

# HODOR: High-level Object Descriptors for Object Re-segmentation in Video Learned from Static Images

Ali Athar<sup>1</sup> Jonathon Luiten<sup>1,2</sup> Alexander Hermans<sup>1</sup> Deva Ramanan<sup>2</sup> Bastian Leibe<sup>1</sup>

<sup>1</sup>RWTH Aachen University, Germany <sup>2</sup>Carnegie Mellon University, USA

{athar, luiten, hermans, leibe}@vision.rwth-aachen.de deva@cs.cmu.edu

## Abstract

Existing state-of-the-art methods for Video Object Segmentation (VOS) learn low-level pixel-to-pixel correspondences between frames to propagate object masks across video. This requires a large amount of densely annotated video data, which is costly to annotate, and largely redundant since frames within a video are highly correlated. In light of this, we propose HODOR: a novel method that tackles VOS by effectively leveraging annotated static images for understanding object appearance and scene context. We encode object instances and scene information from an image frame into robust high-level descriptors which can then be used to re-segment those objects in different frames. As a result, HODOR achieves state-of-the-art performance on the DAVIS and YouTube-VOS benchmarks compared to existing methods trained without video annotations. Without any architectural modification, HODOR can also learn from video context around single annotated video frames by utilizing cyclic consistency, whereas other methods rely on dense, temporally consistent annotations. Source code: <https://github.com/Ali2500/HODOR>

## 1. Introduction

Current state-of-the-art Video Object Segmentation (VOS) methods learn ‘space-time correspondences’ (STC), *i.e.* pixel-to-pixel correspondences, between the image frames in a video. These methods [7, 28, 47] achieve impressive results, but require a large amount of temporally dense annotated video for training. Such datasets require significant human effort, and the annotations are largely redundant since image frames within a video are highly correlated. The largest publicly available VOS dataset [46] contains only a few thousand videos. Single image datasets [14, 20], in contrast, exist with hundreds of thousands of annotated images. In this work, we explore the following question: can VOS be learned with only single-image annotations?

To this end, we propose HODOR: **H**igh-level **O**bject

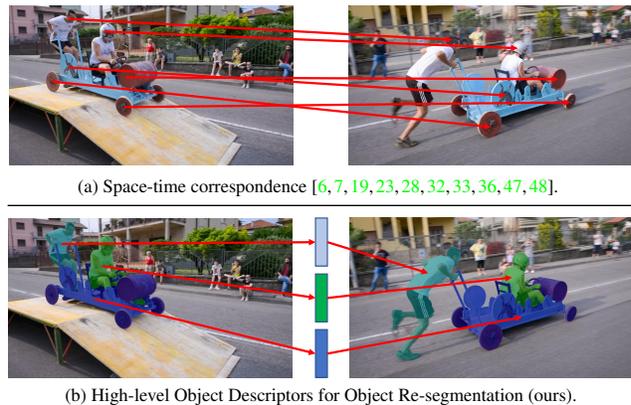


Figure 1. Previous methods (a) learn low-level pixel-pixel correspondence to propagate object masks. HODOR (b) learns high-level object descriptors to re-segment objects in a different frame.

**Descriptors for Object Re-segmentation**, a novel VOS framework which extracts a robust, *high-level descriptor* for the given objects and background in an image. These descriptors are then used to find and segment those objects in another video frame, *i.e.* re-segment them, even if the object moves or changes appearance (Fig. 1b). This differs fundamentally from STC methods which learn low-level, pixel-to-pixel correspondences (Fig. 1a). The underlying idea is that high-level object descriptors can be learned without sequential video data, as this only requires understanding object appearance, and not reasoning about motion. Thus, HODOR can be trained for VOS *using only single images without any video motion augmentation* (Fig 2a), and still be applied to video (Fig 2b). This is inherently not possible with STC methods since learning correspondences requires comparing multiple, different frames.

The key to our approach is that it forces object appearance information to pass through a concise descriptor, *i.e.* an information bottleneck. This prevents the descriptor from trivially summarizing the object mask shape and location. The network thus learns to concisely encode object appearance, and also to match the descriptor to each pixel in order

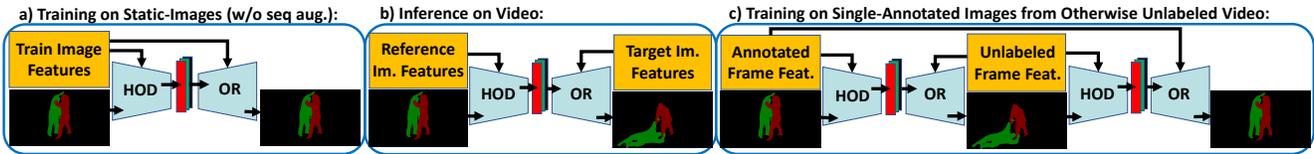


Figure 2. **HODOR train and inference strategies.** HOD: High-level Object Descriptor Encoder. OR: Object Re-segmentation Decoder. Left: HODOR can be trained with single annotated images (without sequence augmentations). Center: HODOR is run on video by feeding features from a different frame to the decoder. Right: Training HODOR can take advantage of unlabeled frames using cycle-consistency.

to re-segment the object in the same image.

If we add sequential augmentation to our single image training strategy to increase the network’s robustness, HODOR out-performs all existing methods trained with similar augmented image sequences on the DAVIS [30] and YouTube-VOS [46] benchmarks. This is because STC methods can only learn correspondences of simple motion from augmented frames, and thus cannot generalize well to the complex motion of real video. HODOR however, being based on high-level object appearance and scene context, is much more resilient to this discrepancy.

HODOR can also be trained using cycle consistency on video where only a single frame is annotated (Fig 2c). Without modifying the approach at all, we can simply propagate masks through unlabeled frames and then in reverse back to the labeled frame to apply the loss. This is enabled by a fully differentiable formulation for attending to soft input masks which allows gradients to flow through multiple frame predictions. Based on this, our network can learn to be more robust to appearance changes that occur in natural video, while only requiring single annotated frames. Current STC methods cannot be trained under this setting.

There are two further advantages: The encoder can process, and model interactions between, an arbitrary number of objects. This improves performance and makes the inference speed largely independent of the number of objects. This is in contrast to many works [6, 7, 28, 47] where part of the network requires separate forward passes per object. (2) The decoder can jointly attend to object descriptors over multiple past frames with negligible overhead. Thus, we can incorporate temporal history during inference even though the method can be trained on just single images.

