

A Simple and Powerful Global Optimization for Unsupervised Video Object Segmentation

Georgy Ponimatkin¹ Nermin Samet¹ Yang Xiao¹ Yuming Du¹ Renaud Marlet^{1,2} Vincent Lepetit¹

¹LIGM, Ecole des Ponts, Univ Gustave Eiffel, CNRS, Marne-la-Vallée, France ²Valeo.ai, Paris, France

georgy.ponimatkin@enpc.fr

Code and supplementary material: <https://ponimatkin.github.io/ssl-vos>

Abstract

We propose a simple, yet powerful approach for unsupervised object segmentation in videos. We introduce an objective function whose minimum represents the mask of the main salient object over the input sequence. It only relies on independent image features and optical flows, which can be obtained using off-the-shelf self-supervised methods. It scales with the length of the sequence with no need for superpixels or sparsification, and it generalizes to different datasets without any specific training. This objective function can actually be derived from a form of spectral clustering applied to the entire video. Our method achieves on-par performance with the state of the art on standard benchmarks (DAVIS2016, SegTrack-v2, FBMS59), while being conceptually and practically much simpler.

1. Introduction

While the two research communities working on unsupervised video object segmentation [46, 1, 86, 45] and on object discovery [71, 82, 37, 62] often remain separated, they share the goal of segmenting objects in visual data without depending on manual labels for these objects. This ability is essential to autonomous systems for evolving and interacting in open world. It is also a fundamental problem for Computer Vision as humans have the ability to learn about new objects without guidance.

Object appearance and motion are important cues to achieve this task. However, many challenges remain: Objects can share similar appearance with the background, their visual appearance may not be uniform, and different parts can move in different directions. As a result, many methods still rely on some sort of supervision, at least for learning to extract visual features.

We present here a simple approach to unsupervised video object segmentation. It leverages recent progress on unsupervised learning in a simple but powerful, novel optimization scheme over the objects' masks. We start from




Figure 1: **(a)** Segmentation obtained by our spectral clustering formulation on the self-supervised image features from DINO [7] in a single frame. **(b)** Segmentation obtained using the same self-supervised image features *and* optical flow from ARFlow [42], but still from a single frame. **(c)** Final segmentation obtained with our complete method, after optimization on the full frame sequence, using the same features and optical flow estimated by ARFlow.

self-supervised features such as DINO [7], MoCo-v3 [11], SWAV [6] or Barlow Twins [89] as the appearance cue, and the optical flow from methods such as RAFT [68] or ARFlow [42]. DINO, MoCo-v3, SWAV, Barlow Twins and ARFlow methods are “self-supervised” or “unsupervised”, in the sense that they do not use any manual annotations, making our method also entirely unsupervised.¹

Used alone, the DINO features can already provide a surprisingly good segmentation of the objects, but still below the state of the art, in particular for videos. Yet, with our optimization scheme, we can reach a performance that is on par or even better than much more sophisticated methods.

Our optimization starts from an initial estimate for the object's mask in each frame of the input sequence. It then optimizes the masks over the whole video sequence. We show that our optimization function can be derived from spectral clustering over the video sequence. Yet, spectral clustering is difficult to make tractable for long sequences as it requires computing one of the eigenvectors of a huge matrix [60, 48]. To obtain a tractable problem, previous methods based on spectral clustering rely on superpixels [23] or graph sparsification [35]. However, superpixels may in-

¹The terminology is still fluctuating, as the DINO method is called “self-supervised” in the original paper, while ARFlow is called “unsupervised”. Throughout the paper, we will use self-supervised and unsupervised in the sense of “not using manual annotations”.




Figure 2: **Generalization to new datasets.** Relying on spectral clustering rather than learning generalizes better to new datasets. (a) The learned method of [82] performs poorly on a new sequence. (b) Spectral clustering on the *same* optical flow performs significantly better, even without optimization. (c) Result of our complete method.

introduce artefacts at object boundaries, and both superpixels and sparsification require careful parameter tuning. Besides, even using fast methods to extract relevant eigenvectors, such as Power Iteration Clustering (PIC) [40], the complexity of these formulations is *a priori* quadratic in the length of the sequence (depending however on sparsity). In contrast, our method only needs eigenvectors computed for each frame independently, based on image features and optical flow for the frame, thus essentially scaling linearly with the number of frames. In short, our approach yields a tractable approximation of spectral clustering over the entire sequence. Concretely, on videos from the DAVIS2016 dataset [55], our method is on average $\sim 170\times$ faster than our full spectral clustering counterpart, *i.e.*, TokenCut [77].

Importantly, for practical applications, our method can be applied to different datasets without retraining. This is in contrast with learning-based segmentation approaches. As illustrated in Fig. 2, they do not generalize well to new data.

To summarize, our contribution is the derivation of a novel objective function from spectral clustering, which results in a simple and efficient optimization method for object discovery and segmentation in videos. Moreover, our method is arguably the first that relies purely on self-supervised features, without the need of any manual annotation. Still, it outperforms previous methods on several standard challenging datasets, including some methods that only rely on supervised training. Last, it can directly be applied to different datasets, without any training or tuning.

2. Related Work

We discuss here works on object discovery and unsupervised video object segmentation, as both are closely related to our work. We also discuss works on unsupervised feature learning, which we rely on for image feature extraction.

2.1. Object Discovery

Object discovery aims to localize the objects in images or videos. Some works rely on a collection of images containing objects of the same class and localize these objects using

clustering [20], image matching [58, 13], topic discovery [59, 63] or optimization for selecting region proposals [71]. One limitation of these works is that such a collection could be difficult to obtain, and its quality would heavily impact performance. Du *et al.* [18] proposes to discover and segment objects from unseen classes based on instance masks predicted by models trained only on seen classes, while our method does not need any supervision.

Recently, some works [44, 5, 24] adopt a bottom-up approach and segment the object in the images by exploiting the similarity among pixel colors or patch features. These works have only demonstrated their performance on synthetic images and can be easily affected by image texture or colors. In contrast, our approach aims to discover objects using unlabeled in-the-wild videos without any supervision.

[82, 37] and [26] can also be considered as object discovery methods, as their goal is to segment the primary objects given some video sequences. [82, 37] propose to use an attention-based network to compute the trajectory embeddings from optical flow only. One major limitation of [82, 37] is that optical flow is not sufficient to segment static objects. Besides, even though [82, 37] do not require any labeled data, they still have to train their network with different strategies on different datasets, which makes it difficult to generalize, as show in Figure 2. As we will show in Section 4, our method does not need any training and can achieve good results on all the benchmarks.

