervised learning, representation learning, visual correspondence, tracking

1 Introduction

Correspondence is considered one of the most fundamental problems in computer vision. At their core, many tasks require learning visual correspondence across space and time, such as video object segmentation [45,14,69], object tracking [38,28,4,36], and optical flow estimation [17,30,51,57,58]. Despite its importance, prior art in visual correspondence has largely relied on supervised learning [67,26,61], which requires costly human annotations that are difficult to obtain at scale. Other works rely on weak supervision from methods like off-the-shelf optical flow estimators, or synthetic training data, which lead to generalization issues when confronted with the long-tailed distribution of real world images.

* Corresponding author.

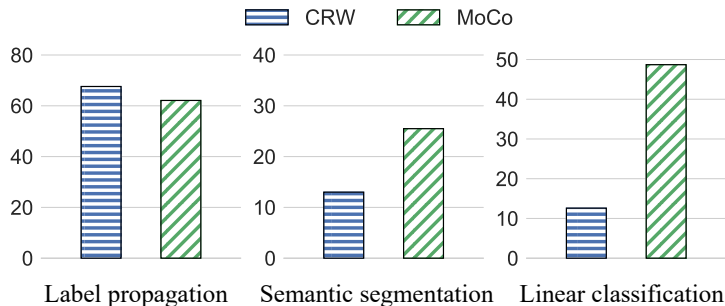


Fig. 1. We compare Contrastive Random Walk (CRW) [32] and MoCo [24] on three different downstream tasks. CRW surpasses MoCo on the label propagation task, but is dramatically outperformed by MoCo on semantic segmentation and image classification (more details are in the appendix).

Recognizing these limitations, many recent works [66,70,37,35,34,32,75] are exploring self-supervision to learn robust representations of spatiotemporal visual correspondence. Aside from creatively leveraging self-supervisory signals across space and time, these works generally share a critical tenet: evaluation on label propagation as an indication of representation quality. Given label information, such as segmentation labels or object keypoints, within an initial frame, the goal is to propagate these labels to subsequent frames based on correspondence.

Let us briefly consider the human visual system and how it performs tracking. Many works have argued that our ability to track objects is rooted in our ability to distinguish and understand differences between said objects [20,19]. We have rough internal models of different objects, which we adjust by attending more closely to local locations that require fine-grained matching [78,3]. In other words, both high-level semantic information and low-level fine-grained information play an important role in correspondence for real visual systems.

However, when we examine current self-supervised correspondence learning methods, we find them lacking under this paradigm. These methods often over-prioritize performance on the label propagation task, and fail to leverage semantic information as well as humans can. In particular, when representations obtained under these methods are transferred to other downstream tasks which require a deeper semantic understanding of images, performance noticeably suffers (see Figure 1). We show that label propagation and tracking-style tasks rely on frame-to-frame differentiation of low-level features, a kind of “shortcut” exploited by the contrastive-based self-supervised algorithms developed so far. Thus, representations learned via these tasks contain limited semantic information, and underperform drastically when used in alternative tasks.

To this end, we propose Semantic-aware Fine-grained Correspondence (SFC), which simultaneously takes into account semantic correspondence and fine-grained correspondence. Firstly, we find that current image-level self-supervised representation learning methods e.g. MoCo [24] force the mid-level convolutional features to implicitly capture correspondences between similar objects or parts. Second,

we design an objective which learns high-fidelity representations of fine-grained correspondence (FC). We do this by extending prior image-level loss functions in self-supervised representation learning to a dense paradigm, thereby encouraging local feature consistency. Crucially, FC does not use temporal information to learn this low-level correspondence, but our ablations show that this extension alone makes our model competitive with previous methods relying on temporal signals in large video datasets for pretraining.

Prior works [71,74] have shown that image-level self-supervision can further facilitate the dense self-supervision in a multitask framework. However, we surprisingly find that our fine-grained training objective and image-level semantic training objectives are inconsistent: each of them requires the model to encode conflicting information about the image, leading to degradation in performance when used in conjunction. We hypothesize that it is necessary to have two independent models, and propose a late fusion operation to combine separately pretrained semantic correspondence and fine-grained correspondence feature vectors. Figure 2 overviews the proposed method. Through our ablations, we categorically verify that low-level fine-grained correspondence and high-level semantic correspondence are complementary, and indeed orthogonal, in the benefits they bring to self-supervised representation learning. The main contributions of our work are as follows:

- We propose to learn semantic-aware fine-grained correspondence (SFC), while most previous works consider and improve the two kinds of correspondence separately.
- We design a simple and effective self-supervised learning method tailored for low-level fine-grained correspondence. Despite using static images and discarding temporal information, we outperform previous methods trained on large-scale video datasets.
- Late fusion is an effective mechanism to prevent conflicting image-level and fine-grained training objectives from interfering with each other.
- Our full model (SFC) sets the new state-of-the-art for self-supervised approaches using convolutional networks on various video label propagation tasks, including video object segmentation, human pose tracking, and human part tracking.

2 Related Work

Self-Supervised Representation Learning Self-supervised representation learning has gained popularity because of its ability to avoid the cost of annotating large-scale datasets. Specifically, methods using instance-level pretext tasks have recently become dominant components in self-supervised learning for computer vision [72,46,27,1,59,24,11,44,9,6,18,50]. Instance-level discrimination aims to pull embeddings of augmented views of the same image (positive pairs) close to each other, while trying to push away embeddings from different images




Fig. 2. Overview of Semantic-aware Fine-grained Correspondence learning framework. By maximizing agreement between positive (similar) image pairs, convolutional representations capture semantic correspondences between similar objects implicitly. By encouraging the spatially close local feature vectors to be consistent, model can learn fine-grained correspondence explicitly. For downstream task, we utilize two kinds of correspondence together to achieve complementary effects.

(negative pairs). Recently, some works [21,12] have discovered that even without negative pairs, self-supervised learning can exhibit strong performance.