To summarize: we propose a novel VOS framework that uses high-level descriptors to propagate objects across video. This enables training using just single images, with or without other unlabeled video frames. Our model processes an arbitrary number of objects simultaneously, and can readily incorporate temporal context during inference. We achieve state-of-the-art results on DAVIS and YouTube-VOS among methods trained without video annotation.

## 2. Related Work

We group existing VOS methods into three categories: pixel-pixel, object-object and object-pixel. Though not all

methods perfectly fit this taxonomy, it is nonetheless useful in comparing our approach to existing works.

**Pixel-pixel Correspondence.** Such approaches learn low-level space-time correspondence between pixels, and use these correspondences to propagate object masks between video frames. Whereas early VOS approaches [8, 13, 16, 29] used pre-computed optical flow as a measure for pixel-pixel correspondence, FEELVOS [36] was the first to learn these correspondences in an end-to-end fashion within the VOS framework, and STM [28] significantly improved upon this. Nearly all subsequent VOS methods [6, 7, 19, 23, 32, 33, 47, 48], including the two current state-of-the-arts (STCN [7] and AOT-L [48]) are based on the space-time correspondence paradigm, with each proposing various novel techniques for improving speed and performance. HODOR diverges from this paradigm by instead learning correspondences between pixels and high-level object descriptors.

**Self-supervised Pixel-pixel Correspondence.** One set of methods learns pixel-pixel correspondences using unlabelled video via self-supervision. To do this, some methods [15, 41] optimize their network with colorization and image reconstruction based training objectives. Other methods [12, 44] learn from cyclic consistency by propagating random image patches through a video sequence. HODOR can also be trained with cyclic consistency, but with the objective of learning high-level object descriptors rather than low-level pixel correspondences.

**Object-object Comparison.** Another common VOS approach involves directly comparing object representations [16, 18, 24, 39, 49]. Such methods first learn object proposals for the target image, and then match these proposals to previously tracked objects. This paradigm is inspired by methods in multi-object tracking [1, 37, 45], and often involves spatial similarity constraints and object ReID vectors [11, 16] for temporal association. Such methods require training for proposal generation on a specific set of object classes, and thus do not generalize well to novel categories.

**Object-pixel Comparison.** An alternative to the pixel-pixel and object-object approaches described above is learning high-level representations for the set of given objects, and then re-segmenting these objects in the target frame by directly comparing the representations to the pixel features in that frame. Early VOS methods followed this paradigm

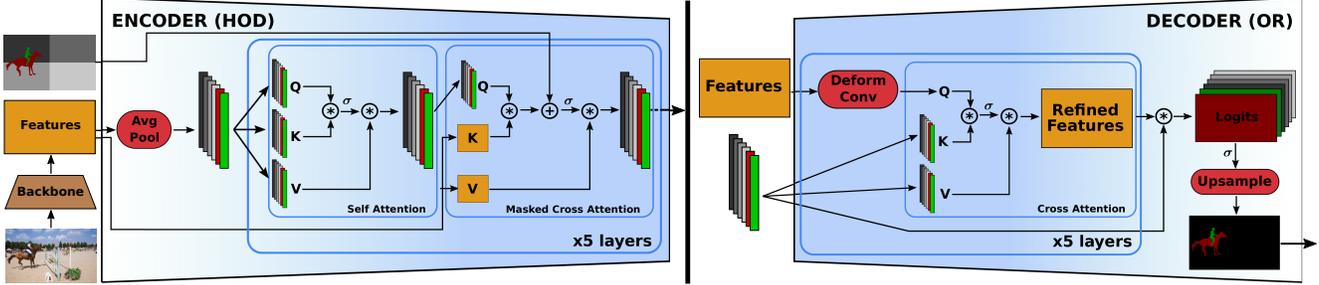


Figure 3. **The HODOR Architecture** consists of a backbone, the HOD encoder, and the OR decoder. Q, K, and V refer to Queries, Keys and Values, respectively. The encoder jointly encodes all objects and background cells (here  $2 \times 2$ ) to descriptors, which are then decoded to masks by the decoder. Some steps are simplified (the final upsampling) or omitted (fc layers, skip connections). See Sec. 3 for details.

by finetuning a segmentation network during inference on the given first frame object masks [2, 25, 29, 38] to embed an object representation in the weights of a network, which is then applied directly to subsequent frames. This is extremely slow and usually achieves poor results. The most similar work to ours is SiamMask [43]. It learns a vector representation for each object which is directly compared to pixel features to determine whether or not the pixels belong to that object. However, this approach trains on large amounts of annotated video data, and compared to HODOR and other existing methods, does not achieve good results. To the best of our knowledge, no other method from this category even achieves competitive results for VOS.

After online finetuning based methods fell out of fashion, leaderboards for VOS benchmarks were dominated by object-object association based methods [16, 24] until the emergence of FEELVOS [36] and STM [28]. Since then, state-of-the-art VOS approaches are almost exclusively based on the pixel-pixel correspondence paradigm.

### 3. Method

The HODOR network architecture consists of three components: (1) a backbone which learns multi-scale image features, (2) a High-level Object Descriptor (HOD) encoder, and (3) an Object Re-segmentation (OR) decoder. With HODOR, we revisit the idea of learning object-level descriptors for VOS which have mostly been replaced in favour of STC-based approaches. To this end, our network architecture enables the essence of an object to be encoded without directly memorizing the object mask’s shape or location. We also introduce attention layers which allow multiple objects to be processed simultaneously, and allow interactions between their descriptors. These attention layers also enable the descriptors to be enriched with image features (in the encoder), and vice versa (in the decoder).

The architecture is illustrated in Fig. 3. Given an RGB image  $I \in \mathbb{R}^{H \times W \times 3}$ , the backbone produces a pair of  $C$ -dimensional feature maps  $\mathcal{F} = \{F^4, F^8\}$  at the  $4 \times$  and  $8 \times$  downsampled input resolution scales, respectively. Assume that the image  $I$  contains  $O$  objects of interest with

segmentation masks  $\mathcal{M}^f = \{M_1^f, \dots, M_O^f\}$ . We first compute a background mask consisting of all the pixels which do not belong to any object. This background mask is then split into  $B$  separate masks  $\mathcal{M}^b = \{M_1^b, \dots, M_B^b\}$  by dividing it into a grid with  $B$  cells.