IKE [26] proposes an iterative refinement approach: In the first stage, the videos are first fed into a graph module for mask initialization. Then, in the second stage, the initialized masks are used to train a segmentation network, whose predictions are used as initialization for the graph module in the first stage. In the first stage, optical flow is used to generate long-term space-time trajectories. Optical flow can be very noisy and thus it may be difficult to obtain long trajectories reliably in general. We will show in the Experiments section that our method can achieve better results compared to [26] while at the same time being simple. Some methods [82, 37, 85] are also presented as self-supervised or unsupervised while they actually rely on supervised optical flow methods such as RAFT [68]. In comparison, our method is able to achieve state-of-the-art performance while only using self-supervised optical flow learned from videos.

2.2. Unsupervised Video Object Segmentation

Given a video sequence, Unsupervised Video Object Segmentation (UVOS) aims to consistently segment and track the most salient objects in the video without any human intervention. The approaches for UVOS can be divided into two categories depending on whether they need labeled data or not. Methods such as [46, 86, 45, 1, 76, 52, 39, 72, 47, 54, 74, 65] use either labeled images or labeled videos

to train their networks for video segmentation. In addition, some works [91, 80, 84, 57, 90, 83, 33, 70, 12, 61, 32, 70] also exploit low-level cues such as object boundaries to get better mask predictions. One main limitation of these methods is that they rely heavily on their large-scale well-annotated training data, which can be hard to get in practice. In contrast, our proposed method is entirely self-supervised and does not need any labeled data.

There are also several works [21, 14, 53, 2, 4, 56, 34, 50, 64, 85, 75, 82, 87] that can segment the primary objects in the video without any labeled data. One major limitation of all these works is that they all consider that the pixels on the objects in the video share similar motion patterns, thus, their methods can fail when the objects are static or move with the same speed as the background, while our proposed method combines the appearance cues and the motion cues and can largely alleviate this issue.

Currently, DyStaB [84] is the best performing unsupervised method for video object segmentation. The method consists of three parts: A static model and a dynamic model (both based on DeepLab backend [8]), together with an inpainter network based on [85]. The three networks are jointly trained via an adversarial loss in an iterative fashion to obtain final segmentation results and a significant post-processing via CRF is utilized. In comparison, our approach does not require training—besides for the image features and optical flow that can be obtained out-of-the-shelf—while an adversarial loss can be difficult to train. More importantly, our approach is much simpler, while achieving a similar performance to DyStaB.

2.3. Self-Supervised Learning

Recent self-supervised approaches [79, 9, 17, 27, 10, 25, 7] propose to train a feature extraction network using an instance classification pipeline, which treats each image as a single class and trains the network to distinguish images cropped from a large image collection without any manual annotation. In particular, DINO [7] introduced an approach which makes a network trained in a self-supervised manner learn class-specific features.

Motivated by this ability, discovering objects in images using self-supervised features recently gained attention [77, 49, 62]. LOST [62] utilizes self-supervised features within a seed selection and expansion strategy and localizes the main object given an image. Closely related to LOST, TokenCut [77] and Deep Spectral Methods (DSM) [49] propose graph-based methods that use self-supervised transformer features to discover and segment salient objects. TokenCut [77] builds a graph where visual tokens are nodes and similarity scores between tokens are edges of a weighted graph. They formulate the segmentation problem as a normalized graph cut and solve it using spectral clustering with eigendecomposition. Similar to

TokenCut [77], inspired from traditional spectral segmentation methods, DSM [49] first constructs a Laplacian matrix which is a combination of color information and self-supervised transformer features. Next, the image is decomposed using the eigenvectors of the Laplacian matrix.

Our method is similar to these methods as we also aim to detect and segment salient objects based on a graph formulation. There are two key differences: (i) We aim to discover objects in videos rather than still frames and in order to ensure the temporal consistency, we extend the affinity matrix with optical flow to establish inter-frame connectivity; (ii) We introduce a novel method for optimizing on the entire video efficiently. Moreover, we show that using power iteration clustering [40] instead of full eigenvector decomposition for spectral clustering as we do results in significantly faster run time (0.1s/frame vs. 17s/frame).

3. Method

3.1. Method Overview

We consider a sequence of N video frames of size $H \times W$, where H and W are the number of rows and columns of a frame. The segmentation of the main salient object, in each frame p , is modeled as a soft, vectorized image mask $\mathbf{x}_p \in \mathbb{R}_+^{HW}$. The actual image mask at frame p is recovered by splitting the pixels into two clusters based on \mathbf{x}_p , separating the object of interest from the background.

We assume that we are given a rough mask estimate $\hat{\mathbf{x}}_p$ in each frame p , as well as functions $w_p^q(\mathbf{x}_p)$ warping a mask \mathbf{x}_p in frame p to frame q using the optical flow from p to q . Then, the objective function we minimize is:

$$\mathcal{L}(\{\mathbf{x}_p\}_p) = \sum_p \lambda \text{CE}(\hat{\mathbf{x}}_p, \mathbf{x}_p) + \text{CE}(\mathbf{x}_{p+1}, w_p^{p+1}(\mathbf{x}_p)) + \text{CE}(\mathbf{x}_p, w_{p+1}^p(\mathbf{x}_{p+1})), \quad (1)$$

where $\text{CE}(\cdot)$ denotes the cross-entropy.

Intuitively, $\text{CE}(\hat{\mathbf{x}}_p, \mathbf{x}_p)$ expresses the deviation between a mask \mathbf{x}_p and the initial mask estimate $\hat{\mathbf{x}}_p$; $\text{CE}(\mathbf{x}_{p+1}, w_p^{p+1}(\mathbf{x}_p))$ expresses the deviation between mask \mathbf{x}_{p+1} and the warping of mask \mathbf{x}_p by flow w_p^{p+1} , and reciprocally for $\text{CE}(\mathbf{x}_p, w_{p+1}^p(\mathbf{x}_{p+1}))$; finally, λ is a constant weight to balance the significance of both kind of deviations (difference from initial mask vs flow discrepancy).

As illustrated in Figure 3, our objective function can be interpreted easily: The optimization starts from a first estimate of the mask for each frame of the sequence; it encourages the mask in frame p to align with the mask in frame $p+1$ after being warped by the optical flow from frame p to frame $p+1$, and vice versa, while keeping the masks close to their initialization.

This approach can in fact be derived from a formulation of the problem in terms of spectral clustering, which also provides a way to compute initial estimates $\hat{\mathbf{x}}_p$. We present this formulation and the derivation of our approach below.

