

# Locality-Aware Inter-and Intra-Video Reconstruction for Self-Supervised Correspondence Learning

Liulei Li<sup>1,6\*</sup>, Tianfei Zhou<sup>2</sup>, Wenguan Wang<sup>3†</sup>, Lu Yang<sup>4</sup>, Jianwu Li<sup>1</sup>, Yi Yang<sup>5</sup>

<sup>1</sup> Beijing Institute of Technology <sup>2</sup> ETH Zurich <sup>3</sup> ReLER, AAIL, University of Technology Sydney

<sup>4</sup> Beijing University of Posts and Telecommunications <sup>5</sup> CCAI, Zhejiang University <sup>6</sup> Baidu Research

<https://github.com/01lililulei/LIIR>

## Abstract

Our target is to learn visual correspondence from unlabeled videos. We develop LIIR, a *locality-aware inter-and intra-video reconstruction method that fills in three missing pieces*, i.e., *instance discrimination, location awareness, and spatial compactness, of self-supervised correspondence learning puzzle*. First, instead of most existing efforts focusing on intra-video self-supervision only, we exploit cross-video affinities as extra negative samples within a unified, inter-and intra-video reconstruction scheme. This enables instance discriminative representation learning by contrasting desired intra-video pixel association against negative inter-video correspondence. Second, we merge position information into correspondence matching, and design a position shifting strategy to remove the side-effect of position encoding during inter-video affinity computation, making our LIIR location-sensitive. Third, to make full use of the spatial continuity nature of video data, we impose a compactness-based constraint on correspondence matching, yielding more sparse and reliable solutions. The learned representation surpasses self-supervised state-of-the-arts on label propagation tasks including objects, semantic parts, and keypoints.

## 1. Introduction

As a fundamental problem in computer vision, correspondence matching facilitates many applications, such as scene understanding [71], object dynamics modeling [27], and 3D reconstruction [19]. However, supervising representation for visual correspondence is not trivial, as obtaining pixel-level manual annotations is costly, and sometimes even prohibitive (due to occlusions and free-form object deformations). Although synthetic data would serve an alternative in some low-level visual correspondence tasks (e.g., optical flow estimation [2]), they limit the generalization to real scenes.

Using natural videos as a source of free supervision, i.e.,

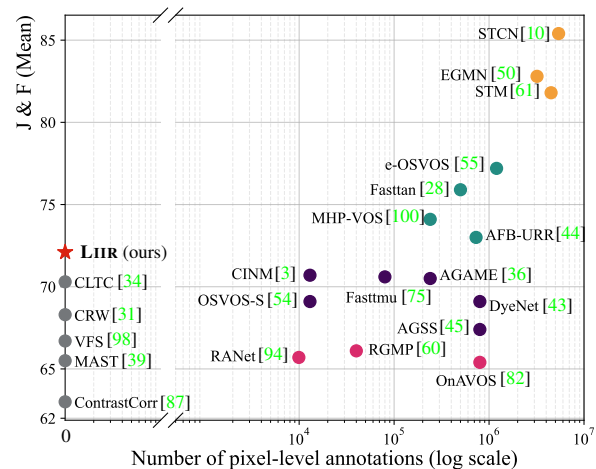


Figure 1. **Performance comparison** over DAVIS<sub>17</sub> val. Our LIIR surpasses all existing self-supervised methods, and is on par with many fully-supervised ones trained with massive annotations.

*self-supervised temporal correspondence learning*, is considered as appealing [40]. This is because videos contain rich realistic appearance and shape variations with almost infinite supply, and deliver valuable supervisory signals from the intrinsic coherence, i.e., correlations among frames.

Along this direction, existing solutions are typically built upon a *reconstruction* scheme (i.e., each pixel from a ‘query’ frame is reconstructed by finding and assembling relevant pixels from adjacent frame(s)) [39, 40, 84], and/or adopt a *cycle-consistent tracking* paradigm (i.e., pixels/patches are encouraged to fall into the same location after one cycle of forward and backward tracking) [31, 42, 52, 86, 90].

Unfortunately, these successful approaches largely neglect three crucial abilities for robust temporal correspondence learning, namely **instance discrimination**, **location awareness**, and **spatial compactness**. First, many of them share a narrow view that only considers intra-video context for correspondence learning. As it is hard to derive a free signal from a single video for identifying different object instances, the learned features are inevitably less instance-discriminative. Second, existing methods are typically built without explicit position representation. Such design seems

\*Work done during an internship at Baidu Research.

†Corresponding author: Wenguan Wang.




Figure 2. **Illustration of different self-supervised architectures for temporal correspondence learning (§2):** (a) reconstruction based, (b) cycle-consistency based, and (c) our LIIR that addresses instance discrimination, location awareness, and spatial compactness.

counter-intuitive, given the extensive evidence that spatial position is encoded in human visual system [25] and plays a vital role when human track objects [62]. Third, as the visual world is continuous and smoothly-varying, both spatial and temporal coherence naturally exist in videos. While numerous strategies are raised to address smoothness on the time axis, far less attention has been paid to the spatial case.

To fill in these three missing pieces to the puzzle of self-supervised correspondence learning, we present a locality-aware inter-and intra-video reconstruction framework – LIIR. **First**, we augment existing intra-video analysis based correspondence learning strategy with *inter-video context*, which is informative for instance-level separation. This leads to an inter-and intra-video reconstruction based training objective, that *inspires* intra-video positive correspondence matching, but *penalizes* unreliable pixel associations within and cross videos. We empirically verify that our inter-and intra-video reconstruction strategy can yield more discriminative features, that encode higher-level semantics beyond low-level intra-instance invariance modeled by previous algorithms. **Second**, to make our LIIR more location-sensitive, we learn to encode position information into the representation. Although position bias is favored for intra-video correspondence matching, it is undesired in the inter-video case. We thus devise a *position shifting* strategy to foster the strength and circumvent the weaknesses of position encoding. We experimentally show that, explicit position embedding benefits correspondence matching. **Third**, we involve a spatial compactness prior in intra-video pixel-wise affinity estimation, resulting in sparse yet compact associations. For each query pixel, the distribution of related pixels is fit by a mixture of Gaussians. This enforces each query pixel to match only a handful of spatially close pixels in an adjacent frame. Our experiments show that such compactness prior not only regularizes training, but also removes outliers during inference.

These three contributions together make LIIR a powerful framework for self-supervised correspondence learning. Without any adaptation, the learned representation is effective for various correspondence-related tasks, *i.e.*, video object segmentation, semantic part propagation, pose tracking. On these tasks, LIIR consistently outperforms unsupervised state-of-the-arts and is comparable to, or even better than, some task-specific fully-supervised methods (*e.g.*, Fig. 1).

## 2. Related Work

**Self-Supervised Temporal Correspondence Learning.** In the video domain, correspondence matching plays a central role in many tasks (*e.g.*, video segmentation [27], flow estimation [15, 29] and object tracking [4]). An emerging line of work tackles this problem in a self-supervised learning paradigm, by exploiting the temporal coherence in videos. One may group these work into two major classes. The first class of methods [39, 40, 84] poses a *colorization* proxy task (Fig. 2(a)), *i.e.*, reconstruct a query frame from an adjacent frame, according to their correspondence. The latter type of methods [31, 42, 52, 86, 90] performs forward and backward tracking and penalizes the inconsistency between the start and end positions of the tracked pixels or regions (Fig. 2(b)). The basic idea – *cycle-consistency* – is also adopted in unsupervised tracking [86, 106], optical flow [56, 111] and depth estimation [33, 103]. Though impressive, these methods miss three key elements for robust correspondence matching: instance discrimination, location awareness, and spatial compactness. In response, LIIR is equipped with three specific modules (Fig. 2(c)). First, for instance discriminative representation learning, it adopts an inter-and intra-video reconstruction scheme that formulates contrast over inter-and intra-video affinities. Second, it involves position encoding into representation learning. Third, it imposes a spatial compactness prior to both correspondence learning and inference.