#### 3.1. Encoder

The encoder accepts as input the set of masks  $\mathcal{M}^f \cup \mathcal{M}^b$  and the image feature map  $F^8$ , and produces a set of object descriptors  $\mathcal{D}^f = \{\mathbf{d}_1^f, \dots, \mathbf{d}_O^f\}$  containing one  $C$ -dimensional descriptor per foreground object, and likewise a set of descriptors  $\mathcal{D}^b = \{\mathbf{d}_1^b, \dots, \mathbf{d}_B^b\}$  containing one  $C$ -dimensional descriptor for each background patch. Intuitively, these descriptors are a concise latent representation for their respective patches (object or background).

Each descriptor is initialized by average pooling the set of pixel features belonging to the corresponding patch. These are then iteratively and jointly refined by a series of transformer-like layers. Each layer consists of multi-head self-attention between the set of descriptors  $\mathcal{D}^f \cup \mathcal{D}^b$ , followed by multi-head cross-attention in which these descriptors absorb patch-specific information from the feature maps  $F^8$  conditioned on the masks  $\mathcal{M}^b \cup \mathcal{M}^f$ .

With some abuse of notation, let us use  $D^{(l)} = \mathcal{D}^f \cup \mathcal{D}^b \in \mathbb{R}^{(O+B) \times C}$  to denote the set of descriptors at the  $l$ -th layer of the encoder and  $M = \mathcal{M}^f \cup \mathcal{M}^b \in \mathbb{R}^{(O+B) \times H \times W}$  for the set of patch masks. The  $l$ -th layer of our encoder can then be described as follows:

$$\begin{aligned}
 D^{(l)} &\leftarrow D^{(l-1)} + \text{SelfAttn}(D^{(l-1)}) \\
 D^{(l)} &\leftarrow D^{(l)} + \text{MaskedCrossAttn}(D^{(l)}, F^8, M) \\
 D^{(l)} &\leftarrow D^{(l)} + \text{FFN}(D^{(l)})
 \end{aligned} \tag{1}$$

We omit the typical LayerNorm for the sake of text clarity (*cf.* [3]). FFN denotes a Feed-forward Network consisting of three fully-connected layers with ReLU activations. SelfAttn denotes multi-head attention [35] wherein queries, keys, and values are produced by applying separate linear projections to the input tensor. MaskedCrossAttn denotes multi-head attention where queries are produced

from the embeddings, but keys and values are produced from the image feature map  $F^8$ . We describe this operation as ‘Masked’ because we condition the pixel features  $F^8$  on the mask  $M$  to enable the descriptors to better focus on their respective patches. This could be done by replacing the dot product affinities between the descriptors  $D$  and  $F^8$  in the Key<sup>T</sup> Query matrix ( $K^T Q$ ) with  $-\infty$  for pixels where the respective mask value for the given patch is zero. However, this operation is not differentiable, and it restricts each descriptor to attend to only those pixels for which the corresponding patch mask value is one. Furthermore, numerical issues can arise for empty masks during training.

**Differentiable Soft Attention-masking.** We propose a better formulation which is differentiable, allows the mask  $M$  to be non-binary (*i.e.* have soft values), and affords the network more flexibility to focus on relevant image features. Given the mask  $M \in [0, 1]$  and a learnable, positive scalar  $\alpha$ , we define the attention operation as follows:

$$\text{softmax}\left(\frac{K^T Q + \alpha M}{\sqrt{C}}\right) \cdot V \quad (2)$$

This is identical to the standard attention operation proposed by Vaswani *et al.* [35] except for the term ‘ $+\alpha M$ ’. In practice, each attention head is assigned a different learnable parameter  $\alpha$  which is optimized during training. Thus, different attention heads attend to pixel features conditioned on different magnitudes of masking. This enables the network to learn descriptors which focus on their respective patch features, but that are also able to capture scene information from other parts of the image if this is beneficial for the training objective. This is inspired by Press *et al.* [31] who used additive offsets in temporal attention in NLP.

The encoder is thus designed to learn descriptors conditioned on object/background masks. This formulation contains an inherent information bottleneck which does not allow the input mask’s shape or location from directly ‘leaking’ into the descriptors. Specifically, in Eq. 2, the mask  $M$  can only influence the  $\text{softmax}(\cdot)$  term, *i.e.* the weights with which the Values ( $V$ ) are summed, but  $M$  cannot directly be copied into the attention operation output.

### 3.2. Decoder

Whereas the encoder produces descriptors  $\mathcal{D}^f \cup \mathcal{D}^b$  by conditioning the image features  $\mathcal{F}$  on patch masks  $\mathcal{M}^f \cup \mathcal{M}^b$ , the decoder does the opposite: it (re)produces the patch masks  $\mathcal{M}^f \cup \mathcal{M}^b$  by conditioning the image features  $\mathcal{F}$  on the descriptors  $\mathcal{D}^f \cup \mathcal{D}^b$ . The architecture is similar to that of the encoder consisting of a series of transformer-like layers with multi-head attention. However, now the image feature map  $F^8$  is updated iteratively by attending to the descriptors. There are two additional differences: (1) the cross-attention does not involve any masking. (2) The self-attention cannot be used for feature maps with large spatial

dimensions due to its quadratic memory complexity, so we instead use a  $3 \times 3$  deformable convolution [9]. Since the purpose of self-attention is to enrich pixel features by allowing them to attend to all other pixels, a deformable convolution can be thought of as having a similar effect where a pixel can interact with a set of other pixels at learned offsets. Although recent works [10, 51] proposed efficient variants of attention for image features, we found that deformable convolutions still require less memory.

Let us denote with  $F^{8(l)}$  the feature map at the  $l$ -th layer of the decoder, and let us use  $D \in \mathbb{R}^{(O+B) \times C}$  to denote the descriptors produced by the encoder. The  $l$ -th decoder layer can then be described as (again omitting LayerNorms):

$$\begin{aligned} F^{8(l)} &\leftarrow F^{8(l-1)} + \text{DeformConv}(F^{8(l-1)}) \\ F^{8(l)} &\leftarrow F^{8(l)} + \text{CrossAttn}(F^{8(l)}, D) \\ F^{8(l)} &\leftarrow F^{8(l)} + \text{FFN}(F^{8(l)}) \end{aligned} \quad (3)$$

For the `CrossAttn`, a linear projection generates the queries from the feature map  $F^{8(l)}$  and the keys and values are two separate linear projections of the descriptors  $D$ . We omit the final `FFN` from Fig. 3 for space reasons.