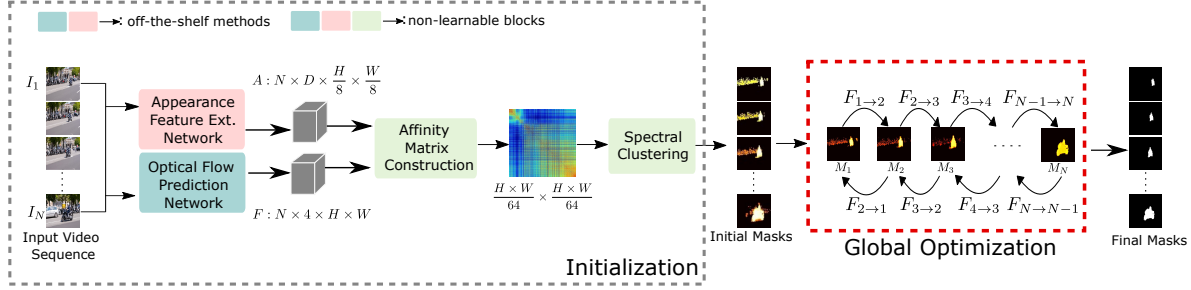


Figure 3: **Overview of our approach.** Given a video sequence, starting from first estimates for the object masks obtained by spectral clustering on each image independently, we optimize the masks so that they remain close to the first estimates while being consistent with the optical flow. The objective function we optimize to retrieve the masks can be derived from spectral clustering applied to the video sequence. Our method can rely on self-supervised visual features only.

3.2. Vanilla Spectral Clustering for Segmentation

We first briefly present spectral clustering, in the context of image segmentation as done in [60, 48], for example. We then discuss how it can be extended to video segmentation. The reader interested in more details about spectral clustering can refer to the tutorial in [73].

For image segmentation with spectral clustering, one considers an affinity matrix we will denote \mathcal{A} . Each row and each column of \mathcal{A} corresponds to an image location. The coefficient $\mathcal{A}_{ij} \geq 0$ on row i and column j should express how likely image location for row i and image location for column j belong to the same cluster. In the following, we will identify the row and column indices i and j with the corresponding image locations for simplicity.

According to spectral clustering theory, a good segmentation mask can be obtained from the second largest eigenvector \mathbf{X}_2 of the normalized affinity matrix \mathcal{W} defined as

$$\mathcal{W} = \mathcal{D}^{-1}\mathcal{A}, \quad (2)$$

where $\mathcal{D} = \text{diag}(\{\sum_j \mathcal{A}_{ij}\}_i)$ is the degree matrix. The coefficients of \mathbf{X}_2 correspond to image locations. By thresholding them, one obtains a binary mask for the segment.

To extend this approach to object segmentation in videos, we also start from an affinity matrix \mathcal{A} . Each row of \mathcal{A} , and each column, corresponds to an image location in a video frame. We will denote the coefficients of \mathcal{A} with $\mathcal{A}_{(pi)(qj)}$: (pi) is a notation for the index of row corresponding to image location i in frame p ; similarly (qj) is the index for column corresponding to image location j in frame q . Coefficients $\mathcal{A}_{(pi)(qj)}$ should express how likely image location i in frame p and image location j in frame q are to correspond to the same object. For this, we will use the similarity between their local image features and optical flow.

However, this results in a very large \mathcal{A} matrix, and thus also a very large \mathcal{W} matrix. of size $NHW \times NHW$, which is in the order of $10^9 \times 10^9$ for typical sequences in standard benchmarks. This is clearly too large in practice both

for storage and for eigenvector computation. Some methods exploit superpixels [23] or graph sparsification [35] to decrease the computational complexity, but it complicates the approach; our method is more direct and more scalable. Another approach [48] is to retrieve the second largest eigenvector \mathbf{X}_2 with the classical problem:

$$\mathbf{X}_2 = \arg \max_{\mathbf{X}} \mathbf{X}^\top \mathcal{W} \mathbf{X} \text{ such that } \|\mathbf{X}\|_2 = 1, \quad (3)$$

which could be approximately but efficiently computed using Power Iteration Clustering (PIC) [40]. However, it still cannot be scaled for common video lengths.

Nevertheless, as described below, we draw on this approach: Instead of building \mathcal{W} explicitly and struggling to compute its second largest eigenvector \mathbf{X}_2 , we compute eigenvectors $\hat{\mathbf{x}}_p$ by PIC for each frame p independently, for scalability. We use data that include inter-frame information to reinforce temporal consistency which leads to greater accuracy. This makes our initialization scheme similar to TokenCut [77], except that we utilize approximate eigenvector extraction instead of full spectral decomposition as well as inclusion of optical flow connectivity between adjacent frames. In Table 1(b), we show differences between using the baseline TokenCut method and our approximation without optical flow. These initial masks are refined using additional inter-frame consistency constraints, resulting in an optimization problem that is much more tractable than the originating spectral clustering problem.

In the following, we introduce a suitable affinity matrix \mathcal{A} , that we use in two ways: to efficiently compute good initial estimates $\hat{\mathbf{x}}_p$, and to derive the formulation in Eq. (1) as a simplification of term $\mathbf{X}^\top \mathcal{W} \mathbf{X}$ we wish to maximize.

3.3. Affinity Matrix for Video Object Segmentation

Our affinity matrix \mathcal{A} relies both on object appearance features and on optical flow features:

Appearance features. For these, we use an image feature extractor that generates an appearance feature vector ψ_p^i of

the object at location i of frame p . The idea is that the features at different locations of the same object should tend to look alike, while differing from the features located on other objects or on the background. Such appearance features can typically be learned using a dataset of annotated objects or, as in our case, using a self-supervised method such as DINO [7] or MoCo [11].

Optical flow features. We use an image flow extractor that yields the optical flow $\phi_{p,q}^i \in \mathbb{R}^2$ at location i between frame p and frame q : What is seen at location i in frame p is also seen at location $j = i + \phi_{p,q}^i$ in frame q . Such an optical flow can typically be obtained using various gradient-based formulations [22], or learned using a dataset of annotated flows [68, 66, 15, 30]. It can also be obtained, as in our case, using a self-supervised method such as ARFlow [42].

Affinity matrix. We define our affinity matrix as follows:

$$\mathcal{A} = \begin{bmatrix} A_1 & \bar{F}_1 & 0 & \dots & 0 \\ F_1 & A_2 & \bar{F}_2 & \ddots & \vdots \\ 0 & F_2 & A_3 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & \bar{F}_{N-1} \\ 0 & \dots & 0 & F_{N-1} & A_N \end{bmatrix}. \quad (4)$$

Each block is a $HW \times HW$ matrix. It contains the affinities between image locations in a frame and image locations in the same frame or in another one. We detail below matrices A_p , F_p and \bar{F}_p , and justify the use of 0 matrices.

Matrices A_p . Each matrix A_p on the block diagonal of \mathcal{A} contains the affinities between two image locations in the same frame p . It is based not only on object appearance features within frame p , but also on information from the preceding and succeeding frames $p-1$ and $p+1$. Concretely, we take the affinity $A_p^{i,j}$ between locations i and j in frame p as the following combined similarity of their object appearance features and their forward and backward flows:

$$A_p^{i,j} = \frac{1}{\alpha_\psi + 2\alpha_\phi} g_s \left(\alpha_\psi \langle \psi_p^i, \psi_p^j \rangle + \alpha_\phi \left(\langle \phi_{p,p+1}^i, \phi_{p,p+1}^j \rangle + \langle \phi_{p,p-1}^i, \phi_{p,p-1}^j \rangle \right) \right), \quad (5)$$

where α_ψ and α_ϕ are relative weights, and $g_s(\cdot)$ is a thresholding function that zeroes values lower than s . We use $s > 0$ to guarantee that the coefficients of A_p are positive.