To obtain better transfer performance to dense prediction tasks such as object detection and semantic segmentation, other works [48,74,71] explore pretext tasks at the pixel level for representation learning. But empirically, we find that these methods fail to leverage fine-grained information well. Our fine-grained correspondence network (FC) is most closely related to PixPro [74] which obtains positive pairs by extracting features from the same pixel through two asymmetric pipelines. Both FC and PixPro can be seen as dense versions of BYOL [21], but the two methods have completely different goals. FC has many design choices tailored for correspondence learning: FC preserves spatial sensitivity by avoiding entirely a pixel propagation module which introduces a certain smoothing effect. Furthermore, we discard color augmentation and use higher resolution feature maps, as we find both modifications are beneficial to the fine-grained correspondence task. Finally, FC can achieve competitive performance to predominant approaches, with compelling computational and data efficiency. In contrast, the transfer performance of PixPro on the correspondence task is far behind its instance-level counterpart [21] and our FC.

We note that DINO [7], a self-supervised Vision Transformer (ViT) [16], exhibits surprisingly strong correspondence properties and competitive performance on DAVIS-2017 benchmark. We speculate that the success of DINO on this task is attributed to the architecture of ViT and much more computation.

Self-Supervised Correspondence Learning Recently, numerous approaches have also been developed for correspondence learning in a self-supervised manner [66,70,68,35,34,37,32,75]. The key idea behind a number of these methods [66,35,34] is to propagate the color of one frame in a video to future frames. TimeCycle [70] relies on a cycle-consistent tracking pretext task. Along this line, CRW [32] cast correspondence as pathfinding on a space-time graph, also using

cycle-consistency as a self-supervisory signal. VFS [75] propose to learn correspondence implicitly by performing image-level similarity learning. Despite the success of these methods, they all rely heavily on temporal information from videos as the core form of self-supervision signal. In our work, we demonstrate that representations with good space-time correspondence can be learned even without videos. Moreover, our framework is an entirely alternative perspective on correspondence learning, which can be flexibly adapted with other video-based methods to further improve performance.

Semantic Correspondence We borrow the notion of semantic correspondence from literature [52,53,54,42,43,39,41,60], which aim to establish dense correspondences across images depicting different instances of the same object categories. Evaluation of these methods exists solely on image datasets with keypoint annotations, which can be more forgiving and translates poorly to the real world. Our semantic correspondence is evaluated on video, which we argue is a much more realistic setting for correspondence. In addition, many supervised semantic correspondence approaches [13,42,29,39] adopt a CNN pre-trained on image classification as their frozen backbone, but we explore a self-supervised pre-trained backbone as an alternative.

3 Method

While our framework is compatible with a wide array of contemporary self-supervised representation methods, we demonstrate its efficacy with two recent approaches: MoCo [24] and BYOL [21], which are reviewed in Section 3.1. Next, in Section 3.2, we argue that image-level methods implicitly learn high-level semantic correspondence. In Section 3.3, we propose our framework to improve fine-grained correspondence learning. Finally, in Section 3.4, we show how to unify these two complementary forms of correspondence to improve performance on video label propagation tasks.

3.1 Background

In image-level self-supervised representation learning, we seek to minimize a distance metric between two random augmentations \mathbf{x}_1 and \mathbf{x}_2 of a single image \mathbf{x} . One popular framework for doing this is contrastive learning [22].

Formally, two augmented views \mathbf{x}_1 and \mathbf{x}_2 are fed into an online encoder and target encoder respectively, where each encoder consists of a backbone \mathbf{f} (e.g. ResNet), and a projection MLP head \mathbf{g} . The l_2 -normalized output global feature vectors for \mathbf{x}_1 and \mathbf{x}_2 can be represented as $\mathbf{z}_1 \triangleq \mathbf{g}_\theta(\mathbf{f}_\theta(\mathbf{x}_1))$ and $\mathbf{z}_2 \triangleq \mathbf{g}_\xi(\mathbf{f}_\xi(\mathbf{x}_2))$, where θ and ξ are parameters of the two respective networks. Let the negative features obtained from K different images be represented by the set $\mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_K\}$. Then contrastive learning uses the InfoNCE [46] to pull \mathbf{z}_1 close to \mathbf{z}_2 while pushing it away from negative features:

$$\mathcal{L}_{\text{InfoNCE}} = -\log \frac{\exp(\mathbf{z}_1 \cdot \mathbf{z}_2 / \tau)}{\exp(\mathbf{z}_1 \cdot \mathbf{z}_2 / \tau) + \sum_{k=1}^K \exp(\mathbf{z}_1 \cdot \mathbf{s}_k / \tau)} \quad (1)$$

where τ is the temperature hyperparameter. While numerous methods have been explored to construct the set of negative samples \mathcal{S} , we choose MoCo [24] for obtaining semantic correspondence representations, which achieves this goal via a momentum-updated queue. In particular, the target encoder’s parameters ξ are the exponential moving average of the online parameters θ :

$$\xi \leftarrow m\xi + (1 - m)\theta, \quad m \in [0, 1] \quad (2)$$

where m is the exponential moving average parameter.

Some recent works [21,12] show that it is not necessary to use negative pairs to perform self-supervised representation learning. One such method, BYOL [21], relies on an additional prediction MLP head \mathbf{q}_θ to transform the output of online encoder $\mathbf{p}_1 \triangleq \mathbf{q}_\theta(\mathbf{z}_1)$. The contrastive objective then reduces to simply minimizing the negative cosine distance between the predicted features \mathbf{p}_1 and the features obtained from the target encoder \mathbf{z}_2 (l_2 -normalized):

$$\mathcal{L}_{\text{global}} = -\frac{\langle \mathbf{p}_1, \mathbf{z}_2 \rangle}{\|\mathbf{p}_1\|_2 \cdot \|\mathbf{z}_2\|_2} \quad (3)$$

Note again that MoCo and BYOL bear striking similarities in their formulation and training objectives. In the following section, we hypothesize that such similarities in frameworks lead to similarities in types of features learned. In particular, we claim that image-level representations in general contain information about semantic correspondences.