The final decoder layer outputs a feature map  $F^{8(L)}$ , which we bilinearly upsample by a factor of 2 and then add to the image feature map  $F^4 \in \mathcal{F}$ . We then apply a  $3 \times 3$  convolution to get  $F^{4(L)}$  and at this scale we compute the per-pixel object logits based on the dot product between  $F^{4(L)}$  and the descriptors  $D$ . The resulting logits are upsampled to the input resolution, before applying a `softmax` over the descriptor dimension, yielding the output masks  $M$ . Formally,  $M \in \mathbb{R}^{H \times W \times (O+B)}$  is calculated as follows:

$$\begin{aligned} F^{4(L)} &\leftarrow \text{Conv}\left(F^4 + \text{upsample}_2(F^{8(L)})\right) \\ M &\leftarrow \text{softmax}\left(\text{upsample}_4(F^{4(L)} \cdot D)\right) \end{aligned} \quad (4)$$

### 3.3. Video Object Segmentation

So far we discussed how the decoder can reproduce the patch masks which were input to the encoder. However, since the descriptors encode a robust representation for the objects in an image, the decoder can re-segment them in any image  $I'$  where these objects exist. Let us use  $I_t$ ,  $\mathcal{F}_t$ ,  $\mathcal{D}_t$  and  $\mathcal{M}_t$  to denote the image frame, feature maps, descriptors, and masks at frame  $t$  of a given video clip, respectively. Given the first frame  $I_1$  of a  $T$ -frame clip, and the segmentation masks  $\mathcal{M}_1^f$  for  $O$  objects in the first frame, we can learn a set of descriptors  $\mathcal{D}_1^f \cup \mathcal{D}_1^b$  which encode these objects as well as the background (*cf.* Sec. 3.1). We can then segment these objects in another video frame  $I_t$  by simply giving the decoder the feature maps  $\mathcal{F}_t$  for that frame and conditioning it on the first-frame descriptors  $\mathcal{D}_1^f \cup \mathcal{D}_1^b$ .

This strategy, however, would not generalize well to lengthy videos with significant scene changes and where

objects intersect and occlude each other. In practice, we therefore propagate the object masks sequentially frame-by-frame:  $t : 1 \rightarrow 2 \rightarrow \dots \rightarrow T$ . At each frame  $t$ , the encoder creates an updated set of object descriptors  $\mathcal{D}_t^f \cup \mathcal{D}_t^b$  from the masks  $\mathcal{M}_{t-1}$  predicted for the previous frame (or the initial input masks when  $t - 1 = 1$ ). The decoder then segments the objects in frame  $t$  using these updated descriptors.

**Temporal History.** To mitigate large object appearance and scene changes in video, existing VOS methods [7, 28, 47] incorporate temporal context from multiple past frames when predicting object masks for the current frame. HODOR also achieves the same feat efficiently: recall that the decoder is conditioned on the set  $\mathcal{D}^f \cup \mathcal{D}^b$  which contains a variable number of object/background descriptors. To incorporate temporal history when predicting masks for frame  $t$ , we simply take the union of the set of descriptors for the past  $T_p$  frames which we want to incorporate, *i.e.*  $\mathcal{D}_{t-T_p}^f \cup \mathcal{D}_{t-T_p}^b \cup \dots \cup \mathcal{D}_{t-1}^f \cup \mathcal{D}_{t-1}^b$ . In the decoder, feature map  $F_t^8$  will be refined by jointly attending to the set of all descriptors in the  $T_p$  frame history. The subsequent dot-product with the descriptors will produce a set of masks  $M_t \in \mathbb{R}^{H \times W \times T_p \times (O+B)}$ . We temporally aggregate over the time dimension to obtain masks for each of the  $O + B$  patches (we use `max` for our method).

This formulation has three advantages: (1) it incurs little computational overhead since we only need the  $O + B$  descriptors for each past frame instead of the full feature maps. (2) We can train with only single images, and still incorporate temporal context during inference without any architectural changes. (3) We can segment an arbitrary number of objects with a single forward pass of the network. This is in contrast to several VOS methods (*e.g.* [7, 28]) which require a per-object forward pass for at least part of the network.

### 3.4. Training

HODOR’s problem formulation makes it versatile with respect to the type of training data it can utilize. For the basic setting, we only need a static image dataset with annotated object masks. However, annotated image sequences, if available, can also be utilized by simply propagating the object (and background) masks over the given sequence.

Furthermore, our problem formulation enables the sequential propagation of object masks over a video to be end-to-end differentiable, *i.e.* even if we only supervise the masks predicted for the last frame of a given clip, the error will be backpropagated over the entire temporal sequence to the first frame. This allows HODOR to also be trained on unlabeled frames from videos with arbitrarily sparse and temporally inconsistent object ID annotations. Given a training clip with  $T$  frames where only frame  $t = 1$  is annotated, we can propagate the given object masks from  $t : 1 \rightarrow T$ , and then further propagate them in reverse tem-

poral order from  $t : T \rightarrow 1$ . We can then use the principle of cyclic consistency [12, 44] for supervision by supervising the predicted masks for  $t = 1$  to be identical to the input masks. The inherent information bottleneck of our method enables it to be trained effectively under this setting without trivially copying the input masks across the sequence.

## 4. Experimental Evaluation

**Datasets.** We evaluate HODOR on the DAVIS’17 [30] and YouTube-VOS 2019 [46] benchmarks. The DAVIS dataset comprises 60, 30, and 30 video sequences for training, validation and testing, respectively. YouTube-VOS is a larger dataset with 3471 videos for training and 507 for validation. For both benchmarks the task is to segment and track an arbitrary number of objects in each video. The ground truth mask for each object is only provided for the first frame in which an object appears. The evaluation measures are the  $\mathcal{J}$  score (Jaccard Index),  $\mathcal{F}$  score ( $F_1$ -score) and the average of the two ( $\mathcal{J}\&\mathcal{F}$ ) is treated as the final measure.

**Implementation Details.** Our backbone network is the ‘Tiny’ variant of the Swin transformer [21] with Feature Pyramid Network (FPN) and both our encoder and decoder consist of 5 layers. For all training settings involving static images, we use the COCO [20] dataset. Whenever training on image sequences, each sequence contains  $T = 3$  frames. We provide between 1 and 4 randomly chosen labeled objects per image/sequence. The encoder and decoder weights are randomly initialized, whereas the backbone is initialized from an off-the-shelf checkpoint trained for object detection [20]. The model is trained using the AdamW optimizer [22] with a batch size of 8 parallelized across 4 Nvidia RTX3090 GPUs. During inference, we use a temporal history comprising 7 past frames. The inference runs at  $\sim 17$  frames/s on an Nvidia RTX3090, independent of the number of instances. The dimensionality of the descriptors  $C = 256$ . See supp. material for further implementation details *e.g.* learning rate schedule, training time.