Matrices F_p and \bar{F}_p . Each matrix F_p stores affinities between image locations in frame p and image locations in frame $p+1$. Here, we only consider the optical flow between frames p and $p+1$, and ignore the information given appearance features. (Using appearance features as well might improve the results, but it would lead to a much more complex optimization problem.) Two locations i and j in

frames p and $p+1$ are likely to both belong to the same object or both lie on the background if the optical flow maps one to the other. In that case, their affinity $F_p^{i,j}$ should be large, close to 1; otherwise, we set it to 0. We thus take:

$$F_p^{i,j} = \begin{cases} 1 & \text{if } i + \lfloor \phi_{p,p+1}^i \rfloor = j \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

where $\lfloor \phi \rfloor$ is the vector of nearest integers to ϕ . \bar{F}_p is defined likewise using the optical flow from frames $p+1$ to p .

If \mathbf{x} denotes a 2D mask in vector form, then $F_p \mathbf{x}$ is “almost” the mask $w_p^{p+1}(\mathbf{x})$ after warping by the optical flow from frame p to frame $p+1$, up to the integer discretization in Eq. (6). It also applies to $\bar{F}_p \mathbf{x}$ and $w_p^p(\mathbf{x})$.

The spectral clustering purist may notice that this expression for F_p and \bar{F}_p makes the affinity matrix \mathcal{A} not symmetric. However, \bar{F}_p is almost equal to F_p^\top ; they are slightly different because (1) optical flow is not exactly a bijection, as some pixels can appear or disappear, and (2) the predicted flow from frame p to $p+1$ is not exactly the inverse of the predicted flow from $p+1$ to p . In practice, we may consider that the affinity matrix \mathcal{A} is “almost” symmetric.

Matrices 0. Matrices that do not belong to the tri-diagonal of matrix \mathcal{A} contain affinities between image locations in two “distant”, *i.e.*, non-consecutive frames. Here, we ignore the information provided by the appearance features and the optical flow, resulting in the zero matrix. Being able to exploit the appearance features would probably slightly improve the final segmentation, but would also make the optimization significantly less scalable and more complex. Ignoring the optical flow between distant frames is a safe thing to do anyway, as its estimation is likely unreliable.

3.4. Deriving the Objective Function

By expanding the term $\mathbf{X}^\top \mathcal{W} \mathbf{X}$ found in Eq. (3), using the definition of the normalized affinity matrix \mathcal{W} in Eq. (2), the notation \mathbf{X} for its second eigenvector, and the definition of the affinity matrix \mathcal{A} in Eq. (4), we obtain:

$$\mathbf{X}^\top \mathcal{W} \mathbf{X} = \sum_p \mathbf{x}_p^\top D_p^{-1} A_p \mathbf{x}_p + \sum_p \mathbf{x}_{p+1}^\top D_{p+1}^{-1} F_p \mathbf{x}_p + \sum_p \mathbf{x}_p^\top D_p^{-1} \bar{F}_p \mathbf{x}_{p+1}, \quad (7)$$

where the D_p matrices are on the block diagonal of \mathcal{D} and are themselves diagonal.

Computing this expression seemingly involves the products between matrices A_p , F_p and vectors \mathbf{x}_p . It would require large amounts of memory and be prohibitively slow as these matrices are very large. However, we show that we do not need to store these matrices, nor do we need to compute these products for estimating masks \mathbf{x}_p .

The terms $\mathbf{x}_p^\top D_p^{-1} A_p \mathbf{x}_p$, when considered independently, lead to a spectral clustering problem per frame. The

second eigenvector $\widehat{\mathbf{x}}_p$ of matrix $W_p = D_p^{-1}A_p$ is in practice a good first estimate of \mathbf{x}_p , thanks to the combination of the object appearance features and the optical flow.

In the supplementary material, we also show that when \mathbf{x}_p is close to $\widehat{\mathbf{x}}_p$, then:

$$\mathbf{x}_p^\top D_p^{-1} A_p \mathbf{x}_p \approx (\widehat{\mathbf{x}}_p)^\top \mathbf{x}_p. \quad (8)$$

Furthermore, as mentioned in Section 3.3, the terms $F_p \mathbf{x}$ and $\overline{F}_p \mathbf{x}$ are approximations of the warpings of vector \mathbf{x} by the optical flow, i.e., $F_p \mathbf{x} \approx w_p^{p+1}(\mathbf{x})$ and $\overline{F}_p \mathbf{x} \approx w_{p+1}^p(\mathbf{x})$. As matrix D_p is diagonal, we can compute efficiently the last two sums in Eq. (7) using

$$\begin{aligned} \mathbf{x}_{p+1}^\top D_{p+1}^{-1} F_p \mathbf{x}_p &= \mathbf{x}_{p+1}^\top (\mathbf{d}_{p+1} \odot w_p^{p+1}(\mathbf{x}_p)), \\ \mathbf{x}_p^\top D_p^{-1} \overline{F}_p \mathbf{x}_{p+1} &= \mathbf{x}_p^\top (\mathbf{d}_p \odot w_{p+1}^p(\mathbf{x}_{p+1})), \end{aligned} \quad (9)$$

where \mathbf{d}_p denotes the vectorized coefficients on the diagonal of matrix D_p^{-1} and \odot is the element-wise product. The terms \mathbf{d}_p weight the warped masks. We noticed during our experiments that their influence was very limited, and we did not keep them in our objective function for simplicity.

Moreover, while spectral clustering is a powerful framework, it actually is an approximation as it relaxes the search for binary masks by looking for continuous values for the coefficients of \mathbf{X} , instead of binary ones. To encourage the optimization to recover binary values, we use the cross-entropy (minimizing it) when measuring the similarity between masks, rather than the dot product (maximizing it).

From Eqs. (3), (7), (8), (9), we finally obtain the formulation in Eq. (1). In practice, to minimize $\mathcal{L}(\{\mathbf{x}_p\}_p)$, and thus maximize the flow consistency while not deviating too much from initial estimates, we first compute the eigenvectors $\widehat{\mathbf{x}}_p$ and use them to initialize the vectors \mathbf{x}_p . Note that since the \mathbf{x}_p vectors remain close to the $\widehat{\mathbf{x}}_p$ vectors, the norm of \mathbf{X} remain approximately constant during optimization and the unit constraint in Eq. (3) is approximately satisfied. After convergence of the minimization, the soft masks \mathbf{x}_p are discretized using K -means with $K = 2$ to separate the object of interest from the background. Implementation details are provided in the supplementary material.