3.2 Semantic Correspondence Learning

Representations learned by current self-supervised correspondence learning methods may contain limited semantic information. To make the representations more neurophysiologically intuitive, we add the crucially missing semantic correspondence learning into our method. Recent image-level self-supervised methods learn representations by imposing invariances to various data augmentations. Two random crops sampled from the same image, followed by strong color augmentation [9] are considered as positive pairs. The augmentation significantly changes the visual appearance of the image but keeps the semantic meaning unchanged. The model can match positive pairs by attending only to the essential part of the representation, while ignoring other non-essential variations. As a result, different images with similar visual concepts are grouped together, inducing a latent space with rich semantic information [62,10,63]. This is evidenced by the results shown in Figure 1, where MoCo [24] achieve high performance on tasks that require a deeper semantic understanding of images. Moreover, previous works [40,75] demonstrate that correspondence naturally emerges in the middle-level convolutional features. Thus we conclude that current self-supervised representation methods can implicitly learn semantic correspondence well.

We utilize one approach, MoCo [24], in our downstream correspondence task. In particular, only the pre-trained online backbone \mathbf{f}_θ is retained, while all other parts of the network, including the online projection head \mathbf{g}_θ and target encoder

$\mathbf{f}_\xi, \mathbf{g}_\xi$, are discarded. We use \mathbf{f}_θ to encode each image as a semantic correspondence feature map: $\mathbf{F} = \mathbf{f}_\theta(\mathbf{x}) \in \mathbb{R}^{H \times W \times C_s}$, where H and W are spatial dimensions. Note also that we can adjust the size of the feature map by changing the stride of residual blocks, offering additional flexibility in the scale of semantic information we wish to imbue our representations with.

Finally, we comment that the emergent mid-level feature behavior extends readily to MoCo, and moreover also to other self-supervised methods like BYOL [21] and SimCLR [9], as the encoders for all such methods are based on ResNet-style architectures. We can thus flexibly swap out the semantic correspondence backbone for any of these image-level self-supervised representations.

3.3 Fine-grained Correspondence Learning

Only considering semantic information is not enough for correspondence learning, which often requires analyzing low-level variables such as object edge, pose, articulation, precise location and so on. Like most previous self-supervised methods, we also incorporate low-level fine-grained correspondence in our approach. BYOL-style methods [21] learn their representations by directly maximizing the similarity of two views of one image (positive pairs) in the feature space. This paradigm naturally connects with our intuitive understanding of correspondence: similar objects, parts and pixels should have similar representations. We are thus inspired to generalize this framework to a dense paradigm to learn fine-grained correspondence specifically.

At a high level, we learn our embedding space by pulling local feature vectors belonging to the same spatial region close together. Specifically, given two augmented views \mathbf{x}_1 and \mathbf{x}_2 of one image, we extract their dense feature maps $\mathbf{F}_1 \triangleq \tilde{\mathbf{f}}_\theta(\mathbf{x}_1) \in \mathbb{R}^{H \times W \times C_f}$ and $\mathbf{F}_2 \triangleq \tilde{\mathbf{f}}_\xi(\mathbf{x}_2) \in \mathbb{R}^{H \times W \times C_f}$ by removing the global pooling layer in the encoders. We adopt a ResNet-style backbone, and we can thus reduce the stride of some residual blocks in order to obtain a higher resolution feature map. In addition, to maintain dense 2D feature vectors, we replace the MLPs in the projection head and prediction head with 1×1 convolution layers. Then we can get dense prediction feature vectors $\mathbf{P}_1 \triangleq \tilde{\mathbf{q}}_\theta(\tilde{\mathbf{g}}_\theta(\mathbf{F}_1)) \in \mathbb{R}^{H \times W \times D}$ and dense projection feature vectors $\mathbf{Z}_2 \triangleq \tilde{\mathbf{g}}_\xi(\mathbf{F}_2) \in \mathbb{R}^{H \times W \times D}$. \mathbf{P}_1^i denotes the local feature vector at the i -th position of \mathbf{P}_1 . Now, a significant question remains: for a given local feature vector \mathbf{P}_1^i , how can we find its positive correspondence local feature vector in \mathbf{Z}_2 ?

Positive Correspondence Pairs Note that after we apply different spatial augmentations (random crop) to the two views, the local feature vectors on the two feature maps are no longer aligned. An object corresponding to a local feature vector in one view may even be cropped in another view. Thus, we only consider feature vectors corresponding to the same cropped region (overlapped areas of two views) and define a small spatial neighborhood around each local feature vector. All the local feature vectors in the spatial neighborhood are designated positive samples.

Specifically, we construct a binary positive mask $\mathbf{M} \in \mathbb{R}^{H \cdot W \times H \cdot W}$ by computing the spatial distance between all pairs of local feature vectors with:

$$\mathbf{M}_{ij} = \begin{cases} 1 & \text{dist}(\Phi(\mathbf{P}_1^i), \Phi(\mathbf{Z}_2^j)) \leq r \\ 0 & \text{dist}(\Phi(\mathbf{P}_1^i), \Phi(\mathbf{Z}_2^j)) > r \end{cases} \quad (4)$$

Φ denotes an operation that translates the coordinates of the local feature vector to the original image space. dist denotes the distance between coordinates of local feature vectors \mathbf{P}_1^i and \mathbf{Z}_2^j in the original image space. r is positive radius, which controls a notion of locality. As we show in the experiment, this is a very important hyperparameter. In summary, all 1s in the i -th row of \mathbf{M} represent the local feature vectors in \mathbf{Z}_2 which are positive samples of the i -th vector in \mathbf{P}_1 . This process is illustrated in Figure 3.