### 4.1. Training Data Versatility

Table 1 shows results on the DAVIS’17 validation set for HODOR trained under different settings. For comparison, we also report results for STCN [7], the current state-of-the-art VOS method, whenever applicable. On just single images, HODOR achieves 61.6  $\mathcal{J}\&\mathcal{F}$ , which is at the level of early VOS approaches [2, 38] which use online fine-tuning. STCN inherently requires an image sequence and cannot be trained in this setting. In row 2, we train on image sequences generated by duplicating the same image  $T$  times without any augmentation. Whereas existing space-time correspondence based methods collapse under this setting by learning to trivially copy the input mask, HODOR achieves 69.4  $\mathcal{J}\&\mathcal{F}$ . Even though this setting does not pro-



Figure 4. **Qualitative results on the DAVIS 2017 validation set:** We omit the given first frame masks and only show results for other frames. Note that fish, pigs, rope, guns, and soapbox carts are not annotated in COCO, which is used to train our model.

Table 1.  $\mathcal{J}\&\mathcal{F}$  scores for various training settings on the DAVIS 2017 validation set. The sequence length  $T = 3$  in all experiments. CC: Cyclic Consistency.

	Training Setting	HODOR	STCN [7]
1	Single image	61.6	-
2	$T \times$ duplicated image (no aug)	69.4	-
3	$T \times$ duplicated image (with aug)	77.5	75.8
4	$T \times$ dupl. video frame (with aug)	79.0	72.8
5	$T$ frames, 1 annotated (with CC)	80.6	-
6	Temporally dense video	81.3	85.4

vide any extra ‘information’ to the network compared to row 1, the  $\mathcal{J}\&\mathcal{F}$  increases significantly ( $61.6 \rightarrow 69.4$ ). This is because the model experiences noisy input masks due to the sequential propagation involved in this setting. During inference, the model can robustly track objects across videos even if the intermediate frame masks are imprecise, as it has encountered similar masks during training.

In row 3, we train on image sequences generated by applying  $T$  random affine transformations to static images. With this setting, our  $\mathcal{J}\&\mathcal{F}$  further improves to 77.5. This is because such augmentations coarsely approximate video motion, thus making the learned object descriptors more robust to object appearance and scene changes. Existing VOS methods also ubiquitously train on such augmented image sequences as a pretraining step, however we out-perform them under this setting (77.5 vs. 75.8). Qualitative results of this model can be seen in Fig. 4 and the supplementary material, including an analysis of the object descriptors.

We then explore how effectively HODOR can leverage single frame annotations that are part of a video sequence. For the next two experiments, we utilize the YouTube-VOS [46] and DAVIS [30] training sets, but assume that only one frame per video (the middle-most frame) is annotated (we only use 3,531 of the 98,797 available video frame annotations). In row 4, we fine-tune models from row 3 by similarly augmenting the selected frames. This further improves the  $\mathcal{J}\&\mathcal{F}$  from 77.5 to 79.0. STCN on the other hand performs worse (72.8), likely because of overfitting.

For row 5, we fine-tune the model from row 3, but this time using cyclic consistency by randomly sampling  $T - 1$  unlabeled frames around the single annotated frame in each video. This improves the  $\mathcal{J}\&\mathcal{F}$  from 77.5 to 80.6. The fact that this is higher than the 79.0  $\mathcal{J}\&\mathcal{F}$  in row 4 shows that HODOR can effectively learn video motion cues from unlabeled frames. Existing supervised STC methods cannot be trained with this strategy since they lack the information bottleneck needed to prevent the network from trivially copying the input mask and also because they cannot back-propagate gradients through the predicted mask.

Finally, we train on temporally dense video annotations with full supervision (row 6), which improves the  $\mathcal{J}\&\mathcal{F}$  from 77.5 to 81.3. STCN out-performs us here (85.4  $\mathcal{J}\&\mathcal{F}$ ) because the same information bottleneck which enables us to train on single images and unlabeled frames with cyclic consistency also has the drawback of limiting the network’s access to fine-grained video motion cues. By contrast, pixel-to-pixel correspondence methods lack such a bottleneck thus enabling them to better leverage dense video data. Nonetheless, to the best of our knowledge, we are the first to surpass 81  $\mathcal{J}\&\mathcal{F}$  on DAVIS’17 using an approach not based on pixel-to-pixel correspondences. Note also that our approach has much better scaling properties since we require only one frame annotation per video as opposed to the dense annotations required by existing methods.

## 4.2. Comparison to State-of-the-art

In Table 2, we report results for existing VOS methods categorized by the type of training data used. Results for HODOR are given for two settings: (1) when trained on augmented image sequences from COCO [20] (*cf.* Table 1, row 3), and (2) after fine-tuning with cyclic consistency using just one labeled frame per training set video (*cf.* Table 1, row 5). We use the same model checkpoint for all three benchmarks. For the sake of completeness we also list results for methods that do not require any annotations, and also those trained on densely annotated video.

In the ‘Labeled Images’ category, HODOR trained on COCO achieves 77.5  $\mathcal{J}\&\mathcal{F}$  on DAVIS’17, outperforming all existing methods. This includes earlier methods [2, 25,

Table 2. Quantitative results on the DAVIS and YouTube-VOS datasets. For YouTube-VOS we focus on the 2019 validation set, but substitute 2018 validation set results when only those are available (slightly higher, highlighted in grey). As is common, we evaluate unseen (us) and seen (s) object classes separately for Youtube-VOS, UI<sup>†</sup>: Unlabeled Images, OL: Online Fine-tuning, \*: retrained by us.