4. Experiments

In this section, we first introduce the implementation details of our method. Then we describe the datasets we use for comparison with other methods and make a comprehensive analysis of each component in our approach.

4.1. Datasets

We evaluate our model on three standard benchmarks in unsupervised video object segmentation: DAVIS2016, SegTrack-v2, and FBMS59. **DAVIS2016** [55] is a densely-annotated video object segmentation dataset, featuring 50

sequences and 3455 frames in total, captured in 1080p resolution at 24 FPS with precise annotation of a primary moving object at 480p. **SegTrack-v2** [38] is a densely-annotated dataset of 14 sequences with 976 frames in total. It sometimes features multiple objects in a single video and has multiple challenges such as motion blur, deformations, interactions, and objects being static. **FBMS59** [51, 4] is a dataset of 59 sequences with every 20-th frame being annotated, which yields 720 annotated frames in total. The dataset may involve multiple objects (some of which can be static), occlusions, and other challenging conditions. Since SegTrack-v2 and FBMS might feature multiple objects in one scene and since our method is interested in the main object segmentation, we merge the segmentation masks of the individual objects into one, similar to [31, 85, 82].

4.2. Metrics

Jaccard (\mathcal{J}). For all datasets, we evaluate them with the Jaccard metric \mathcal{J} , which measures the intersection-over-union between predicted masks and ground-truth masks.

Contour accuracy (\mathcal{F}). For the DAVIS2016 dataset, we also report contour accuracy. This measure treats the mask as a set of closed contours to calculate their precision and recall with respect to the annotation. The contour accuracy is then taken as $\mathcal{F} = \frac{2P_c R_c}{P_c + R_c}$ where P_c is the contour precision and R_c is the contour recall.

Computational cost. On an Nvidia V100, our initialization and optimization runs on average in ~ 0.5 s/frame, where initial eigenvector extraction takes ~ 0.1 s/frame and the optimization takes ~ 0.4 s/frame.

Optimization consumes ~ 8 GB of VRAM for typical sequences in DAVIS2016, and ~ 20 MFLOPS/frame.

4.3. Ablation and Parameter Study

In order to understand the different factors that contribute to the performance of our method, we conduct a series of ablation studies. Table 1 reports the importance of the different components of our pipeline. Note that we do not apply any post-processing in our ablation experiments to show direct gains by our method. We evaluate the different aspects of our method as detailed below.

Initial masks $\widehat{\mathbf{x}}_p$ from appearance only. We study applying spectral clustering to appearance features only. We consider recent self-supervised features: DINO [7], MoCo-v3 [11], SWAV [6] and Barlow Twins (BT) [89]. Table 1(a) compares how these appearance features affect our results. We obtain the best performance with DINO. The fact that DINO largely outperforms other self-supervised features, as also noted in [77, 49], remains to be understood, but is out of scope. We also note that our method can exploit any image features. Better feature extractors in the future could even improve our results further. Table 1(a) also shows that, in every case our optimization method improves the \mathcal{J} met-

(a) Different appearance features	(\mathcal{J} \uparrow)	(\mathcal{F} \uparrow)
(BT [RN-50] + ARFlow)	53.5	30.3
(BT [RN-50] + ARFlow) + Optimization (Eq. (1))	59.5 (+6.0)	48.2 (+17.9)
(SWAV [RN-50] + ARFlow)	48.0	27.2
(SWAV [RN-50] + ARFlow) + Optimization (Eq. (1))	53.6 (+5.6)	46.0 (+18.8)
(MoCo-v3 [ViT] + ARFlow)	58.0	35.3
(MoCo-v3 [ViT] + ARFlow) + Optimization (Eq. (1))	64.2 (+6.2)	61.1 (+25.8)
(DINO [ViT] + ARFlow)	72.1	72.5
(DINO [ViT] + ARFlow) + Optimization (Eq. (1))	76.8 (+4.7)	77.0 (+4.5)
(b) Different optical flows	(\mathcal{J} \uparrow)	(\mathcal{F} \uparrow)
DINO	61.2	65.8
TokenCut [77]	62.7	62.3
DINO + Optimization (Eq. (1))	66.7 (+5.5)	70.4 (+4.6)
(DINO + RAFT)	70.7	72.9
(DINO + RAFT) + Optimization (Eq. (1))	75.3 (+4.6)	76.2 (+3.3)
(DINO + ARFlow)	72.1	72.5
(DINO + ARFlow) + Optimization (Eq. (1))	76.8 (+4.7)	77.0 (+4.5)

Table 1: **Ablation experiments on DAVIS2016 for different optical flow and appearance features for the frame-based initialization and after optimization.** (a) Effects of different self-supervised methods for appearance features. (b) Effects of optical flow and different optical flow methods on performance. We obtain all initial results using Spectral Clustering on each frame independently. “+Optimization” denotes the results after optimizing our objective function. No post processing is applied to the results. Our optimization approach consistently improves the segmentation results by a large margin for all appearance features, with or without optical flow for initializing the masks.

ric by the average value of **+5.6%** and the \mathcal{F} metric by the average value of **+16.6%**. This ablation also shows one potential application of our method, which could serve as a way to benchmark fine-grained representation capacity of self-supervised features [78, 19]: Given frozen features and strictly predefined optical flow, one can “plug-in” new features and estimate how good the representation capacity is.

As it has the most similar (graph-based) formulation to ours, we also compare to TokenCut [77] applied to each frame independently, using the same DINO features. On David2016, our mask extraction from single frames (i.e., our initial masks \hat{x}_p) is 1.5 pt behind TokenCut on the \mathcal{J} metric, while outperforming it on \mathcal{F} by 3.5 pts, which shows that our initialization is better at detecting object boundaries on individual images. More importantly, we are $\sim 170\times$ faster than TokenCut [77] (0.1 s/frame vs. 17 s/frame).

Initial masks \hat{x}_p from both appearance and flow. We apply spectral clustering on the combination of appearance features and optical flow, as in Eq. (5). We consider supervised RAFT and self-supervised ARFlow. Effects of adding the optical flow to our approach can be seen in Table 1(b). The addition of optical flow to the appearance features greatly improves single-frame clustering performance and gives us a good initialization for the object masks. (Two successive frames are used for computing the optical flow.)

Masks x_p optimized from different initial masks \hat{x}_p .