We are not the first to explore inter-video context. In [52], Lu *et al.* raise an unsupervised learning objective, *i.e.*, discriminate between a set of surrogate video classes, but they compute this over video-level embeddings, not sufficient for pixel-wise correspondence learning. In [87], Wang *et al.* also simultaneously consider inter-and intra-video representation associations, but requiring pre-aligned *patch pairs*. In addition, they use three loss terms to address the desired intra-inter constraints, which are, however, formulated as a single unified training objective in our case. Moreover, we take a further step by addressing location awareness and spatial compactness, instead inter-video context only. In [98], Xu *et al.* revisit the idea of image-level similarity learning and consider frame pairs from the same videos are positive samples and pairs from different videos are negative, however, they find negative samples (*i.e.*, inter-video context) hurt the

performance of their model. In contrast, we formulate inter- and intra-video context through a unified, pixel-wise affinity framework that boosts instance-level discrimination without sacrificing intra-instance invariance. Our results suggest that how to make a good use of negative samples for temporal correspondence learning is still an intriguing question.

**Self-Supervised Video Representation Learning.** Correspondence learning approaches that use unlabeled video data fall in a broad field of self-supervised video representation learning. Towards learning transferable video representation, diverse pretext tasks are proposed to explore different intrinsic properties of videos as free supervisory signals, including temporal sequence ordering [17, 58, 95], predicting motion patterns [1, 18, 65, 78, 85], solving space-time cubic puzzles [38], anticipating future representations [49, 83], and temporally aligning videos [16]. The learned representations are compact video descriptors that can be generalized to various downstream tasks (*e.g.*, action recognition [21, 22, 69, 96], video captioning [74, 110], video retrieval [57]), while in this work we specifically focus on learning fine-grained visual features for pixel-level correspondence matching.

**Self-Supervised Image Representation Learning.** Basically, self-supervised approaches for image representation learning share a similar key idea with the counterparts for videos – design pretext tasks so as to mine supervision signals from the inherent information inside images [102]. Examples of such tasks include estimating spatial context [13], predicting image rotations [20], solving jigsaw puzzles [59], among many others [41, 66, 92]. Recent efforts were mainly devoted to improving deep metric learning techniques in large scale, *i.e.*, adopt an instance discrimination pretext task where the cross-entropy objective is used to discriminate each image from other different images (*i.e.*, negative samples) [7, 8, 23, 26, 63]. In this work, we assimilate the idea of contrasting positive samples (intra-video affinities) over numerous negative ones (inter-video affinities) for robust temporal correspondence matching, and formulate this within a unified inter- and intra-video reconstruction framework.

**Video Mask Propagation.** This task, a.k.a. semi-automatic video segmentation, aims at propagating first-frame object masks over the whole video sequence [77, 88]. It addresses classic automatic video segmentation techniques’ [51, 64, 89, 90, 108, 109] lack of flexibility in defining target objects [91]. Depending on how they make use of the first-frame supervision, existing mask propagation models can be classified into three groups: i) *online fine-tuning* based methods that fine-tune a generic segmentation network with masks during inference [5, 82]; and ii) *matching* based methods that directly condition the segmentation network on the first-frame masks and/or previous segments [32, 50, 61, 67, 80, 104]. As shown in Fig. 1, extensive annotations are typically required to train such systems, *i.e.*, first pre-train on ImageNet [70] and then fine-tune with COCO [46], DAVIS [68], Youtube-

VOS [99], *etc.* In contrast, we pursue a more annotation-efficient solution; similar to previous self-supervised correspondence learning methods [39, 40, 52, 84], LIIR is trained using unlabeled videos only. Once trained, it can be directly applied for mask propagation, without adaptation.

### 3. Our Approach

We present LIIR, a self-supervised framework that learns dense correspondence from raw videos. Before elaborating on our model design (*cf.* §3.2), we first review the classic reconstruction based temporal correspondence learning strategy (*cf.* §3.1), which serves as the basis of our LIIR.

#### 3.1. Preliminary: Learning Temporal Correspondence through Frame Reconstruction

Due to the appearance continuity in video, one can consider pixels in a ‘query’ frame as being *copied* from some locations of other ‘reference’ frames. In light of this, a few studies [40, 84] raise a reconstruction-based correspondence learning scheme: each query pixel struggles to find pixels in a reference frame that can best reconstruct itself.

Formally, let  $I_q, I_r \in \mathbb{R}^{H \times W \times 3}$  respectively denote a query frame and a reference frame from the same video. They are projected into a pixel embedding space by a ConvNet encoder (*e.g.*, ResNet [24])  $\phi: \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^{h \times w \times c}$ , such that  $\mathbf{I}_q, \mathbf{I}_r = \phi(I_q), \phi(I_r)$ . The *copy* operator can be approximated as an inter-frame affinity matrix  $A \in [0, 1]^{hw \times hw}$ :

$$A(i, j) = \frac{\exp(\mathbf{I}_q(i) \cdot \mathbf{I}_r(j))}{\sum_{j'} \exp(\mathbf{I}_q(i) \cdot \mathbf{I}_r(j'))}, \quad i, j \in \{1, \dots, hw\} \quad (1)$$

where  $A(i, j) \in [0, 1]$  refers to  $(i, j)$ -th element in  $A$ , signifying the similarity between pixel  $i$  in  $I_q$  and pixel  $j$  in  $I_r$ , and ‘ $\cdot$ ’ stands for the dot product. In this way,  $A$  gives the strength of all the pixel pair-wise correspondence between  $I_q$  and  $I_r$ , according to which pixel  $i$  in  $I_q$  can be reconstructed by a weighted sum of pixels in  $I_r$ :

$$\hat{I}_q(i) = \sum_j A(i, j) I_r(j). \quad (2)$$

The training objective of  $\phi$  is hence defined as a reconstruction loss:

$$\mathcal{L}_{\text{res}} = \|I_q - \hat{I}_q\|_2. \quad (3)$$

In practice, to avoid trivial solutions caused by information leakage, an *information bottleneck* is adopted over training samples, *e.g.*, RGB2gray operation [84], channel-wise dropout in RGB [40] or Lab [39] colorspace. After training, the representation encoder  $\phi$  is used for correspondence matching: similar to Eq. 2, the affinity  $A$  is estimated and used to *propagate* desired pixel-level entities (*e.g.*, instance masks, key-point maps), from a reference frame to a query frame.