		DAVIS val 17			DAVIS test-dev 17			YouTube-VOS val 18 /val 19					
		OL	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}_{us}$	$\mathcal{F}_{us}$	$\mathcal{J}_s$	$\mathcal{F}_s$
UI <sup>†</sup>	DINO [4]		71.4	67.9	74.9	-	-	-	-	-	-	-	-
Labeled images	OSVOS [2]	✓	60.3	61.6	69.1	49.0	45.6	52.5	-	-	-	-	-
	OnAVOS [38]	✓	63.6	61.0	66.1	57.7	54.8	60.5	-	-	-	-	-
	OSVOS <sup>S</sup> [25]	✓	68.0	64.7	71.3	57.5	-	-	-	-	-	-	-
	STM (5x Mix) [28]		60.0	57.9	62.1	-	-	-	69.1	-	-	-	-
	DMN+AOA (COCO) [18]		67.9	65.8	70.0	-	-	-	-	-	-	-	-
	KMN (5x Mix) [32]		68.9	67.1	70.8	-	-	-	-	-	-	-	-
	STCN (5x Mix) [7]		75.8	73.1	78.6	-	-	-	-	-	-	-	-
	CFBI (COCO) [47] *		57.6	55.8	59.4	54.0	51.7	56.3	68.4	64.7	70.9	68.2	69.7
	STCN (COCO) [7] *		55.0	52.3	57.6	51.7	49.5	54.0	69.4	66.4	73.9	67.8	69.3
	HODOR (Ours, COCO)		77.5	74.7	80.2	65.0	62.6	67.4	71.7	<b>65.5</b>	<b>71.4</b>	73.7	76.0
HODOR (Ours, COCO + CC)		<b>80.6</b>	<b>77.8</b>	<b>83.4</b>	<b>66.0</b>	<b>63.6</b>	<b>68.4</b>	<b>72.4</b>	63.9	70.5	<b>75.8</b>	<b>79.3</b>	
Unlabeled videos	MAST [15]		65.5	63.3	67.6	-	-	-	64.9	61.5	68.4	64.3	65.3
	STC-CRW [12]		67.6	64.8	70.2	-	-	-	-	-	-	-	-
	MAMP [27] (uses optical flow)		<b>70.4</b>	<b>68.7</b>	<b>72.0</b>	-	-	-	<b>68.2</b>	<b>65.4</b>	<b>73.7</b>	<b>66.3</b>	<b>67.5</b>
Labeled videos	FEELVOS [36]		71.5	69.1	74.0	57.8	55.2	60.5	-	-	-	-	-
	AFB-URR [19]		74.6	73.0	76.1	-	-	-	79.6	74.1	82.6	78.8	83.1
	e-OSVOS [26]	✓	77.2	74.4	80.0	64.8	60.9	68.6	71.4	74.3	74.3	71.7	66.0
	STM [28]		81.8	79.2	84.3	-	-	-	79.4	72.8	80.9	79.7	84.2
	CFBI [47]		81.9	79.1	84.6	74.8	71.1	78.5	81.0	75.2	83.0	80.6	85.1
	EG-VOS [23]		82.8	80.2	85.2	-	-	-	80.2	74.0	80.9	80.7	85.1
	KMN [32]		82.8	80.0	85.6	77.2	74.1	80.3	81.4	75.3	83.3	81.4	85.6
	DMN+AOA [18]		84.5	81.4	87.5	78.3	74.8	81.7	82.7	76.7	84.8	82.6	87.0
	HMMN [33]		84.7	81.9	87.5	78.6	74.7	82.5	82.5	77.3	85.0	81.7	86.1
	STCN [7]		<b>85.4</b>	82.2	<b>88.6</b>	76.5	73.1	79.6	82.7	78.2	85.9	81.1	85.4
AOT-L [48]		<b>85.4</b>	<b>82.4</b>	88.4	<b>81.2</b>	<b>77.3</b>	<b>85.1</b>	<b>84.5</b>	<b>78.4</b>	<b>86.7</b>	<b>84.0</b>	<b>88.8</b>	

29, 38] that perform online fine-tuning (best score: 68.0  $\mathcal{J}\&\mathcal{F}$  from OSVOS<sup>S</sup> [25]), and also current state-of-the-art methods which pre-train on similar augmented image sequences. The best performing method among these is STCN (75.8  $\mathcal{J}\&\mathcal{F}$ ) which is 1.7  $\mathcal{J}\&\mathcal{F}$  lower than our 77.5. It is worth noting that while DMN+AOA [18] use COCO images for this training step, STM [28], KMN [32] and STCN [7] use a collection of 5 image datasets [5, 17, 34, 42, 50] ('5x Mix' in the table). To verify that this discrepancy does not disadvantage other methods, we retrained STCN and CFBI on COCO images using their respective training code. These results are marked with '\*' in the table. We see that HODOR performs better than both, but the difference for YouTube-VOS is smaller than that for DAVIS. A possible explanation for the large performance difference of these methods on the two datasets is that the augmentations applied to static images are quite aggressive, which make objects undergo significant movement across frames. This better approximates YouTube-VOS videos where objects also frequently undergo large motions. By contrast, object motion in DAVIS videos is milder, and because these

methods learn pixel-to-pixel correspondences, they do not perform well during inference if the nature of object motion is different from what was encountered during training.

Finally, we also report our result after fine-tuning on cyclic consistency using only the middle-most annotated frame from each video in the YouTube-VOS and DAVIS training set. This improves the  $\mathcal{J}\&\mathcal{F}$  by 3.1, 1.0 and 0.7 points on the DAVIS validation, DAVIS test and YouTube-VOS validation sets, respectively.

### 4.3. Ablations

We perform ablations to investigate our design choices and report the results in Table 3. Additional ablations and failure cases are given in supplementary.

**Multi-instance.** Unlike most other VOS methods, HODOR handles all objects in a single forward pass. This enables it to effectively utilize multi-object context and learn better descriptors. For experiment (1), we train and infer with only one foreground instance, and merge multiple forward passes by running a pixel-wise `argmax` over the

object logits during inference. This reduces the  $\mathcal{J}\&\mathcal{F}$  from 77.5 to 71.5, clearly highlighting the benefit of our multi-instance approach which also increases inference speed by negating the need for per-object forward passes.

**Attention Masking in Encoder.** Recall from Sec. 3.1 that we condition the descriptors on their respective patch masks with our proposed soft attention mechanism. Experiment (2) shows the result without attention masking. Thus, the only cue for the descriptors to specialize to their respective targets is their initialization (average pooling over the target pixel features); this reduces the  $\mathcal{J}\&\mathcal{F}$  to 74.4. For experiment (3), we apply hard attention masking by thresholding the masks at 0.5 and setting the  $K^TQ$  matrix entries inside the attention operation to  $-\infty$  for pixels where the mask is zero. This strategy yields a similarly reduced  $\mathcal{J}\&\mathcal{F}$  of 74.5. The performance increase from using our learned soft attention masking shows that it helps the encoder to better condition the descriptors on the given object/background masks.