Method	Training on videos	Optical flow	Fully unsuperv.	DAVIS		STv2	FBMS59
				\mathcal{J} \uparrow	\mathcal{F} \uparrow	\mathcal{J} \uparrow	\mathcal{J} \uparrow
<i>(Mostly) Unsupervised methods:</i>							
CUT [34]		LDOF [3]	✓	55.2	55.2	54.3	57.2
FTS [53]		LDOF [67]	✓	55.8	51.1	47.8	47.7
AMD [43]	✓		✓	57.8	-	57.0	47.5
MoSeg [82]	✓	RAFT [68]		68.3	61.1	58.6	53.1
CIS [85]	✓	PWCNet [66]		71.5	70.5	62.0	63.5
DS [87]	✓	RAFT [68]		79.1	-	72.1	71.8
DyStaB [84]	✓	PWCNet [66]		80.0	-	74.2	73.2
Ours		ARFlow [42]	✓	80.2	77.5	74.9	70.0
<i>Supervised methods:</i>							
NLC [21]		SIFTFlow [41]		55.1	52.3	67.2	51.5
SFL [12]	✓	FlowNetS [16]		67.4	66.7	-	-
FSEG [31]	✓			70.7	65.3	61.4	68.4
LVO [70]	✓	MP-Net [69]		75.9	72.1	57.3	65.1
ARP [36]		CPMFlow [28]		76.2	70.6	57.2	59.8
MSgStP [29]				77.6	-	70.1	60.8
MATNet [91]	✓	PWCNet [66]		82.4	80.7	-	-
DyStaB [84]	✓	PWCNet [66]		82.8	-	74.2	75.8
3DC-Seg [46]	✓			84.3	84.7	-	-
ViTAE [81]	✓			89.2	90.4	-	-
BATMAN [88]	✓	RAFT [68]		90.7	94.2	-	-

Table 2: **Results of our approach.** We show the performance of our approach on three standard benchmarks for video object segmentation (DAVIS2016, SegTrack-v2, and FBMS59), where our approach achieves state-of-the-art results despite. We also provide a comparison with some supervised methods, where we achieve a performance comparable to some of them without using any supervision. Note that recent unsupervised methods (MoSeg, CIS, DyStaB and IKE) use supervised optical flow estimators such as RAFT [68], PWCNet [66] or FlowNet [16, 30].

We present the effects of our global optimization on different appearance features and optical flows in Table 1. Our optimization gives an average boost of **+4.9%** over the single frame clustering, validating our approach in all cases.

Number of flow steps. We found (see full results in supp. mat.) that one flow step achieves the best performance, which supports our assumption of using a tridiagonal global affinity matrix: Even in the single flow step case, all frames are tied together via optical flow. Conversely, more steps might complicate the optimization due to the potentially noisier flow estimated between distant frames because of the larger displacements and additional occlusions.

Cross-entropy vs. dot product. In our objective function, we use the cross-entropy rather than the dot product between the masks. We compared the two options experimentally and found that using the cross-entropy indeed improves the masks that we recover. We provide the full results in the supplementary material.

4.4. Comparison to the State of the Art

Here, we use our best configuration (DINO [ViT] + ARFlow + Optimization). Following CIS [85] and DyStaB [84], we use a CRF as a post-processing step. Table 2 presents the performances of our method and several estab-

lished state-of-the-art unsupervised and supervised methods. Overall, our method performs on-par with the state-of-the-art methods on DAVIS2016 and it achieves the best performance on the STv2 dataset. Our method also has high contour accuracy \mathcal{F} , which shows that our approach achieves high quality boundaries. Among unsupervised methods, our method outperforms all previous methods by achieving 80.2 \mathcal{J} and 74.9 \mathcal{J} scores on DAVIS2016 and STv2, respectively. Our method outperforms the SOTA method DyStaB albeit slightly, with a much simpler approach. Also note that, in contrast to recent unsupervised methods, our approach does not use any supervised component in the pipeline, including optical flow.

Table 2 shows the generalization ability of our method: Rather than training on a target dataset, we leverage general images features obtained from a network pretrained on a large dataset with no supervision. While it may not be optimal in some contexts due to different data distributions, we achieve an excellent performance on three different benchmarks, without training and reusing the exact same networks for feature extraction and flow computation.

Comparing ViT and CNN based methods is not straightforward as each of them has their own advantages. Although other methods do not exploit ViTs, the most recent ones do use advanced CNN networks, *e.g.*, DeepLabv3 (used in DyStaB), one of the SOTA architectures for segmentation. Besides, while we rely only on pretrained networks, although possibly on large datasets (*e.g.*, ImageNet for appearance, Sintel for flow), a number of other methods (supervised or unsupervised) are advantaged by the fact they train on the target datasets, hence on the task itself and using a data distribution closer to the test sets. Note that other methods in Table 2 also use extra data beyond the evaluated dataset, *e.g.*, DyStaB uses ImageNet-pretrained weights to initialize its network, while AMD uses Youtube-VOS (a large video dataset) pretrained weights.

Figure 5 shows some failure cases of our approach. More examples are provided in the supplementary material. Note that some of those failure cases can be removed by further postprocessing, but we do not use it to keep our approach simple. We show examples of segmentation results by our method for qualitative visual inspection in Figure 4.

5. Conclusion

Our method consists of minimizing an objective function that is intuitive, simple to implement, and can be optimized efficiently. It can be derived from spectral clustering, which gives it solid theoretical ground. It could also be used to evaluate fine-grained capabilities of modern self-supervised representations, which is still a very active area of research [78, 19]. We hope that the simplicity of our method and its connection to spectral clustering will provide others with insights for future development.




Figure 4: **Qualitative results** obtained with our approach. The quality of segmentation improves with each component. In the first row, partially segmented object is recovered; the second and third rows show that we successfully recover the small and occluded objects; in the fourth row our method removes the background residuals and finally in the fifth and sixth rows, we show that our method recovers and segments multiple close objects.



Figure 5: **Failure cases.** Our approach has 3 main failure modes: *Oversegmentation* which is most often caused by multiple similar objects in the scene, imprecise masks due to *occlusion*, and *undersegmentation* in scenes with multiple objects which leads to the parts of the primary object not being segmented. In the last example, note that when we use DINO features, our clusters tend to group a single semantic class, which is the desired behavior for object discovery, rather than segment anything that moves. This is thus a failure case w.r.t. the benchmark, but not to the goal of object discovery.

Acknowledgments

This work was granted access to the HPC resources of IDRIS under the allocation 2021-AD011012896 made by GENCI and supported in part by the Chistera IPalm project.