#### 3.2. LIIR: Locality-Aware Inter- and Intra-Video Reconstruction Framework

Building upon the reconstruction-by-copy scheme, LIIR is empowered with three crucial yet long overlooked abili-

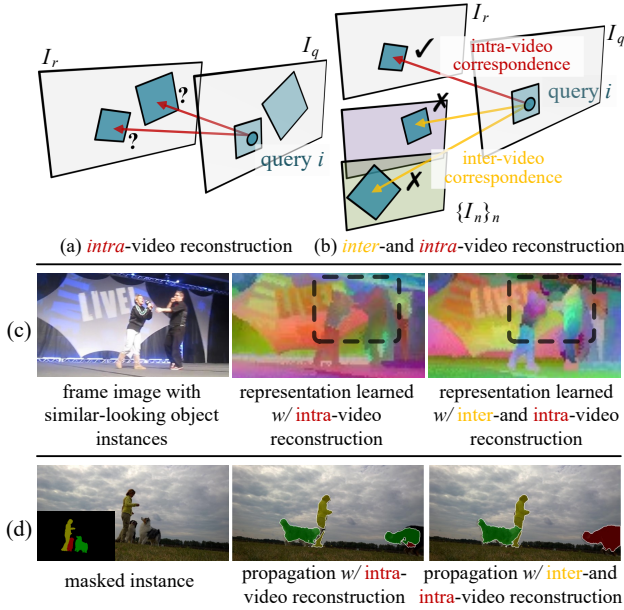


Figure 3. **Inter-and intra-video reconstruction** (§3.2). (a) Previous intra-video reconstruction based approaches struggle to offer supervisory signal for distinguishing between different instances. (b) In our inter-and intra-video reconstruction, each query pixel is forced to distinguish intra-video correspondence ( $\rightarrow$ ) from negative inter-video association ( $\leftrightarrow$ ), enabling cross-instance discrimination. (c)-(d) Representation learned with inter-and intra-video reconstruction is more robust for multiple object instances.

ties for robust correspondence learning: instance discrimination, location awareness, and spatial compactness.

**Inter-and Intra-Video Reconstruction.** With the computation of the intra-video affinity  $A$  (Eq. 1), each query pixel is forced to distinguish its counterpart (positive) reference pixels from unrelated (negative) ones *within a same video*, with the indicator of the reconstruction quality  $\mathcal{L}_{\text{res}}$  (Eqs. 2-3). As both the positive and negative samples are sourced from the same video, there is less evidence for distinction among similar object instances with intra-video appearance only (Fig. 3(a)). As one single video only contains limited content, conducting correspondence matching within videos is less challenging, and inevitably hinders the discrimination potential of the learned representation [87]. These insights motivate us to improve the intra-video affinity based reconstruction scheme by further accounting for negative *across-video* correspondence. Concretely, given the query ( $I_q$ ) and reference ( $I_r$ ) frames from the same video, an *intra-inter video affinity*  $A' \in [0, 1]^{hw \times hw}$  is computed:

$$A'(i, j) = \frac{\exp(I_q(i) \cdot I_r(j))}{\underbrace{\sum_{j'} \exp(I_q(i) \cdot I_r(j'))}_{\text{intra-video correspondence}} + \underbrace{\sum_n \sum_k \exp(I_q(i) \cdot I_n(k))}_{\text{inter-video correspondence}}} \quad (4)$$

where  $\{I_n\}_n$  refer to a collection of frames, which are sampled from the whole training dataset, except the source video of  $I_q$  ( $I_r$ ). By additionally considering other irrele-

vant videos during affinity computation, both the quantity and diversity of negative samples are greatly improved, allowing us to derive a more challenging *inter-and intra-video reconstruction* scheme (Fig. 3(b)):

$$\hat{I}_q(i) = \sum_j A'(i, j) I_r(j). \quad (5)$$

With Eqs. 4-5, each pixel  $i$  in the query frame  $I_q$  is required to distinguish its counterpart pixels from massive unrelated ones, which are from not only the reference frame  $I_r$  in current video, but a huge amount of irrelevant frames  $\{I_n\}_n$  in other videos. This powerful idea, yet, is elegantly achieved by the same training objective as in Eq. 3. Note that Eq. 4 normalizes intra-video correspondence over both inter-and intra-video pixel-to-pixel relevance, while Eq. 5 only uses the pixels in the reference frame  $I_r$  for reconstruction. The rationale here is that, even if the encoder  $\phi$  wrongly matches a query pixel  $i$  with a negative but similar-looking pixel  $k$  in  $I_n$ , i.e.,  $\exp(I_q(i) \cdot I_n(k))$  will be large and  $\sum_j A'(i, j) \ll 1$ , the synthesized color  $\hat{I}_q(i)$  will be *still* very different to  $I_q(i)$  and  $\phi$  will receive a large gradient from Eq. 3. Thus  $\phi$  is driven to mine more high-level semantics and context-related clues, hence reinforcing the instance-level discrimination ability (Fig. 3(c)). Fig. 3(d) shows that the representation learned with our inter-and intra-video reconstruction strategy can distinguish similar-looking dogs nearby.

Although [87] also addresses inter-video analysis based reconstruction, it conducts embedding association on *patch level*, relying on a pre-trained tracker for patch alignment. Besides, the method consumes three loss terms for supervision, which is much more complicated than ours. Further, it separately conducts inter-and intra-video affinity based reconstruction. This is problematic; when both the reference frame in current video and an irrelevant frame from other videos contain query-like pixels, there is no explicit supervision signal to determine which one should be matched.

**Position Encoding and Position Shifting.** Plenty of literature in neuroscience has revealed that human visual system encodes both appearance and position information when we perceive and track objects [25, 48, 62]. Yet existing unsupervised correspondence methods put all focus on improving appearance based representation by ConvNets, ignoring the value of position information. Although [30, 37] suggest that ConvNets can implicitly capture position information by utilizing image boundary effects, explicit position encoding has already been a core of full attention networks (e.g., Transformer [79]), and facilitated a variety of tasks (e.g., instance segmentation [93], tracking [47], video segmentation [11]). All these indicate that position encoding deserves more attention in the field of temporal correspondence learning.

Along this direction, LIIR explicitly injects a position encoding map  $P \in \mathbb{R}^{h' \times w' \times c'}$  into the feature encoder  $\phi$ :

$$I = \phi(I, P), \quad (6)$$

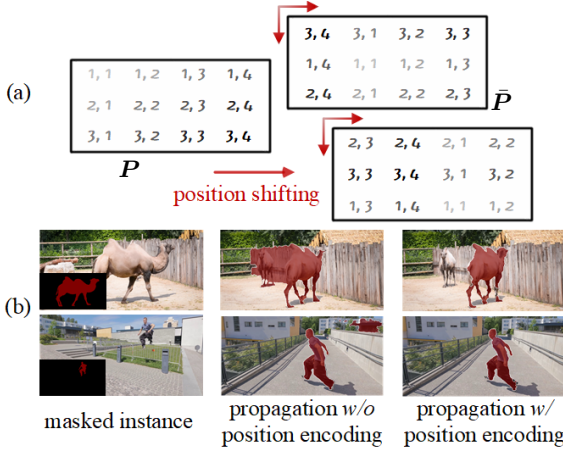


Figure 4. (a) Illustration of **position shifting**. (b) Effect of **position encoding**. See §3.2 for details.

where  $P$  is added with the output feature of the first conv layer of  $\phi$  and has the same size and dimension as the conv feature. We explore three **position encoding strategies**:

- **2D Sinusoidal Position Embedding (2DSPE)**:  $P$  is given with a family of pre-defined sinusoidal functions, without introducing new trainable parameters:

$$P(x, y, 2u) = \sin(x \cdot \varepsilon \frac{4u}{c^t}), P(x, y, 2u+1) = \cos(x \cdot \varepsilon \frac{4u}{c^t}),$$

$$P(x, y, 2v + \frac{c'}{2}) = \sin(y \cdot \varepsilon \frac{4v}{c^t}), P(x, y, 2v+1 + \frac{c'}{2}) = \cos(y \cdot \varepsilon \frac{4v}{c^t}),$$

where  $x \in [0, w']$ ,  $y \in [0, h']$  specify the horizontal and vertical positions,  $u, v \in [0, c'/4]$  specify the dimension, and  $\varepsilon = 10^{-4}$ . The horizontal (vertical) positions are encoded in the first (second) half of the dimensions. 2DSPE naturally handles resolutions that are unseen during training.