**Encoder/Decoder Layers.** For experiments (4-7), we ablate the depth of the encoder and decoder. Reducing the depth from 5 to 3 in either of them reduces the  $\mathcal{J}\&\mathcal{F}$  by  $\sim 1$ . For the zero-layer case, the 5 encoder layers are replaced by a single MLP consisting of 3 fully-connected layers and the 5 decoder layers are replaced with two  $3 \times 3$  convolutions. For the zero-layer decoder, the  $\mathcal{J}\&\mathcal{F}$  reduces to 74.4, whereas the zero-layer encoder reduces the  $\mathcal{J}\&\mathcal{F}$  to 72.8. This shows that while both components play an important role in the overall performance, the encoder has a larger impact. One reason could be that whereas self-attention in the encoder allows object descriptors to interact, such interactions are not as profound in the decoder where deformable convolutions are used instead of self-attention.

**Deformable Convolution in Decoder.** Due to memory constraints, we use a  $3 \times 3$  deformable convolution [9] instead of the self-attention operation (*cf.* Sec. 3.2). In experiment (8), we instead use a regular  $3 \times 3$  convolution and observe a reduction in  $\mathcal{J}\&\mathcal{F}$  from 77.5 to 75.1. This highlights the importance of substituting self attention with an operation that is able to attend to far-away spatial locations.

**Background Descriptors.** We use nine background descriptors initialized by dividing the image into a  $3 \times 3$  grid and average pooling the background pixel features in each cell. This gives HODOR more flexibility to model the background. For experiment (9) we instead use a single background descriptor, reducing the  $\mathcal{J}\&\mathcal{F}$  from 77.5 to 76.2.

## 5. Discussion

**Limitations.** Aside from our performance on dense video data (discussed in Sec. 4.1), another limitation of our method is that when there are distractor objects in the scene with similar appearances, HODOR sometimes compels it-

Table 3. Several ablation results on the DAVIS 2017 validation set.

	Setting	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
1	Single foreground instance	71.5	69.2	73.9
2	No masking in encoder	74.4	71.5	77.2
3	Hard masking in encoder	74.5	71.8	77.1
4	# layers in encoder: 5 $\rightarrow$ 0	72.8	70.5	75.2
5	# layers in encoder: 5 $\rightarrow$ 3	76.6	73.9	79.4
6	# layers in decoder: 5 $\rightarrow$ 0	74.4	71.7	77.1
7	# layers in decoder: 5 $\rightarrow$ 3	76.4	73.6	79.3
8	Regular convolution in decoder	75.1	72.0	78.2
9	1x background descriptor	76.2	73.7	78.7
	HODOR	77.5	74.7	80.2

self to segment an object even if that object has moved out of the video scene. Since such cases arise more frequently in YouTube-VOS videos, this is one reason why our  $\mathcal{J}\&\mathcal{F}$  score for YouTube-VOS is lower than that for DAVIS. In future work, improved training strategies could be formulated to better optimize the model for such challenging cases.

**Ethical Considerations.** As with most computer vision methods, the dual-use dilemma can and should not be ignored. However, it is unlikely that our approach could be utilized to facilitate negative use-cases (*e.g.* population tracking or surveillance) more effectively than dedicated approaches for these applications. Another important ethical aspect is that dataset annotation is often performed by an exploited labor force deprived of minimum wage and/or legally binding benefits. Reducing the need for such annotations can thus be seen as a positive aspect of our approach.

## 6. Conclusion

We proposed a novel VOS approach which uses high-level descriptors for encoding and propagating objects across video. Our approach contains an information bottleneck which enables training on single images and unlabeled frames using cyclic consistency. Thus, unlike existing STC based methods which train on dense video data, HODOR can be trained on static images, or on videos with arbitrarily sparse, temporally inconsistent frame annotations. Since annotating single frames is easier than dense video, HODOR has strong potential for scaling up performance by learning from large-scale video datasets with sparse, or even automatically generated frame annotations [40].

**Acknowledgements.** This project was funded in part by ERC Consolidator Grant DeeVise (ERC-2017-COG-773161). Compute was granted by RWTH Aachen University under project ‘supp0003’. We thank Paul Voigtlaender, Jonas Schult, Istvan Sarandi, Sabarinath Mahadevan, Markus Knoche, Christian Schmidt, Jason Zhang, Gautam Gare and Tarasha Khurana for helpful discussions.

## References

- [1] P. Bergmann, T. Meinhardt, and L. Leal-Taixe. Tracking without bells and whistles. In *ICCV*, 2019. 2
- [2] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-Shot Video Object Segmentation. In *CVPR*, 2017. 3, 5, 6, 7
- [3] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. End-to-End Object Detection with Transformers. In *ECCV*, 2020. 3
- [4] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *ICCV*, 2021. 7
- [5] H. K. Cheng, J. Chung, Y.-W. Tai, and C.-K. Tang. Cascadepsp: toward class-agnostic and very high-resolution segmentation via global and local refinement. In *CVPR*, 2020. 7
- [6] H. K. Cheng, Y.-W. Tai, and C.-K. Tang. Modular Interactive Video Object Segmentation: Interaction-to-Mask, Propagation and Difference-Aware Fusion. In *CVPR*, 2021. 1, 2
- [7] H. K. Cheng, Y.-W. Tai, and C.-K. Tang. Rethinking Space-Time Networks with Improved Memory Coverage for Efficient Video Object Segmentation. In *NeurIPS*, 2021. 1, 2, 5, 6, 7
- [8] J. Cheng, Y.-H. Tsai, S. Wang, and M.-H. Yang. Segflow: Joint learning for video object segmentation and optical flow. In *ICCV*, 2017. 2
- [9] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei. Deformable convolutional networks. In *ICCV*, 2017. 4, 8
- [10] X. Dai, Y. Chen, J. Yang, P. Zhang, L. Yuan, and L. Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *ICCV*, 2021. 4
- [11] A. Hermans\*, L. Beyer\*, and B. Leibe. In Defense of the Triplet Loss for Person Re-Identification. *arXiv:1703.07737*, 2017. 2
- [12] A. Jabri, A. Owens, and A. A. Efros. Space-Time Correspondence as a Contrastive Random Walk. In *NeurIPS*, 2020. 2, 5, 7
- [13] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele. Lucid Data Dreaming for Video Object Segmentation. *IJCV*, 127(9):1175–1197, 2019. 2
- [14] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov, et al. The open images dataset v4. *IJCV*, 128(7):1956–1981, 2020. 1
- [15] Z. Lai, E. Lu, and W. Xie. MAST: A Memory-Augmented Self-supervised Tracker. In *CVPR*, 2020. 2, 7
- [16] X. Li and C. C. Loy. Video object segmentation with joint re-identification and attention-aware mask propagation. In *ECCV*, 2018. 2, 3
- [17] X. Li, T. Wei, Y. P. Chen, Y.-W. Tai, and C.-K. Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *CVPR*, 2020. 7
- [18] S. Liang, X. Shen, J. Huang, and X.-S. Hua. Video Object Segmentation With Dynamic Memory Networks and Adaptive Object Alignment. In *ICCV*, 2021. 2, 7
- [19] Y. Liang, X. Li, N. Jafari, and J. Chen. Video Object Segmentation with Adaptive Feature Bank and Uncertain-Region Refinement. In *NeurIPS*, 2020. 1, 2, 7
- [20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common Objects in Context. In *ECCV*, 2014. 1, 5, 6
- [21] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 5
- [22] I. Loshchilov and F. Hutter. Decoupled Weight Decay Regularization. In *ICLR*, 2019. 5
- [23] X. Lu, W. Wang, M. Danelljan, T. Zhou, J. Shen, and L. Van Gool. Video Object Segmentation with Episodic Graph Memory Networks. In *ECCV*, 2020. 1, 2, 7
- [24] J. Luiten, P. Voigtlaender, and B. Leibe. PReMVOS: Proposal-Generation, Refinement and Merging for Video Object Segmentation. In *ACCV*, 2018. 2, 3
- [25] K.-K. Maninis, S. Caelles, Y. Chen, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. Video Object Segmentation Without Temporal Information. *PAMI*, 41(6):1515–1530, 2018. 3, 6, 7
- [26] T. Meinhardt and L. Leal-Taixé. Make One-Shot Video Object Segmentation Efficient Again. In *NeurIPS*, 2020. 7