References

- [1] Ali Athar, Sabarinath Mahadevan, Aljosa Osep, Laura Leal-Taixé, and Bastian Leibe. Stem-Seg: Spatio-Temporal EM-beddings For Instance Segmentation in Videos. In *European Conference on Computer Vision (ECCV)*, 2020.
- [2] Pia Bideau and Erik Learned-Miller. It’s Moving! a Probabilistic Model for Causal Motion Segmentation in Moving Camera Videos. In *European Conference on Computer Vision (ECCV)*, 2016.
- [3] Thomas Brox and Jitendra Malik. Large Displacement Optical Flow: Descriptor Matching in Variational Motion Estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2010.
- [4] Thomas Brox and Jitendra Malik. Object Segmentation by Long Term Analysis of Point Trajectories. In *European Conference on Computer Vision (ECCV)*, 2010.
- [5] Christopher P. Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised Scene Decomposition and Representation. In *arXiv Preprint*, 2019.
- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. EMerging Properties in Self-Supervised Vision Transformers. In *International Conference on Computer Vision (ICCV)*, 2021.
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2017.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations. In *International Conference on Machine Learning (ICML)*, 2020.
- [10] Xinlei Chen and Kaiming He. Exploring Simple Siamese Representation Learning. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [11] Xinlei Chen, Saining Xie, and Kaiming He. An EMpirical Study of Training Self-Supervised Vision Transformers. In *arXiv Preprint*, 2021.
- [12] Jingchun Cheng, Yi-Hsuan Tsai, Shengjin Wang, and Ming-Hsuan Yang. SegFlow: Joint Learning for Video Object Segmentation and Optical Flow. In *International Conference on Computer Vision (ICCV)*, 2017.
- [13] Minsu Cho, Suha Kwak, Cordelia Schmid, and Jean Ponce. Unsupervised Object Discovery and Localization in the Wild: Part-Based Matching with Bottom-Up Region Proposals. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [14] Ioana Croitoru, Simion-Vlad Bogolin, and Marius Leordeanu. Unsupervised Learning from Video to Detect Foreground Objects in Single Images. In *International Conference on Computer Vision (ICCV)*, 2017.
- [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 *International Conference on Computer Vision (ICCV)*, 2015.
- [16] 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 *International Conference on Computer Vision (ICCV)*, 2015.
- [17] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative Unsupervised Feature Learning with Convolutional Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [18] Yuming Du, Yang Xiao, and Vincent Lepetit. Learning to Better Segment Objects from Unseen Classes with Unlabeled Videos. In *International Conference on Computer Vision (ICCV)*, 2021.
- [19] Linus Ericsson, Henry Gouk, and Timothy M. Hospedales. How Well Do Self-Supervised Models Transfer? In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [20] Alon Faktor and Michal Irani. Clustering by Composition”: Unsupervised Discovery of Image Categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 36(6), 2013.
- [21] Alon Faktor and Michal Irani. Video Segmentation by Non-Local Consensus Voting. In *British Machine Vision Conference (BMVC)*, 2014.
- [22] D. Fleet and Y. Weiss. *Optical Flow Estimation*. Springer US, 2006.
- [23] Fabio Galasso, Roberto Cipolla, and Bernt Schiele. Video Segmentation with Superpixels. In *Asian Conference on Computer Vision (ACCV)*, 2012.
- [24] Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-Object Representation Learning with Iterative Variational Inference. In *International Conference on Machine Learning (ICML)*, 2019.
- [25] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Dohersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [26] Emanuela Haller, Adina Magda Florea, and Marius Leordeanu. Iterative Knowledge Exchange Between Deep Learning and Space-Time Spectral Clustering for Unsupervised Segmentation in Videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2021.
- [27] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Rep-

- resentation Learning. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [28] Yinlin Hu, Rui Song, and Yunsong Li. Efficient Coarse-To-Fine Patchmatch for Large Displacement Optical Flow. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [29] Yuan-Ting Hu, Jia-Bin Huang, and Alexander G. Schwing. Unsupervised Video Object Segmentation Using Motion Saliency-Guided Spatio-Temporal Propagation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [30] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [31] Suyog Jain, Bo Xiong, and Kristen Grauman. FusionSeg: Learning to Combine Motion and Appearance for Fully Automatic Segmentation of Generic Objects in Videos. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [32] Suyog Dutt Jain, Bo Xiong, and Kristen Grauman. Fusionseg: Learning to Combine Motion and Appearance for Fully Automatic Segmentation of Generic Objects in Videos. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [33] Ge-Peng Ji, Keren Fu, Zhe Wu, Deng-Ping Fan, Jianbing Shen, and Ling Shao. Full-Duplex Strategy for Video Object Segmentation. In *International Conference on Computer Vision (ICCV)*, 2021.
- [34] Margret Keuper, Bjoern Andres, and Thomas Brox. Motion Trajectory Segmentation via Minimum Cost Multicuts. In *International Conference on Computer Vision (ICCV)*, 2015.
- [35] Anna Khoreva, Fabio Galasso, Matthias Hein, and Bernt Schiele. Classifier-Based Graph Construction for Video Segmentation. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [36] Yeong Jun Koh and Chang-Su Kim. Primary Object Segmentation in Videos Based on Region Augmentation and Reduction. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [37] Hala Lamdouar, Weidi Xie, and Andrew Zisserman. Segmenting Invisible Moving Objects. In *British Machine Vision Conference (BMVC)*, 2021.
- [38] Fuxin Li, Taeyoung Kim, Ahmad Humayun, David Tsai, and James M. Rehg. Video Segmentation by Tracking Many Figure-Ground Segments. In *International Conference on Computer Vision (ICCV)*, 2013.
- [39] Xiaoxiao Li and Chen Change Loy. Video Object Segmentation with Joint Re-Identification and Attention-Aware Mask Propagation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [40] Frank Lin and William W. Cohen. Power Iteration Clustering. In *International Conference on Machine Learning (ICML)*, 2010.
- [41] Ce Liu and Others. *Beyond Pixels: Exploring New Representations and Applications for Motion Analysis*. PhD thesis, Massachusetts Institute of Technology, 2009.
- [42] Liang Liu, Jiangning Zhang, Ruifei He, Yong Liu, Yabiao Wang, Ying Tai, Donghao Luo, Chengjie Wang, Jilin Li, and Feiyue Huang. Learning by Analogy: Reliable Supervision from Transformations for Unsupervised Optical Flow Estimation. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [43] Runtao Liu, Zhirong Wu, Stella Yu, and Stephen Lin. The EMergence of Objectness: Learning Zero-Shot Segmentation from Videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [44] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-Centric Learning with Slot Attention. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [45] 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 *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [46] Sabarinath Mahadevan, Ali Athar, Aljoša Ošep, Sebastian Hennen, Laura Leal-Taixé, and Bastian Leibe. Making a Case for 3D Convolutions for Object Segmentation in Videos. In *arXiv Preprint*, 2020.
- [47] 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 Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 41(6), 2018.
- [48] Marina Meila and Jianbo Shi. Learning Segmentation by Random Walks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2001.
- [49] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Deep Spectral Methods: A Surprisingly Strong Baseline for Unsupervised Semantic Segmentation and Localization. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [50] Peter Ochs and Thomas Brox. Object Segmentation in Video: A Hierarchical Variational Approach for Turning Point Trajectories into Dense Regions. In *International Conference on Computer Vision (ICCV)*, 2011.
- [51] Peter Ochs, Jitendra Malik, and Thomas Brox. Segmentation of Moving Objects by Long Term Video Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2014.
- [52] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video Object Segmentation Using Space-Time Memory Networks. In *International Conference on Computer Vision (ICCV)*, 2019.
- [53] Anestis Papazoglou and Vittorio Ferrari. Fast Object Segmentation in Unconstrained Video. In *International Conference on Computer Vision (ICCV)*, 2013.
- [54] Federico Perazzi, Anna Khoreva, Rodrigo Benenson, Bernt Schiele, and Alexander Sorkine-Hornung. Learning Video Object Segmentation from Static Images. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