Learning Objectives We construct a pairwise similarity matrix \mathbf{S} , where $\mathbf{S} \in \mathbb{R}^{H \cdot W \times H \cdot W}$ with:

$$\mathbf{S}_{ij} = \text{sim}(\mathbf{P}_1^i, \mathbf{Z}_2^j) \quad (5)$$

$\text{sim}(\mathbf{u}, \mathbf{v}) = \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\|_2 \cdot \|\mathbf{v}\|_2}$ denotes the cosine similarity between two vectors. We multiply the similarity matrix \mathbf{S} and the positive mask \mathbf{M} to get the masked similarity matrix $\tilde{\mathbf{S}} = \mathbf{S} \odot \mathbf{M}$. Finally, the loss function seeks to maximize each element in the masked similarity matrix $\tilde{\mathbf{S}}$:

$$\mathcal{L}_{\text{local}} = - \frac{\sum_{i=1}^{H \cdot W} \sum_{j=1}^{H \cdot W} \tilde{\mathbf{S}}_{ij}}{\sum_{i=1}^{H \cdot W} \sum_{j=1}^{H \cdot W} \mathbf{M}_{ij}} \quad (6)$$

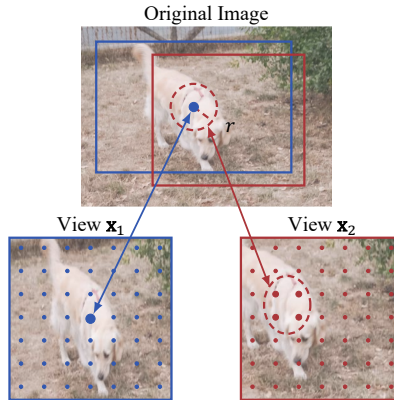


Fig. 3. For a feature vector in one view \mathbf{x}_1 , we designate all the feature vectors in view \mathbf{x}_2 which belonging to the same spatial region as positive pairs.

3.4 Fusion of Correspondence Signals

To combine semantic correspondence (Sec. 3.2) and fine-grained correspondence (Sec. 3.3) representations, one intuitive approach is simultaneously train with both semantic-level and fine-grained level losses, like [71,74,2]. However, our investigations reveal that jointly using both these objectives may not be sensible, as the representations fundamentally conflict, in two main ways. 1)*receptive fields*. We find that fine-grained correspondence relies heavily on a higher resolution feature map (see the appendix). But trivially increasing the feature resolution of a semantic-level method like MoCo [24] during training causes performance on the label propagation task to drop a lot. This is because low-level fine-grained information needs small receptive fields while relatively large receptive fields are necessary to encode global high-level semantic information. 2)*data augmentation*. Similar to VFS [75], we find that color augmentation (e.g. color distortion and grayscale conversion) is harmful to learning fine-grained correspondence,

since fine-grained correspondence heavily relies on low-level color and texture details. In contrast, image-level self-supervised learning methods learn semantic representations by imposing invariances on various data transformations. For example, as seen in the augmentations ablation for SimCLR (Fig. 5 in [9]), removing color augmentation leads to severe performance issues.

We conclude that an end-to-end framework utilizing multiple levels of supervision does not always work, especially when these modes of supervision have different requirements on both the model and data sides (see Sec. 4.4 for experimental evidence). We argue it is necessary to decouple the two models, which is consistent with how humans also attend very differently when re-identifying an object’s main body versus its accurate pixel boundary. Inspired by Two-Stream ConvNets[55], which use a late fusion to combine two kinds of complementary information, and hypercolumns [23], which effectively leverage information across different layers of CNNs, we implement a similar mechanism to fuse our orthogonal correspondences.

For a given image, suppose we have two networks, one which produces a semantic correspondence feature map $\mathbf{F}_s \triangleq \mathbf{f}_\theta(\mathbf{x}) \in \mathbb{R}^{H \times W \times C_s}$ and one which produces a fine-grained correspondence feature map $\mathbf{F}_f \triangleq \tilde{\mathbf{f}}_\theta(\mathbf{x}) \in \mathbb{R}^{H \times W \times C_f}$. Note that these two feature maps can have different channel dimensions. We consider channel-wise concatenation as a simple and intuitive way to fuse these feature maps:

$$\mathbf{F} = [\text{L2Norm}(\mathbf{F}_s), \lambda \cdot \text{L2Norm}(\mathbf{F}_f)] \quad (7)$$

where L2Norm denotes an l_2 normalization of local feature vectors in every spatial location. This ameliorates issues of scale, considering that the two feature maps are obtained under different training objectives which attend to features of different scales. λ is a hyperparameter to balance two feature maps. Note that \mathbf{F} also needs to be re-normalized when it is employed in downstream tasks, like label propagation.

3.5 Implementation Details

Any off-the-shelf image-level self-supervised pre-trained network can serve as our semantic correspondence backbone. In our implementation, we use MoCo as the default network, with ResNet-50 [25] as the base architecture and pre-trained on the 1000-class ImageNet [15] training set with strong data augmentation.

As for our fine-grained correspondence network, we use YouTube-VOS [76] as our pre-training dataset for direct comparison with previous works [34]. It contains 3471 videos totalling 5.58 hours of playtime, much smaller than Kinetics400 [8] (800 hours). Although Youtube-VOS is a video dataset, we treat it as a conventional image dataset and randomly sample individual frames during training (equivalent to 95k images). Crucially, this discards temporal information and correspondence signals our model would otherwise be able to exploit. We use *cropping-only* augmentation. Following [32,37,70], we adopt ResNet-18 as the backbone. Please see the appendix for augmentation, architecture and optimization details.

4 Experiments

We evaluate the learned representation without fine-tuning on several challenging video propagation tasks involving objects, human pose and parts. We will first introduce our detailed evaluation settings and baselines, then we conduct the comparison with the state-of-the-art self-supervised algorithms. Finally, we perform extensive ablations on different elements for SFC.

4.1 Experimental Settings

Label propagation Ideally, a model with good space-time correspondence should be able to track an arbitrary user-annotated target object throughout a video. Previous works formulate this kind of tracking task as video label propagation [70,37,32,75]. We follow the same evaluation protocol as prior art [32] for consistency and equitable comparison. At a high level, we use the representation from our pre-trained model as a similarity function. Given the ground-truth labels in the first frame, a recurrent inference strategy is applied to propagate the labels to the rest of the frames. See the appendix for detailed description.

We compare with state-of-the-art algorithms on DAVIS-2017 [49], a widely-used publicly-available benchmark for video object segmentation. To see whether our method can generalize to more visual correspondence tasks, we further evaluate our method on JHMDB benchmark [33], which involves tracking 15 human pose keypoints, and on the Video Instance Parsing (VIP) benchmark [79], which involves propagating 20 parts of the human body. We use the same settings as [32,37] and report the standard metrics, namely region-based similarity \mathcal{J} and contour-based accuracy \mathcal{F} [47] for DAVIS, probability of a correct pose (PCK) metric [77] for JHMDB and mean intersection-over-union (IoU) for VIP.