- **1D Absolute Position Embedding (1DAPE)**: 1DAPE is the most heavy-weight strategy: the whole  $P$  is a learnable parameter matrix without any constraint.
- **2D Absolute Position Embedding (2DAPE)**: As in [14], two separate parameter sets:  $X \in \mathbb{R}^{w' \times c'/2}$  and  $Y \in \mathbb{R}^{h' \times c'/2}$ , are learned for encoding the horizontal and vertical positions, respectively, and combined to generate  $P$ .

With our intra-inter video affinity (Eq. 4), exploiting position information in intra-video correspondence matching, *i.e.*,  $\{\exp(I_q(i) \cdot I_r(j))\}_j$ , addresses local continuity resides in videos. However, for inter-video pixel relevance computation, *i.e.*,  $\{\exp(I_q(i) \cdot I_n(k))\}_{n,k}$ , such position prior is undesirable, as it inspires the query pixel  $i$  in  $I_q$  to prefer matching these pixels with similar positions in other irrelevant videos  $\{I_n\}_n$ . To eliminate such position encoding induced bias from inter-video correspondence matching, we design a **position shifting** strategy (Fig. 4(a)). During training, for  $I_n$  from other videos, we circularly shift the position encoding vectors in  $P$  by a random step in horizontal and vertical axes, respectively. The reason why we adopt random shifting with circular boundary conditions, instead of random shuffling, is to preserve the spatial layout in the modulated position encoding map  $\bar{P}$ . Then  $\bar{P}$  and  $I_n$  are fed into  $\phi$  for

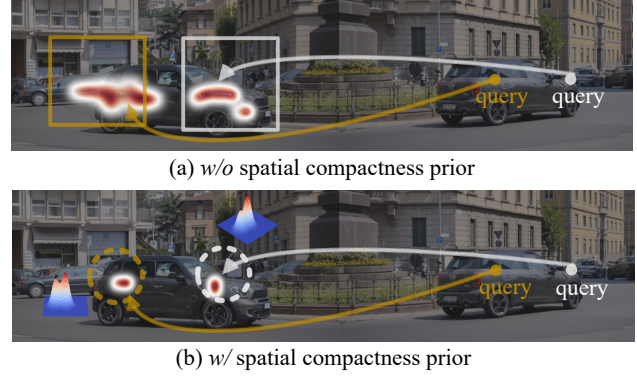


Figure 5. Illustration of **spatial compactness prior** (§3.2).

inter-video correspondence matching, and  $\bar{P}$  related gradients are abandoned if learnable 1DAPE or 2DAPE is used. Note that the standard position encoding  $P$  is applied for the query ( $I_q$ ) and reference ( $I_r$ ) frames and updated normally. Fig. 4(b) intuitively shows that merging position information into visual representation can enable robust correspondence matching even with confusing background and fast motion. In §4.4, we will quantitatively verify that 1DAPE is more favored and indeed boosts the performance.

**Spatial Compactness Prior.** As the visual world is continuous and smoothly-varying, it is reasonable to assume appearances in video data change smoothly both in spatial and temporal dimensions. For correspondence learning, the temporal coherence has been extensively studied, while the spatial continuity received far less attention. To reduce search region, some existing methods [39, 40, 84] restrict correspondence matching within a local window, considering spatial regularities in a simple way. To make a better use of the spatial continuity, we augment the original reconstruction objective with an additional prior, termed as *spatial compactness*. Such a prior poses constraints on the spatial distribution of associated pixels, leading to *sparse* and *coherent* solutions. Specifically, given the query ( $I_q$ ) and reference ( $I_r$ ) frames, we expect **i)** each query pixel  $i$  to be only matched with a small number of reference pixels, and **ii)** the matched reference pixels to be clustered. For a query pixel  $i$  and its matching ‘heatmap’:  $A_i = [A(i, j)]_{j \in [0, 1]^{h \times w}}$ , w.r.t.  $I_r$ , we assume  $A_i$  follows a mixture of  $M$  2D Gaussian distributions:

$$P(x, y) = \sum_{m=1}^M \omega_m \mathcal{N}(x, y | \mu_m, \Sigma_m), \quad (7)$$

where  $(x, y)$  specifies the coordinate of a pixel location. We set  $\{\mu_m = [\mu_{x,m}, \mu_{y,m}]\}_m$  as the coordinates of top- $M$  scores in  $A_i$ , and set  $M = 2$  to address sparse robust matching. Other parameters, *i.e.*,  $\{\Sigma_m = [\sigma_{x,m}^2, 0; 0, \sigma_{y,m}^2]\}_m$ ,  $\{\omega_m\}_m$ , can be estimated efficiently from  $A_i$  without incurring high computational cost. In this way, we can derive a ‘compact’ matching heatmap  $\tilde{A}_i \in [0, 1]^{h \times w}$  for each query pixel  $i$ , and eventually have  $\tilde{A} = [\tilde{A}_i]_{i \in [0, 1]^{hw \times hw}}$ . Such spatial compactness prior  $\tilde{A}$  is fully aware of **i)** and **ii)**, and used to regularize our representation learning:

Method	Backbone	Supervised	Dataset (size)	$\mathcal{J}$ & $\mathcal{F}$ (Mean) $\uparrow$	$\mathcal{J}$ (Mean) $\uparrow$	$\mathcal{J}$ (Recall) $\uparrow$	$\mathcal{F}$ (Mean) $\uparrow$	$\mathcal{F}$ (Recall) $\uparrow$
Colorization[84]	ResNet-18		Kinetics (-, 800 hours)	34.0	34.6	34.1	32.7	26.8
CorrFlow[40]	ResNet-18		OxUvA (-, 14 hours)	50.3	48.4	53.2	52.2	56.0
TimeCycle[90]	ResNet-50		VLOG (-, 344 hours)	48.7	46.4	50.0	50.0	48.0
UVC[42]	ResNet-50		C + Kinetics (30k, 800 hours)	60.9	59.3	68.8	62.7	70.9
MuG[52]	ResNet-18		OxUvA (-, 14 hours)	54.3	52.6	57.4	56.1	58.1
MAST[39]	ResNet-18		Youtube-VOS (-, 5.58 hours)	65.5	63.3	73.2	67.6	77.7
CRW[31]	ResNet-18		Kinetics (-, 800 hours)	68.3	65.5	78.6	71.0	82.9
ContrastCorr[87]	ResNet-18		C + TrackingNet (30k, 300 hours)	63.0	60.5	-	65.5	-
VFS[98]	ResNet-18		Kinetics (-, 800 hours)	66.7	64.0	-	69.4	-
JSTG[105]	ResNet-18		Kinetics (-, 800 hours)	68.7	65.8	77.7	71.6	84.3
CLTC[34] <sup>†</sup>	ResNet-18		Youtube-VOS (-, 5.58 hours)	70.3	67.9	78.2	72.6	83.7
DINO[6]	ViT-B/8		I (1.28M, -)	71.4	67.9	-	<b>74.9</b>	-
<b>LIIR</b>	ResNet-18		Youtube-VOS (-, 5.58 hours)	<b>72.1</b>	<b>69.7</b>	<b>81.4</b>	<b>74.5</b>	<b>85.9</b>
ResNet[24]	ResNet-18	✓	I (1.28M, -)	62.9	60.6	69.9	65.2	73.8
OSVOS[5]	VGG-16	✓	I+D (1.28M, 10k)	60.3	56.6	63.8	63.9	73.8
FEELVOS[81]	Xception-65	✓	I + C + D + Youtube-VOS (1.28M, 663k)	71.5	69.1	79.1	74.0	83.8
STM[61]	ResNet-50	✓	I + D + Youtube-VOS (1.28M, 164k)	81.8	79.2	88.7	84.3	91.8

<sup>†</sup>: using task-specific model weights and architectures. I: ImageNet[12]. C: COCO[46]. D: DAVIS<sub>17</sub>[68].