- [55] 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 *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [56] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J. Black. Competitive Collaboration: Joint Unsupervised Learning of Depth, Camera Motion, Optical Flow and Motion Segmentation. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [57] Sucheng Ren, Wenxi Liu, Yongtuo Liu, Haoxin Chen, Guoqiang Han, and Shengfeng He. Reciprocal Transformations for Unsupervised Video Object Segmentation. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [58] Michael Rubinstein, Armand Joulin, Johannes Kopf, and Ce Liu. Unsupervised Joint Object Discovery and Segmentation in Internet Images. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [59] Bryan C. Russell, William T. Freeman, Alexei A. Efros, Josef Sivic, and Andrew Zisserman. Using Multiple Segmentations to Discover Objects and Their Extent in Image Collections. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [60] Jianbo Shi and Jitendra Malik. Normalized Cuts and Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2000.
- [61] Mennatullah Siam, Chen Jiang, Steven Lu, Laura Petrich, Mahmoud Gamal, Mohamed Elhoseiny, and Martin Jagersand. Video Object Segmentation Using Teacher-Student Adaptation in a Human Robot Interaction (hri) Setting. In *International Conference on Robotics and Automation (ICRA)*, 2019.
- [62] Oriane Siméoni, Gilles Puy, Huy V. Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing Objects with Self-Supervised Transformers and No Labels. In *British Machine Vision Conference (BMVC)*, 2021.
- [63] Josef Sivic, Bryan C. Russell, Andrew Zisserman, William T. Freeman, and Alexei A. Efros. Unsupervised Discovery of Visual Object Class Hierarchies. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [64] Josef Sivic, Frederik Schaffalitzky, and Andrew Zisserman. Object Level Grouping for Video Shots. In *European Conference on Computer Vision (ECCV)*, 2004.
- [65] Hongmei Song, Wenguan Wang, Sanyuan Zhao, Jianbing Shen, and Kin-Man Lam. Pyramid Dilated Deeper ConvLstm for Video Salient Object Detection. In *European Conference on Computer Vision (ECCV)*, 2018.
- [66] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs For Optical Flow Using Pyramid, Warping, and Cost Volume. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [67] Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. Dense Point Trajectories by Gpu-Accelerated Large Displacement Optical Flow. In *European Conference on Computer Vision (ECCV)*, 2010.
- [68] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *European Conference on Computer Vision (ECCV)*, 2020.
- [69] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning Motion Patterns in Videos. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [70] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning Video Object Segmentation with Visual Memory. In *International Conference on Computer Vision (ICCV)*, 2017.
- [71] Huy V. Vo, Francis Bach, Minsu Cho, Kai Han, Yann Lecun, Patrick Pérez, and Jean Ponce. Unsupervised Image Matching and Object Discovery as Optimization. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [72] Paul Voigtlaender and Bastian Leibe. Online Adaptation of Convolutional Neural Networks for Video Object Segmentation. In *arXiv Preprint*, 2017.
- [73] Ulrike Von Luxburg. A Tutorial on Spectral Clustering. *Stat. Comput.*, 17(4), 2007.
- [74] Wenguan Wang, Xiankai Lu, Jianbing Shen, David J. Crandall, and Ling Shao. Zero-Shot Video Object Segmentation via Attentive Graph Neural Networks. In *International Conference on Computer Vision (ICCV)*, 2019.
- [75] Wenguan Wang, Jianbing Shen, and Fatih Porikli. Saliency-Aware Geodesic Video Object Segmentation. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [76] 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 *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [77] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James Crowley, and Dominique Vaufreydaz. Self-Supervised Transformers for Unsupervised Object Discovery Using Normalized Cut. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [78] Zhongdao Wang, Hengshuang Zhao, Ya-Li Li, Shengjin Wang, Philip Torr, and Luca Bertinetto. Do Different Tracking Tasks Require Different Appearance Models? In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [79] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised Feature Learning via Non-Parametric Instance Discrimination. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [80] Christopher Xie, Yu Xiang, Zaid Harchaoui, and Dieter Fox. Object Discovery in Videos as Foreground Motion Clustering. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [81] Yufei Xu, Qiming ZHANG, Jing Zhang, and Dacheng Tao. Vitae: Vision transformer advanced by exploring intrinsic inductive bias. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [82] Charig Yang, Hala Lamdouar, Erika Lu, Andrew Zisserman, and Weidi Xie. Self-Supervised Video Object Segmentation

- by Motion Grouping. In *International Conference on Computer Vision (ICCV)*, 2021.
- [83] Shu Yang, Lu Zhang, Jinqing Qi, Huchuan Lu, Shuo Wang, and Xiaoxing Zhang. Learning Motion-Appearance Co-Attention for Zero-Shot Video Object Segmentation. In *International Conference on Computer Vision (ICCV)*, 2021.
- [84] Yanchao Yang, Brian Lai, and Stefano Soatto. DyStaB: Unsupervised Object Segmentation via Dynamic-Static Bootstrapping. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [85] Yanchao Yang, Antonio Loquercio, Davide Scaramuzza, and Stefano Soatto. Unsupervised Moving Object Detection via Contextual Information Separation. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [86] Zhao Yang, Qiang Wang, Luca Bertinetto, Weiming Hu, Song Bai, and Philip HS Torr. Anchor Diffusion for Unsupervised Video Object Segmentation. In *International Conference on Computer Vision (ICCV)*, 2019.
- [87] Vickie Ye, Zhengqi Li, Richard Tucker, Angjoo Kanazawa, and Noah Snavely. Deformable sprites for unsupervised video decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2657–2666, June 2022.
- [88] Ye Yu, Jialin Yuan, Gaurav Mittal, Li Fuxin, and Mei Chen. BATMAN: Bilateral Attention Transformer in Motion-Appearance Neighboring Space for Video Object Segmentation. In *ECCV*, 2022.
- [89] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In *arXiv Preprint*, 2021.
- [90] Kaihua Zhang, Zicheng Zhao, Dong Liu, Qingshan Liu, and Bo Liu. Deep Transport Network for Unsupervised Video Object Segmentation. In *International Conference on Computer Vision (ICCV)*, 2021.
- [91] Tianfei Zhou, Shunzhou Wang, Yi Zhou, Yazhou Yao, Jianwu Li, and Ling Shao. Motion-Attentive Transition for Zero-Shot Video Object Segmentation. In *American Association for Artificial Intelligence Conference (AAAI)*, 2020.