Baselines We compare with the following baselines:

Instance-Level Pre-Trained Representations: We consider supervised and self-supervised pre-trained models (MoCo, BYOL, SimSiam, etc.) on ImageNet. We also compare with two recent video-based self-supervised representation learning baselines: VINCE [18] and VFS [75]. We evaluate VFS pre-trained model using our label propagation implementation (official CRW [32] evaluation code).

Pixel-Level Pre-Trained Representations: We evaluate representations trained with pixel-level self-supervised proxy tasks: PixPro[74], DetCo[73], DenseCL[71].

Task-Specific Temporal Correspondence Representations: There are many self-supervised methods designed specifically for visual correspondence learning and evaluated on label propagation. We include these for a more comprehensive analysis: Colorization [66], CorrFlow [35], MAST [34], TimeCycle [70], UVC [37], CRW [32].

4.2 Comparison with State-Of-The-Art

We compare our method against previous self-supervised methods in Table 1. In summary, our results strongly validate the design choices in our model. In

Table 1. Video object segmentation results on the DAVIS-2017 val set. *Dataset* indicates dataset(s) used for pre-training, including: I=ImageNet, V=ImageNet-VID, C=COCO, D=DAVIS-2017, P=PASCAL-VOC, J=JHMDB. \star indicates that the method uses its own label propagation algorithm.

| Method | Supervised | Dataset (Size) | DAVIS | | | JHMDB | | VIP |
|--------------------|--------------|-------------------------------|------------------------------|-----------------|-----------------|-------------|-------------|------|
| | | | $\mathcal{J}\&\mathcal{F}_m$ | \mathcal{J}_m | \mathcal{F}_m | PCK@0.1 | PCK@0.2 | |
| Rand.Init | \times | - | 32.5 | 32.4 | 32.6 | 50.8 | 72.3 | 18.6 |
| Supervised[25] | \checkmark | I (1.28M) | 66.9 | 64.5 | 69.4 | 59.7 | 81.2 | 38.6 |
| InstDis[72] | \times | I (1.28M) | 66.4 | 63.9 | 68.9 | 58.5 | 80.2 | 32.5 |
| MoCo[24] | \times | I (1.28M) | 65.9 | 63.4 | 68.4 | 59.4 | 80.9 | 33.1 |
| SimCLR[9] | \times | I (1.28M) | 66.9 | 64.4 | 69.4 | 59.0 | 80.8 | 35.3 |
| BYOL[21] | \times | I (1.28M) | 66.5 | 64.0 | 69.0 | 58.8 | 80.9 | 34.8 |
| SimSiam[12] | \times | I (1.28M) | 67.2 | 64.8 | 68.8 | 59.9 | 81.6 | 33.8 |
| VINCE[18] | \times | Kinetics (800 hours) | 65.2 | 62.5 | 67.8 | 58.8 | 80.4 | 35.3 |
| VFS*[75] | \times | Kinetics (800 hours) | 68.9 | 66.5 | 71.3 | 60.9 | 80.7 | 43.2 |
| DetCo[73] | \times | I (1.28M) | 65.7 | 63.3 | 68.1 | 57.1 | 79.3 | 35.5 |
| DenseCL[71] | \times | I (1.28M) | 61.4 | 60.0 | 62.9 | 58.7 | 81.4 | 32.9 |
| PixPro[74] | \times | I (1.28M) | 57.5 | 56.6 | 58.3 | 57.8 | 80.8 | 29.6 |
| Colorization* [66] | \times | Kinetics (800 hours) | 34.0 | 34.6 | 32.7 | 45.2 | 69.6 | - |
| CorrFlow* [35] | \times | OxUvA (14 hours) | 50.3 | 48.4 | 52.2 | 58.5 | 78.8 | - |
| MAST* [34] | \times | YT-VOS (5.58 hours) | 65.5 | 63.3 | 67.6 | - | - | - |
| TimeCycle [70] | \times | VLOG (344 hours) | 48.7 | 46.4 | 50.0 | 57.3 | 78.1 | 28.9 |
| UVC [37] | \times | Kinetics (800 hours) | 60.9 | 59.3 | 62.7 | 58.6 | 79.6 | 34.1 |
| CRW [32] | \times | Kinetics (800 hours) | 67.6 | 64.8 | 70.2 | 58.8 | 80.3 | 37.6 |
| FC(Ours) | \times | YT-VOS (5.58 hours) | 67.7 | 64.7 | 70.5 | 59.3 | 80.8 | 34.0 |
| SFC(Ours) | \times | YT-VOS, I(5.58 hours + 1.28M) | 71.2 | 68.3 | 74.0 | 61.9 | 83.0 | 38.4 |
| OSVOS [5] | \checkmark | I/D (1.28M + 10k) | 60.3 | 56.6 | 63.9 | - | - | - |
| OnAVOS [65] | \checkmark | I/C/P/D (1.28M + 517k) | 65.4 | 61.6 | 69.1 | - | - | - |
| FEEVOS [64] | \checkmark | I/C/D/YT-VOS (1.28M + 663k) | 71.5 | 69.1 | 74.0 | - | - | - |
| PAAP [31] | \checkmark | I/J (1.28M + 32K) | - | - | - | 51.6 | 73.8 | - |
| Thin-Slicing [56] | \checkmark | I/J (1.28M + 32K) | - | - | - | 68.7 | 92.1 | - |
| ATEN [79] | \checkmark | VIP (20k) | - | - | - | - | - | 37.9 |

particular, the full semantic-aware fine-grained correspondence network (SFC), achieves state-of-the-art performance on all tasks investigated. SFC significantly outperforms other methods that learn only semantic correspondence (MoCo, 65.9 \rightarrow 71.2 on DAVIS-2017) or only fine-grained correspondence (FC, 67.7 \rightarrow 71.2 on DAVIS-2017). SFC even outperforms several supervised baselines specially designed for video object segmentation and human part tracking.