- [27] B. Miao, M. Bennamoun, Y. Gao, and A. Mian. Self-Supervised Video Object Segmentation by Motion-Aware Mask Propagation. *arXiv:2107.12569*, 2021. [7](#)
- [28] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim. Video Object Segmentation using Space-Time Memory Networks. In *ICCV*, 2019. [1](#), [2](#), [3](#), [5](#), [7](#)
- [29] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *CVPR*, 2017. [2](#), [3](#), [6](#)
- [30] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 DAVIS Challenge on Video Object Segmentation. *arXiv:1704.00675*, 2017. [2](#), [5](#), [6](#)
- [31] O. Press, N. A. Smith, and M. Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv:2108.12409*, 2021. [4](#)
- [32] H. Seong, J. Hyun, and E. Kim. Kernelized Memory Network for Video Object Segmentation. In *ECCV*, 2020. [1](#), [2](#), [7](#)
- [33] H. Seong, S. W. Oh, J.-Y. Lee, S. Lee, S. Lee, and E. Kim. Hierarchical Memory Matching Network for Video Object Segmentation. In *ICCV*, 2021. [1](#), [2](#), [7](#)
- [34] J. Shi, Q. Yan, L. Xu, and J. Jia. Hierarchical image saliency detection on extended cssd. *PAMI*, 2015. [7](#)
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [3](#), [4](#)
- [36] P. Voigtlaender, Y. Chai, F. Schroff, H. Adam, B. Leibe, and L.-C. Chen. FEELVOS: Fast End-To-End Embedding Learning for Video Object Segmentation. In *CVPR*, 2019. [1](#), [2](#), [3](#), [7](#)
- [37] P. Voigtlaender, M. Krause, A. Osep, J. Luiten, B. B. G. Sekar, A. Geiger, and B. Leibe. MOTs: Multi-Object Tracking and Segmentation. In *CVPR*, 2019. [2](#)
- [38] P. Voigtlaender and B. Leibe. Online Adaptation of Convolutional Neural Networks for Video Object Segmentation. In *BMVC*, 2017. [3](#), [5](#), [6](#), [7](#)
- [39] P. Voigtlaender, J. Luiten, P. H. Torr, and B. Leibe. Siam R-CNN: Visual Tracking by Re-Detection. In *CVPR*, 2020. [2](#)
- [40] P. Voigtlaender, L. Luo, C. Yuan, Y. Jiang, and B. Leibe. Reducing the annotation effort for video object segmentation datasets. In *WACV*, 2021. [8](#)
- [41] C. Vondrick, A. Shrivastava, A. Fathi, S. Guadarrama, and K. Murphy. Tracking Emerges by Colorizing Videos. In *ECCV*, 2018. [2](#)
- [42] L. Wang, H. Lu, Y. Wang, M. Feng, D. Wang, B. Yin, and X. Ruan. Learning to detect salient objects with image-level supervision. In *CVPR*, 2017. [7](#)
- [43] Q. Wang, L. Zhang, L. Bertinetto, W. Hu, and P. H. Torr. Fast Online Object Tracking and Segmentation: A Unifying Approach. In *CVPR*, 2019. [3](#)
- [44] X. Wang, A. Jabri, and A. A. Efros. Learning Correspondence from the Cycle-Consistency of Time. In *CVPR*, 2019. [2](#), [5](#)
- [45] N. Wojke, A. Bewley, and D. Paulus. Simple online and realtime tracking with a deep association metric. In *ICIP*, 2017. [2](#)
- [46] N. Xu, L. Yang, Y. Fan, D. Yue, Y. Liang, J. Yang, and T. Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv:1809.03327*, 2018. [1](#), [2](#), [5](#), [6](#)
- [47] Z. Yang, Y. Wei, and Y. Yang. Collaborative Video Object Segmentation by Foreground-Background Integration. In *ECCV*, 2020. [1](#), [2](#), [5](#), [7](#)
- [48] Z. Yang, Y. Wei, and Y. Yang. Associating Objects with Transformers for Video Object Segmentation. In *NeurIPS*, 2021. [1](#), [2](#), [7](#)
- [49] X. Zeng, R. Liao, L. Gu, Y. Xiong, S. Fidler, and R. Urtasun. Dmm-net: Differentiable mask-matching network for video object segmentation. In *ICCV*, 2019. [2](#)
- [50] Y. Zeng, P. Zhang, J. Zhang, Z. Lin, and H. Lu. Towards high-resolution salient object detection. In *ICCV*, 2019. [7](#)
- [51] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2020. [4](#)