Table 1. **Quantitative results for video object segmentation** (§4.1) on DAVIS<sub>17</sub>[68] val. For size of datasets, we report (#raw images, length of raw videos) for self-supervised methods and (#image-level annotations, #pixel-level annotations) for supervised methods.




Figure 6. **Qualitative results for video object segmentation** (§4.1), on DAVIS<sub>17</sub>[68] val (left) and Youtube-VOS[99] val (right).

$$\mathcal{L}_{\text{com}} = \|\tilde{A} - A\|_2. \quad (8)$$

Note that  $\mathcal{L}_{\text{com}}$  is only applied for intra-video correspondence matching. Moreover, we replace  $A$  with  $\tilde{A}$  during inference, which eliminates outliers effectively. Two recent fully supervised video segmentation methods also explore the local continuity via single-Gaussian locality prior [72] or top- $k$  matching filtering [9], while both of them can be viewed as specific instances of our mixture-Gaussian based compactness prior, despite our different task settings. Fig. 5 shows that our spatial compactness prior helps build reliable correspondence by inspiring sparse and compact solutions. Related experiments can be found in §4.4.

### 3.3. Implementation Details

**Network Configuration.** For fair comparison, our feature encoder  $\phi$  is implemented as ResNet-18[24] as in [31, 84, 98]. Following [34, 39, 40],  $\times 2$  downsampling is only made in the third residual block. Thus  $\phi$  finally outputs 256 feature maps of 1/4 size of the input, *i.e.*,  $h = \frac{H}{4}$ ,  $w = \frac{W}{4}$ ,  $c = 256$ . The position embedding is added to feature after the first  $7 \times 7$  Conv-BN-ReLU layer, *i.e.*,  $h' = \frac{H}{2}$ ,  $w' = \frac{W}{2}$ ,  $c' = 64$ .

**Training:** LIIR is trained from scratch on two NVIDIA RTX-3090 GPUs and only uses the raw videos from Youtube-VOS [99]. Each training image is resized into  $256 \times 256$  and channel-wise dropout in Lab colorspace [39] is adopted as the information bottleneck. Adam optimizer is

used. At the initial 30 epochs, only intra-video reconstruction learning is adopted for warm-up, with a learning rate of  $10^{-3}$  and batch size of 32. Then we conduct inter-video reconstruction learning with spatial compactness based regularization at the next 5 epochs, with a learning rate of  $10^{-4}$  and batch size of 12. We online maintain a memory bank of 1,440 frames from different videos. For each query pixel, we sample 4 feature points from each stored frame, *i.e.*, a total of  $1,440 \times 4$  negative samples are used for the inter-video correspondence computation, and we employ the moving average strategy [23, 76, 97] for parameter updating.

**Testing:** Once LIIR finishes training, there is no fine-tuning when applied to downstream tasks. Note that we utilize the compactness prior enhanced inter-frame affinity  $\tilde{A}$  for mask propagation. As in [39, 61], we take multiple frames as reference for the full use of temporal context: at time step  $t$ , previous frames  $I_0$ ,  $I_5$ ,  $I_{t-5}$ ,  $I_{t-3}$ , and  $I_{t-1}$  (if applicable) are referred for current-frame mask propagation.

## 4. Experiment

We evaluate the learned representation on diverse video label propagation tasks, *i.e.*, video object segmentation (§4.1), body part propagation (§4.2), and pose keypoint tracking (§4.3). As in conventions [31, 84, 98], all these tasks are to propagate the first frame annotation to the whole video sequence, and we use our model to compute inter-frame dense




Figure 7. Qualitative results for part propagation (§4.2) and pose tracking (§4.3), on VIP [107] val (left) and JHMDB [35] val (right).

Methods	Sup.	Overall	Seen		Unseen	
			$\mathcal{J} \uparrow$	$\mathcal{F} \uparrow$	$\mathcal{J} \uparrow$	$\mathcal{F} \uparrow$
Colorization [84]		38.9	43.1	38.6	36.6	37.4
CorrFlow [40]		46.6	50.6	46.6	43.8	45.6
MAST [39]		64.2	63.9	64.9	60.3	67.7
CLTC [34] <sup>†</sup>		67.3	66.2	67.9	63.2	71.7
<b>LIIR</b>		<b>69.3</b>	<b>67.9</b>	<b>69.7</b>	<b>65.7</b>	<b>73.8</b>
OSVOS [5]	✓	58.8	59.8	60.5	54.2	60.7
PreMVOS [53]	✓	66.9	71.4	75.9	56.5	63.7
STM [61]	✓	79.4	79.7	84.2	72.8	80.9

<sup>†</sup>: using task-specific model weights and architectures.

Table 2. Quantitative results for video object segmentation (§4.1) on Youtube-VOS [99] val.

correspondences. In §4.4, we conduct a set of ablative studies to examine the efficacy of our essential model designs.

#### 4.1. Results for Video Object Segmentation

**Dataset.** We first test our method on val sets of two popular video object segmentation datasets, *i.e.*, DAVIS<sub>17</sub> [68] and YouTube-VOS [99]. There are 30 and 474 videos in DAVIS<sub>17</sub> and YouTube-VOS val sets, respectively.

**Evaluation Metric.** Following the official protocol [68], we use the region similarity ( $\mathcal{J}$ ) and contour accuracy ( $\mathcal{F}$ ) as the evaluation metrics. Note that the scores on YouTube-VOS are respectively reported for *seen* and *unseen* categories, obtained from the official evaluation server.

**Performance on DAVIS<sub>17</sub>:** As illustrated in Table 1, our LIIR consistently outperforms all existing self-supervised methods across all the evaluation metrics. For example, it surpasses current best-performing self-supervised method, *i.e.*, CLTC [34], in terms of mean  $\mathcal{J}\&\mathcal{F}$  (72.1 vs. 70.3). In addition, even without using *any* manual annotations for training, LIIR achieves very competitive segmentation performance in comparison with some famous supervised models [5, 81] trained with massive pixel-wise annotations.

**Performance on YouTube-VOS val.** Table 2 reports performance comparison of LIIR against four self-supervised competitors on YouTube-VOS val. It can be observed that LIIR sets new state-of-the-art. In particular, LIIR yields an overall score of 69.3%, surpassing the second-best (*i.e.*, CLTC [34]) and third-best (*i.e.*, MAST [39]) approaches by 2.0% and 5.1%, respectively. Further, LIIR even outperforms some famous supervised methods (*i.e.*, OSVOS [5] and PreMVOS [53]), especially for the *unseen* categories, clearly demonstrating its remarkable generalization ability.