Note also that our fine-grained correspondence network (FC) can achieve comparable performance on DAVIS and JHMDB with methods like CRW, despite training with far less data and discarding temporal information. The performance of FC on VIP is lower, but it may be further improved by exploiting more inductive bias, e.g., temporal context or viewpoint changes in videos.

We show the results on DAVIS-2017 of FC using different pre-training datasets in the appendix. FC achieves 67.9 $\mathcal{J}\&\mathcal{F}_m$ when pre-trained on ImageNet. This suggests that a larger dataset offers marginal benefits for fine-grained correspondence learning, which is largely different from learning semantic correspondence. When replacing YouTube-VOS pre-trained FC with ImageNet pre-trained one, SFC still achieves 71.3 $\mathcal{J}\&\mathcal{F}_m$. This indicates that the performance gain of SFC doesn't come from the extra YouTube-VOS dataset. We use YouTube-VOS for faster training and fair comparisons of other correspondence learning methods.

We also report results of SFC on *semantic segmentation* and ImageNet-1K *linear probing* in the appendix. Our SFC achieves improved results on all con-




Fig. 4. Qualitative results for label propagation. Given ground-truth labels in the first frame (outlined in blue), our method can propagate them to the rest of frames. For more results, please refer to the appendix.

sidered tasks, showing strong generalization ability and the flexibility of our core contribution.

4.3 Visualization

Figure 4 shows samples of video label propagation results. We further visualize the learned correspondences of our model in Figure 5, compared with its components, MoCo and FC. We notice that the correspondence map of MoCo tends to scatter across the entire visual object, indicating that it focuses more on object-level semantics instead of low-level fine-grained features. On the contrary, the correspondence map of FC is highly concentrated, but sometimes loses track of the source pixel, indicating a failure to capture high-level semantics. By balancing semantics and fine-grained correspondences, our proposed method SFC is able to overcome their respective drawbacks and give the most accurate correspondence.




Fig. 5. Correspondence map of SFC, compared with MoCo and FC. Given the source frame with one pixel highlighted in red, we calculate the feature similarity between the target frame and this pixel. Red regions indicate high similarity.

4.4 Ablative Analysis

In this section, we investigate our results on video object segmentation using DAVIS-2017 in more detail, and outline several ablations on important design choices throughout our model architectures and pipelines.

Fusion Strategy We perform experiments by combining the FC training objective with a global image-level loss, resulting in an end-to-end multi-task framework. But we find the two losses fail to boost performance synergistically. For example, when we add a BYOL loss to FC for joint optimization (see the appendix for details), the performance on DAVIS-2017 drops a little ($67.7 \rightarrow 67.2$). The reason is that the two losses need different receptive fields and augmentations. The optimal configuration of FC model will induce a sub-optimal solution under the image-level loss, and vice versa. Thus, it is sensible to train two independent models and use concatenation to fuse the two different kinds of representations.

One may expect the concatenation operation is some form of model ensemble. Does combining an arbitrary two networks lead to any reasonable improvement in performance? To answer this, we conduct experiments on two sets of models: in the first set, all models have two semantic correspondence networks; while in the second set, all models have two fine-grained correspondence networks. Results are shown in Table 2. We observe that if two networks have the same type of correspondence, their combination leads to unremarkable increases in performance.

In the appendix, we show that we can flexibly replace semantic correspondence backbone (MoCo \rightarrow InstDis, SimCLR, BYOL, etc.) and still maintain strong performance on DAVIS. This strongly confirms our hypothesis that image-level self-supervised representations in general contain information about semantic correspondence. It also supports our framing of semantic correspondence and fine-grained correspondence as orthogonal sources of information.

Next, we mainly conduct a series of ablation studies on our fine-grained network (FC).

Crop Size and Positive Radius When we apply cropping for an image, a random patch is selected, with an area uniformly sampled between γ_1 (lower bound) and γ_2 (upper bound) of that of the original image. In Figure 6, we plot FC model performance on different ratios of crop size area, by varying γ_1 : $\{0, 0.08, 0.2, 0.3\}$ and fixing γ_2 to 1. Simultaneously, for every lower bound γ_1 , we investigate how different positive radii r can also affect performance on correspondence learning.

Table 2. Fusion of two networks with the same kind of correspondence.

FC denotes our fine-grained correspondence network, it achieves $67.7 \mathcal{J} \& \mathcal{F}_m$ on DAVIS-2017. Other single model results can be found in Table 1.

| Type | Combination | $\mathcal{J} \& \mathcal{F}_m$ | \mathcal{J}_m | \mathcal{F}_m |
|--------------------------------|----------------|--------------------------------|-----------------|-----------------|
| Semantic Correspondence | InstDis + MoCo | 67.4 | 65.0 | 69.7 |
| | SimCLR + MoCo | 67.2 | 64.7 | 69.6 |
| | BYOL + MoCo | 67.5 | 65.0 | 70.0 |
| | Simsiam + MoCo | 67.4 | 65.1 | 69.6 |
| | VINCE + MoCo | 66.7 | 64.2 | 69.2 |
| | VFS + MoCo | 68.1 | 66.1 | 70.2 |
| Fine-grained Correspondence | TimeCycle + FC | 67.8 | 65.2 | 70.5 |
| | UVC + FC | 61.5 | 59.8 | 63.3 |
| | CRW + FC | 68.8 | 65.7 | 71.9 |




Fig. 6. Effect of random crop size and positive radius.

We find that as the lower bound γ_1 increases, mode performance worsens. $\gamma_1 = 0$ yields relatively strong performance under a wide range of positive radius r . We conjecture that using a small lower bound γ_1 results in larger scale and translation variations between two views of one image, which induces strong spatial augmentation and thus allows our correspondence learning to rely on scale-invariant representations. We also observe that an appropriate positive radius r is crucial for fine-grained correspondence learning. On the DAVIS dataset, we show that a large (smooth) or small (sharp) r is demonstrably harmful to performance. Finally, for different γ_1 , the optimal value of r is different.