Methods	Sup.	VIP		JHMDB	
		mIoU $\uparrow$	AP $\uparrow$	PCK@0.1 $\uparrow$	PCK@0.2 $\uparrow$
TimeCycle [90]		28.9	15.6	57.3	78.1
UVC [42]		34.1	17.7	58.6	79.6
CRW [31]		38.6	-	59.3	80.3
ContrastCorr [87]		37.4	21.6	61.1	80.8
VFS [98]		39.9	-	60.5	79.5
CLTC [34] <sup>†</sup>		37.8	19.1	60.5	82.3
JSTG [105]		40.2	-	<b>61.4</b>	<b>85.3</b>
<b>LIIR</b>		<b>41.2</b>	<b>22.1</b>	60.7	81.5
ResNet [24]	✓	31.9	12.6	53.8	74.6
ATEN [73]	✓	37.9	24.1	-	-
TSN [107]	✓	-	-	68.7	92.1

<sup>†</sup>: using task-specific model weights and architectures.

Table 3. Quantitative results for part propagation (§4.2) and pose tracking (§4.3), on VIP [107] val and JHMDB [35] val.

**Qualitative Results.** Fig. 6 depicts visual results on representative videos in the datasets. As seen, LIIR is able to establish accurate correspondences under various challenging scenarios, *e.g.*, scale changes, small objects and occlusions.

#### 4.2. Results for Body Part Propagation

**Dataset.** We next evaluate our model performance for body part propagation. Experiments are conducted on VIP val [107], which contains 50 videos with annotations of 19 human semantic part categories (*e.g.*, hair, face, dress).

**Evaluation Metric.** As suggested by VIP [107], we adopt mean intersection-over-union (mIoU) and mean Average Precision (mAP) metrics for evaluation of semantic-level and instance-level parsing, respectively.

**Performance.** As shown in Table 3, LIIR achieves state-of-the-art performance on both semantic-level and instance-level parsing. This indicates that LIIR can generate strong representation which models both cross-instance discrimination and intra-instance invariance well. Fig. 7 depicts some visualization results on two representative videos. LIIR achieves temporally stable results and shows robustness to typical challenges (*e.g.*, pose variations, occlusions).

#### 4.3. Results for Pose Keypoint Tracking

**Dataset.** We then examine the model performance on human keypoint tracking, using JHMDB [35] val. JHMDB val has 268 videos. For each person, a total of 15 body joints, *e.g.*, torso, head, shoulder, elbow, are annotated.

**Evaluation Metric.** We use *probability of correct keypoint*

#	#Negative Samples	DAVIS		VIP		Position Encoding	DAVIS		VIP		Spatial Compactness		DAVIS		VIP	
		$\mathcal{J}\&\mathcal{F}_m \uparrow$	mIoU $\uparrow$	$\mathcal{J}\&\mathcal{F}_m \uparrow$	mIoU $\uparrow$		$\mathcal{J}\&\mathcal{F}_m \uparrow$	mIoU $\uparrow$	training	inference	$\mathcal{J}\&\mathcal{F}_m \uparrow$	mIoU $\uparrow$	training	inference	$\mathcal{J}\&\mathcal{F}_m \uparrow$	mIoU $\uparrow$
1	0	69.2	38.4	<i>w/o</i> PE	70.9	40.3	N/A	71.3	40.6					69.8	39.6	
2	480	70.9 (+1.7)	39.8 (+1.4)	2DSPE	70.6 (-0.3)	40.2 (-0.1)	shuffling	71.8 (+0.5)	41.0 (+0.4)	✓				71.5 (+1.7)	40.8 (+1.2)	
3	960	71.7 (+2.5)	41.0 (+2.6)	1DAPE	72.1 (+1.2)	41.2 (+0.9)	shifting	72.1 (+0.8)	41.2 (+0.6)		✓			70.8 (+1.0)	40.3 (+0.7)	
4	1,440	72.1 (+2.9)	41.2 (+2.8)	2DAPE	71.9 (+1.0)	41.1 (+0.8)				✓	✓			72.1 (+2.3)	41.2 (+1.6)	

(a) Inter-and Intra-Video Recons.

(b) Position Encoding

(c) Position Shifting

(d) Spatial Compactness Prior

Table 5. A set of ablation studies on DAVIS<sub>17</sub> [68] val and VIP [107] val. See §4.4 for details.

#	Inter-and Intra-Video Recons.	Position Encoding	Spatial Compactness	DAVIS $\mathcal{J}\&\mathcal{F}_m \uparrow$	VIP mIoU $\uparrow$
1				65.3	35.2
2	✓			68.7 (+3.4)	38.4 (+3.2)
3		✓		66.9 (+1.6)	37.0 (+1.8)
4			✓	68.4 (+3.1)	37.2 (+2.0)
5	✓	✓	✓	72.1 (+6.8)	41.2 (+6.0)

Table 4. Detailed analysis of essential components of LIIR on DAVIS<sub>17</sub> [68] val and VIP [107] val. See §4.4 for details.

(PCK) [101] to measure the accuracy between each tracking result and corresponding ground-truth with a threshold  $\tau$ .

**Performance.** Table 3 shows that LIIR exhibits compelling overall performance. Note that CLTC [34] uses different checkpoints and model architectures for different tasks and datasets, while we only use a single model for evaluation. The visual results in Fig. 7 also demonstrate the strong capability of LIIR in establishing precise correspondence.

#### 4.4. Diagnostic Experiment

For further detailed analysis, we conduct a series of ablation studies on DAVIS<sub>17</sub> [68] val and VIP [107] val sets. **Key Component Analysis.** We first examine the efficacy of essential components of LIIR, *i.e.*, inter-and intra-video reconstruction, position encoding, and spatial compactness. The results are summarized in Table 4, where position encoding is implemented as 1DAPE, and compactness prior is used during both training and inference stages. When separately comparing row #2 - #4 with the baseline (MAST [39]) in row #1, we can observe that each individual module indeed boosts the performance. For example, on DAVIS<sub>17</sub> val, inter-and intra-video reconstruction, position encoding, and spatial compactness prior respectively bring 3.4%, 1.6%, and 3.1%  $\mathcal{J}\&\mathcal{F}$  gains. This verifies our core insight that these three elements are crucial for correspondence learning. Finally, in row #5, we combine all the three components together – LIIR, and obtain the best performance. This suggests that these modules are complementary to each other, and confirms the effectiveness of our whole design.

**Inter-and Intra-Video Reconstruction.** We next study the impact of increasing the number of negative samples, *i.e.*, frames from other irrelevant videos used for inter-video correspondence computation (Eq. 4). In Table 5a, row #1 gives scores of learning without considering inter-video correspondence. In this case, the results are unsatisfactory. When more negative samples are involved (*i.e.*, 0→1,440), better performance can be achieved (*i.e.*, 69.2→72.1 on DAVIS<sub>17</sub>

val, 38.4→41.2 on VIP val). Finally we use 1,440 negative samples for inter-video reconstruction based learning, which is the maximum number allowed by our GPU.

**Position Encoding.** To determine the effect of our position encoding module, we then report the performance with different encoding strategies in Table 5b. As seen, the non-learnable strategy, 2DSPE, hinders the performance, while the learnable alternatives, *i.e.*, 1DAPE and 2DAPE, lead to better results. Compared with 2DAPE, 1DAPE is more favored, probably due to its high flexibility and capacity.

**Position Shifting.** We further study the influence of our position shifting strategy on performance. As shown in Table 5c, we consider two alternatives, *i.e.*, ‘NAN’ and ‘position shuffling’. ‘NAN’ refers to using the normal position encoding map  $\mathcal{P}$  without any modulation during inter-video correspondence matching. Compared with ‘position shifting’, ‘NAN’ suffers from performance degradation (*i.e.*, 72.1→71.3 on DAVIS<sub>17</sub> val, 41.2→40.6 on VIP val), showing the negative effect of the position-induced bias. The other baseline, ‘position shuffling’, *i.e.*, randomly shuffling the position encoding map for inter-video affinity computation, though better than ‘NAN’, is still worse than ‘position shifting’. This is because it destroys the spatial layouts.

**Spatial Compactness Prior.** The spatial compactness prior (Eq. 7) is used to regularize intra-video correspondence matching during both training and inference stages. In Table 5d, we quantitatively identify the performance contribution of our spatial compactness prior in different stages.

## 5. Conclusion

We presented a self-supervised temporal correspondence learning approach, LIIR, that makes contributions in three aspects. First, going beyond the popular intra-video analysis based learning scheme, we further enforce separation between intra- and inter-video pixel associations, enhancing instance-level feature discrimination. Second, with a clever position shifting strategy, we bring the advantages of position encoding into full play, while avoiding its undesirable impact at the same time. Third, a spatial compactness prior is introduced to regularize representation learning and improve correspondence inference. The effectiveness was thoroughly validated over various label propagation tasks.

**Acknowledgements** This work was supported in part by the Beijing Natural Science Foundation under Grant L191004, CCF-Baidu Open Fund, and ARC DECRA DE220101390.



## References

- [1] Pulkit Agrawal, Joao Carreira, and Jitendra Malik. Learning to see by moving. In *ICCV*, 2015. 3
- [2] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *IJCV*, 92(1):1–31, 2011. 1
- [3] Linchao Bao, Baoyuan Wu, and Wei Liu. Cnn in mrf: Video object segmentation via inference in a cnn-based higher-order spatio-temporal mrf. In *CVPR*, 2018. 1
- [4] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *ECCV*, 2016. 2
- [5] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *CVPR*, 2017. 3, 6, 7
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021. 6
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 3
- [8] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 3
- [9] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion. In *CVPR*, 2021. 6
- [10] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *NeurIPS*, 2021. 1
- [11] Hai Ci, Chunyu Wang, and Yizhou Wang. Video object segmentation by learning location-sensitive embeddings. In *ECCV*, 2018. 4
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6
- [13] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 3
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5
- [15] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015. 2
- [16] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal cycle-consistency learning. In *CVPR*, 2019. 3
- [17] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *CVPR*, 2017. 3
- [18] Chuang Gan, Boqing Gong, Kun Liu, Hao Su, and Leonidas J Guibas. Geometry guided convolutional neural networks for self-supervised video representation learning. In *CVPR*, 2018. 3
- [19] Andreas Geiger, Julius Ziegler, and Christoph Stiller. Stereoscan: Dense 3d reconstruction in real-time. In *2011 IEEE Intelligent Vehicles Symposium*, 2011. 1
- [20] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 3
- [21] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *ECCV*, 2020. 3
- [22] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *NeurIPS*, 2020. 3
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 3, 6
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3, 6, 7
- [25] Robert F Hess and Ian E Holliday. The coding of spatial position by the human visual system: effects of spatial scale and contrast. *Vision Research*, 32(6):1085–1097, 1992. 2, 4
- [26] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR*, 2019. 3
- [27] Yuan-Ting Hu, Jia-Bin Huang, and Alexander G Schwing. Videomatch: Matching based video object segmentation. In *ECCV*, 2018. 1, 2
- [28] Xuhua Huang, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Fast video object segmentation with temporal aggregation network and dynamic template matching. In *CVPR*, 2020. 1
- [29] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 2
- [30] Md Amirul Islam, Sen Jia, and Neil DB Bruce. How much position information do convolutional neural networks encode? In *ICLR*, 2020. 4
- [31] Allan Jabri, Andrew Owens, and Alexei A Efros. Space-time correspondence as a contrastive random walk. In *NeurIPS*, 2020. 1, 2, 6, 7
- [32] Varun Jampani, Raghudeep Gadde, and Peter V. Gehler. Video propagation networks. In *CVPR*, 2017. 3
- [33] Joel Janai, Fatma Guney, Anurag Ranjan, Michael Black, and Andreas Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *ECCV*, 2018. 2

- [34] Sangryul Jeon, Dongbo Min, Seungryong Kim, and Kwanghoon Sohn. Mining better samples for contrastive learning of temporal correspondence. In *CVPR*, 2021. 1, 6, 7, 8
- [35] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *ICCV*, 2013. 7
- [36] Joakim Johnander, Martin Danelljan, Emil Brissman, Fahad Shahbaz Khan, and Michael Felsberg. A generative appearance model for end-to-end video object segmentation. In *CVPR*, 2019. 1
- [37] Osman Semih Kayhan and Jan C van Gemert. On translation invariance in cnns: Convolutional layers can exploit absolute spatial location. In *CVPR*, 2020. 4
- [38] Dahun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *AAAI*, 2019. 3
- [39] Zihang Lai, Erika Lu, and Weidi Xie. Mast: A memory-augmented self-supervised tracker. In *CVPR*, 2020. 1, 2, 3, 5, 6, 7, 8
- [40] Zihang Lai and Weidi Xie. Self-supervised learning for video correspondence flow. In *BMVC*, 2019. 1, 2, 3, 5, 6, 7
- [41] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *ECCV*, 2016. 3
- [42] Xueting Li, Sifei Liu, Shalini De Mello, Xiaolong Wang, Jan Kautz, and Ming-Hsuan Yang. Joint-task self-supervised learning for temporal correspondence. In *NeurIPS*, 2019. 1, 2, 6, 7
- [43] Xiaoxiao Li and Chen Change Loy. Video object segmentation with joint re-identification and attention-aware mask propagation. In *ECCV*, 2018. 1
- [44] Yongqing Liang, Xin Li, Navid Jafari, and Qin Chen. Video object segmentation with adaptive feature bank and uncertain-region refinement. In *NeurIPS*, 2020. 1
- [45] Huaijia Lin, Xiaojuan Qi, and Jiaya Jia. Agss-vos: Attention guided single-shot video object segmentation. In *CVPR*, 2019. 1
- [46] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 3, 6
- [47] Yuan Liu, Ruoteng Li, Yu Cheng, Robby T Tan, and Xiubao Sui. Object tracking using spatio-temporal networks for future prediction location. In *ECCV*, 2020. 4
- [48] Margaret S Livingstone and David H Hubel. Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience*, 7(11):3416–3468, 1987. 4
- [49] William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. *arXiv preprint arXiv:1605.08104*, 2016. 3
- [50] Xiankai Lu, Wenguan Wang, Martin Danelljan, Tianfei Zhou, Jianbing Shen, and Luc Van Gool. Video object segmentation with episodic graph memory networks. In *ECCV*, 2020. 1, 3
- [51] Xiankai Lu, Wenguan Wang, Chao Ma, Jianbing Shen, Ling Shao, and Fatih Porikli. See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In *CVPR*, 2019. 3
- [52] Xiankai Lu, Wenguan Wang, Jianbing Shen, Yu-Wing Tai, David J Crandall, and Steven CH Hoi. Learning video object segmentation from unlabeled videos. In *CVPR*, 2020. 1, 2, 3, 6
- [53] Jonathon Luiten, Paul Voigtlaender, and Bastian Leibe. Premevos: Proposal-generation, refinement and merging for video object segmentation. In *ACCV*, 2018. 7
- [54] K-K Maninis, Sergi Caelles, Yuhua Chen, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. Video object segmentation without temporal information. *IEEE TPAMI*, 2018. 1
- [55] Tim Meinhardt and Laura Leal-Taixé. Make one-shot video object segmentation efficient again. *NeurIPS*, 2020. 1
- [56] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *AAAI*, 2018. 2
- [57] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, 2020. 3
- [58] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016. 3
- [59] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 3
- [60] Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, and Seon Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *CVPR*, 2018. 1
- [61] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *ICCV*, 2019. 1, 3, 6, 7
- [62] Lauri Oksama and Jukka Hyönä. Position tracking and identity tracking are separate systems: Evidence from eye movements. *Cognition*, 146:393–409, 2016. 2, 4
- [63] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [64] Anestis Papazoglou and Vittorio Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, 2013. 3
- [65] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *CVPR*, 2017. 3
- [66] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 3
- [67] Federico Perazzi, Anna Khoreva, Rodrigo Benenson, Bernt Schiele, and Alexander Sorkine-Hornung. Learning video object segmentation from static images. In *CVPR*, 2017. 3
- [68] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 3, 6, 7, 8