Data Augmentation VFS [75] has pointed out that color augmentation jeopardizes fine-grained correspondence learning. To systematically study the effects of individual data augmentations, we investigate the performance of our FC model on DAVIS when applying random cropping and another common augmentation (random flip, color jittering, etc.). We report the results in Table 3. Among all color data augmentations, the one that has the greatest negative impact on fine-grained correspondence learning is actually color dropping (grayscale conversion). This is in contrast to image-level self-supervised learning, where strong color augmentation [9] is crucial for learning good representations. We adopt random crop as the only augmentation in our best-performing models.

Table 3. Effect of data augmentation. The augmentation strategy follows MoCo v2 [11].

| Augmentation | $\mathcal{J} \& \mathcal{F}_m$ | \mathcal{J}_m | \mathcal{F}_m |
|-------------------------------|--------------------------------|-----------------|-----------------|
| Random crop | 67.6 | 64.7 | 70.5 |
| Random crop & Random flip | 67.7 | 64.8 | 70.6 |
| Random crop & Color jittering | 65.5 | 62.6 | 68.4 |
| Random crop & Gaussian blur | 65.9 | 63.2 | 68.6 |
| Random crop & Color dropping | 61.9 | 58.9 | 64.9 |

5 Conclusion and Discussions

We have developed a novel framework to learn both semantic and fine-grained correspondence from still images alone. We demonstrate that these two forms of correspondence offer complementary information, thereby facilitating a simple yet intuitive fusion scheme which leads to state-of-the-art results on a number of downstream correspondence tasks. In this work, we mainly explore the correspondence properties of ConvNet. Whether ViT [16] also benefits from dense fine-grained self-supervision and combination of two kinds of correspondence is an interesting open question left to future exploration.

Acknowledgements.

This work is supported by the Ministry of Science and Technology of the People’s Republic of China, the 2030 Innovation Megaprojects “Program on New Generation Artificial Intelligence” (Grant No. 2021AAA0150000). This work is also supported by a grant from the Guoqiang Institute, Tsinghua University.

References

1. Bachman, P., Hjelm, R.D., Buchwalter, W.: Learning representations by maximizing mutual information across views. arXiv preprint arXiv:1906.00910 (2019)
2. Bai, Y., Chen, X., Kirillov, A., Yuille, A., Berg, A.C.: Point-level region contrast for object detection pre-training. arXiv preprint arXiv:2202.04639 (2022)
3. Ballard, D.H., Zhang, R.: The hierarchical evolution in human vision modeling. *Topics in Cognitive Science* **13**(2), 309–328 (2021)
4. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: European conference on computer vision. pp. 850–865. Springer (2016)
5. Caelles, S., Maninis, K.K., Pont-Tuset, J., Leal-Taixé, L., Cremers, D., Van Gool, L.: One-shot video object segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 221–230 (2017)
6. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. arXiv preprint arXiv:2006.09882 (2020)
7. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9650–9660 (2021)
8. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
10. Chen, T., Luo, C., Li, L.: Intriguing properties of contrastive losses. arXiv preprint arXiv:2011.02803 (2020)
11. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
12. Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021)
13. Chen, Y.C., Huang, P.H., Yu, L.Y., Huang, J.B., Yang, M.H., Lin, Y.Y.: Deep semantic matching with foreground detection and cycle-consistency. In: Asian Conference on Computer Vision. pp. 347–362. Springer (2018)
14. Cheng, H.K., Tai, Y.W., Tang, C.K.: Rethinking space-time networks with improved memory coverage for efficient video object segmentation. arXiv preprint arXiv:2106.05210 (2021)
15. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
16. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
17. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: Flownet: Learning optical flow with convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2758–2766 (2015)

18. Gordon, D., Ehsani, K., Fox, D., Farhadi, A.: Watching the world go by: Representation learning from unlabeled videos. arXiv preprint arXiv:2003.07990 (2020)
19. Gould, S., Arfvidsson, J., Kaehler, A., Sapp, B., Messner, M., Bradski, G.R., Baumstarck, P., Chung, S., Ng, A.Y., et al.: Peripheral-foveal vision for real-time object recognition and tracking in video. In: *Ijcai*. vol. 7, pp. 2115–2121 (2007)
20. Grabner, H., Matas, J., Van Gool, L., Cattin, P.: Tracking the invisible: Learning where the object might be. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 1285–1292. IEEE (2010)
21. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., et al.: Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint arXiv:2006.07733 (2020)
22. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). vol. 2, pp. 1735–1742. IEEE (2006)
23. Hariharan, B., Arbeláez, P., Girshick, R., Malik, J.: Hypercolumns for object segmentation and fine-grained localization. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 447–456 (2015)
24. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9729–9738 (2020)
25. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
26. Held, D., Thrun, S., Savarese, S.: Learning to track at 100 fps with deep regression networks. In: *European conference on computer vision*. pp. 749–765. Springer (2016)
27. Henaff, O.: Data-efficient image recognition with contrastive predictive coding. In: *International Conference on Machine Learning*. pp. 4182–4192. PMLR (2020)
28. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. *IEEE transactions on pattern analysis and machine intelligence* **37**(3), 583–596 (2014)
29. Huang, S., Wang, Q., Zhang, S., Yan, S., He, X.: Dynamic context correspondence network for semantic alignment. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2010–2019 (2019)
30. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: Flownet 2.0: Evolution of optical flow estimation with deep networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2462–2470 (2017)
31. Iqbal, U., Garbade, M., Gall, J.: Pose for action-action for pose. In: 2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017). pp. 438–445. IEEE (2017)
32. Jabri, A., Owens, A., Efros, A.A.: Space-time correspondence as a contrastive random walk. *Advances in Neural Information Processing Systems* (2020)
33. Jhuang, H., Gall, J., Zuffi, S., Schmid, C., Black, M.J.: Towards understanding action recognition. In: *Proceedings of the IEEE international conference on computer vision*. pp. 3192–3199 (2013)
34. Lai, Z., Lu, E., Xie, W.: Mast: A memory-augmented self-supervised tracker. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6479–6488 (2020)
35. Lai, Z., Xie, W.: Self-supervised learning for video correspondence flow. arXiv preprint arXiv:1905.00875 (2019)

36. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8971–8980 (2018)
37. Li, X., Liu, S., De Mello, S., Wang, X., Kautz, J., Yang, M.H.: Joint-task self-supervised learning for temporal correspondence. arXiv preprint arXiv:1909.11895 (2019)
38. Liu, C., Yuen, J., Torralba, A., Sivic, J., Freeman, W.T.: Sift flow: Dense correspondence across different scenes. In: European conference on computer vision. pp. 28–42. Springer (2008)
39. Liu, Y., Zhu, L., Yamada, M., Yang, Y.: Semantic correspondence as an optimal transport problem. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4463–4472 (2020)
40. Long, J.L., Zhang, N., Darrell, T.: Do convnets learn correspondence? Advances in neural information processing systems **27**, 1601–1609 (2014)
41. Min, J., Cho, M.: Convolutional hough matching networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2940–2950 (2021)
42. Min, J., Lee, J., Ponce, J., Cho, M.: Hyperpixel flow: Semantic correspondence with multi-layer neural features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3395–3404 (2019)
43. Min, J., Lee, J., Ponce, J., Cho, M.: Learning to compose hypercolumns for visual correspondence. In: Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16. pp. 346–363. Springer (2020)
44. Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6707–6717 (2020)
45. Oh, S.W., Lee, J.Y., Xu, N., Kim, S.J.: Video object segmentation using space-time memory networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9226–9235 (2019)
46. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
47. Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 724–732 (2016)
48. Pinheiro, P.O., Almahairi, A., Benmalek, R.Y., Golemo, F., Courville, A.: Unsupervised learning of dense visual representations. arXiv preprint arXiv:2011.05499 (2020)
49. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675 (2017)
50. Purushwalkam, S., Gupta, A.: Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. arXiv preprint arXiv:2007.13916 (2020)
51. Ranjan, A., Black, M.J.: Optical flow estimation using a spatial pyramid network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4161–4170 (2017)
52. Rocco, I., Arandjelovic, R., Sivic, J.: Convolutional neural network architecture for geometric matching. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6148–6157 (2017)

53. Rocco, I., Arandjelović, R., Sivic, J.: End-to-end weakly-supervised semantic alignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6917–6925 (2018)
54. Rocco, I., Cimpoi, M., Arandjelović, R., Torii, A., Pajdla, T., Sivic, J.: Neighbourhood consensus networks. arXiv preprint arXiv:1810.10510 (2018)
55. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems* **27** (2014)
56. Song, J., Wang, L., Van Gool, L., Hilliges, O.: Thin-slicing network: A deep structured model for pose estimation in videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4220–4229 (2017)
57. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8934–8943 (2018)
58. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: European conference on computer vision. pp. 402–419. Springer (2020)
59. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*. pp. 776–794. Springer (2020)
60. Truong, P., Danelljan, M., Timofte, R.: Glu-net: Global-local universal network for dense flow and correspondences. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6258–6268 (2020)
61. Valmadre, J., Bertinetto, L., Henriques, J., Vedaldi, A., Torr, P.H.: End-to-end representation learning for correlation filter based tracking. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2805–2813 (2017)
62. Van Gansbeke, W., Vandenhende, S., Georgoulis, S., Proesmans, M., Van Gool, L.: Scan: Learning to classify images without labels. In: *European Conference on Computer Vision*. pp. 268–285. Springer (2020)
63. Van Gansbeke, W., Vandenhende, S., Georgoulis, S., Van Gool, L.: Revisiting contrastive methods for unsupervised learning of visual representations. arXiv preprint arXiv:2106.05967 (2021)
64. Voigtlaender, P., Chai, Y., Schroff, F., Adam, H., Leibe, B., Chen, L.C.: Feelvos: Fast end-to-end embedding learning for video object segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9481–9490 (2019)
65. Voigtlaender, P., Leibe, B.: Online adaptation of convolutional neural networks for video object segmentation. arXiv preprint arXiv:1706.09364 (2017)
66. Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos. In: Proceedings of the European conference on computer vision (ECCV). pp. 391–408 (2018)
67. Wang, N., Yeung, D.Y.: Learning a deep compact image representation for visual tracking. *Advances in neural information processing systems* (2013)
68. Wang, N., Song, Y., Ma, C., Zhou, W., Liu, W., Li, H.: Unsupervised deep tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1308–1317 (2019)
69. Wang, Q., Zhang, L., Bertinetto, L., Hu, W., Torr, P.H.: Fast online object tracking and segmentation: A unifying approach. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1328–1338 (2019)
70. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2566–2576 (2019)

71. Wang, X., Zhang, R., Shen, C., Kong, T., Li, L.: Dense contrastive learning for self-supervised visual pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3024–3033 (2021)
72. Wu, Z., Xiong, Y., Yu, S., Lin, D.: Unsupervised feature learning via non-parametric instance-level discrimination. arXiv preprint arXiv:1805.01978 (2018)
73. Xie, E., Ding, J., Wang, W., Zhan, X., Xu, H., Sun, P., Li, Z., Luo, P.: Detco: Unsupervised contrastive learning for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8392–8401 (2021)
74. Xie, Z., Lin, Y., Zhang, Z., Cao, Y., Lin, S., Hu, H.: Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16684–16693 (2021)
75. Xu, J., Wang, X.: Rethinking self-supervised correspondence learning: A video frame-level similarity perspective. arXiv preprint arXiv:2103.17263 (2021)
76. Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., Huang, T.: Youtube-vos: A large-scale video object segmentation benchmark. arXiv preprint arXiv:1809.03327 (2018)
77. Yang, Y., Ramanan, D.: Articulated pose estimation with flexible mixtures-of-parts. In: CVPR 2011. pp. 1385–1392. IEEE (2011)
78. Zhang, R., Saran, A., Liu, B., Zhu, Y., Guo, S., Niekum, S., Ballard, D., Hayhoe, M.: Human gaze assisted artificial intelligence: a review. In: IJCAI: Proceedings of the Conference. vol. 2020, p. 4951. NIH Public Access (2020)
79. Zhou, Q., Liang, X., Gong, K., Lin, L.: Adaptive temporal encoding network for video instance-level human parsing. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 1527–1535 (2018)