- [69] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *CVPR*, 2021. 3
- [70] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015. 3
- [71] Scott Satkin and Martial Hebert. 3dnn: Viewpoint invariant 3d geometry matching for scene understanding. In *ICCV*, 2013. 1
- [72] Hongje Seong, Junhyuk Hyun, and Euntai Kim. Kernelized memory network for video object segmentation. In *ECCV*, 2020. 6
- [73] Jie Song, Limin Wang, Luc Van Gool, and Otmar Hilliges. Thin-slicing network: A deep structured model for pose estimation in videos. In *CVPR*, 2017. 7
- [74] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *ICCV*, 2019. 3
- [75] Mingjie Sun, Jimin Xiao, Eng Gee Lim, Bingfeng Zhang, and Yao Zhao. Fast template matching and update for video object tracking and segmentation. In *CVPR*, 2020. 1
- [76] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 6
- [77] David Tsai, Matthew Flagg, and James M Rehg. Motion coherent tracking using multi-label MRF optimization. In *BMVC*, 2010. 3
- [78] Hsiao-Yu Fish Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NeurIPS*, 2017. 3
- [79] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4
- [80] Carles Ventura, Miriam Bellver, Andreu Girbau, Amaia Salvador, Ferran Marques, and Xavier Giro-i Nieto. Rvos: End-to-end recurrent network for video object segmentation. In *CVPR*, 2019. 3
- [81] Paul Voigtlaender, Yuning Chai, Florian Schroff, Hartwig Adam, Bastian Leibe, and Liang-Chieh Chen. Feelvos: Fast end-to-end embedding learning for video object segmentation. In *CVPR*, 2019. 6, 7
- [82] Paul Voigtlaender and Bastian Leibe. Online adaptation of convolutional neural networks for video object segmentation. 2017. 1, 3
- [83] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *CVPR*, 2016. 3
- [84] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by coloring videos. In *ECCV*, 2018. 1, 2, 3, 5, 6, 7
- [85] Jiangliu Wang, Jianbo Jiao, Linchao Bao, Shengfeng He, Yunhui Liu, and Wei Liu. Self-supervised spatio-temporal representation learning for videos by predicting motion and appearance statistics. In *CVPR*, 2019. 3
- [86] Ning Wang, Yibing Song, Chao Ma, Wengang Zhou, Wei Liu, and Houqiang Li. Unsupervised deep tracking. In *CVPR*, 2019. 1, 2
- [87] Ning Wang, Wengang Zhou, and Houqiang Li. Contrastive transformation for self-supervised correspondence learning. In *AAAI*, 2021. 1, 2, 4, 6, 7
- [88] Wenguan Wang, Jianbing Shen, Fatih Porikli, and Ruigang Yang. Semi-supervised video object segmentation with super-trajectories. *IEEE TPAMI*, 41(4):985–998, 2018. 3
- [89] Wenguan Wang, Jianbing Shen, Ruigang Yang, and Fatih Porikli. Saliency-aware video object segmentation. 40(1):20–33, 2017. 3
- [90] Wenguan Wang, Hongmei Song, Shuyang Zhao, Jianbing Shen, Sanyuan Zhao, Steven CH Hoi, and Haibin Ling. Learning unsupervised video object segmentation through visual attention. In *CVPR*, 2019. 1, 2, 3, 6, 7
- [91] Wenguan Wang, Tianfei Zhou, Fatih Porikli, David Crandall, and Luc Van Gool. A survey on deep learning technique for video segmentation. *arXiv preprint arXiv:2107.01153*, 2021. 3
- [92] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015. 3
- [93] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. In *NeurIPS*, 2020. 4
- [94] Ziqin Wang, Jun Xu, Li Liu, Fan Zhu, and Ling Shao. Ranet: Ranking attention network for fast video object segmentation. In *CVPR*, 2019. 1
- [95] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *CVPR*, 2018. 3
- [96] Haiping Wu and Xiaolong Wang. Contrastive learning of image representations with cross-video cycle-consistency. In *ICCV*, 2021. 3
- [97] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 6
- [98] Jiarui Xu and Xiaolong Wang. Rethinking self-supervised correspondence learning: A video frame-level similarity perspective. In *ICCV*, 2021. 1, 2, 6, 7
- [99] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *ECCV*, 2018. 3, 6, 7
- [100] Shuangjie Xu, Daizong Liu, Linchao Bao, Wei Liu, and Pan Zhou. Mhp-vos: Multiple hypotheses propagation for video object segmentation. In *CVPR*, 2019. 1
- [101] Yi Yang and Deva Ramanan. Articulated human detection with flexible mixtures of parts. *IEEE TPAMI*, 2012. 8
- [102] Yi Yang, Yueting Zhuang, and Yunhe Pan. Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, 22(12):1551–1558, 2021. 3
- [103] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, 2018. 2

- [104] Jae Shin Yoon, François Rameau, Jun-Sik Kim, Seokju Lee, Seunghak Shin, and In So Kweon. Pixel-level matching for video object segmentation using convolutional neural networks. In *ICCV*, 2017. 3
- [105] Zixu Zhao, Yueming Jin, and Pheng-Ann Heng. Modelling neighbor relation in joint space-time graph for video correspondence learning. In *ICCV*, 2021. 6, 7
- [106] Jilai Zheng, Chao Ma, Houwen Peng, and Xiaokang Yang. Learning to track objects from unlabeled videos. In *ICCV*, 2021. 2
- [107] Qixian Zhou, Xiaodan Liang, Ke Gong, and Liang Lin. Adaptive temporal encoding network for video instance-level human parsing. In *ACMMM*, 2018. 7, 8
- [108] Tianfei Zhou, Jianwu Li, Xueyi Li, and Ling Shao. Target-aware object discovery and association for unsupervised video multi-object segmentation. In *CVPR*, 2021. 3
- [109] Tianfei Zhou, Shunzhou Wang, Yi Zhou, Yazhou Yao, Jianwu Li, and Ling Shao. Motion-attentive transition for zero-shot video object segmentation. In *AAAI*, 2020. 3
- [110] Linchao Zhu and Yi Yang. Actbert: Learning global-local video-text representations. In *CVPR*, 2020. 3
- [111] Yuliang Zou, Zelun Luo, and Jia-Bin Huang. Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In *ECCV*, 2018. 2