

A Transductive Approach for Video Object Segmentation

Yizhuo Zhang^{2*} Zhirong Wu^{1*}
¹Microsoft Research Asia

Houwen Peng¹ Stephen Lin¹
²Carnegie Mellon University

Abstract

Semi-supervised video object segmentation aims to separate a target object from a video sequence, given the mask in the first frame. Most of current prevailing methods utilize information from additional modules trained in other domains like optical flow and instance segmentation, and as a result they do not compete with other methods on common ground. To address this issue, we propose a simple yet strong transductive method, in which additional modules, datasets, and dedicated architectural designs are not needed. Our method takes a label propagation approach where pixel labels are passed forward based on feature similarity in an embedding space. Different from other propagation methods, ours diffuses temporal information in a holistic manner which take accounts of long-term object appearance. In addition, our method requires few additional computational overhead, and runs at a fast ~ 37 fps speed. Our single model with a vanilla ResNet50 backbone achieves an overall score of 72.3% on the DAVIS 2017 validation set and 63.1% on the test set. This simple yet high performing and efficient method can serve as a solid baseline that facilitates future research. Code and models are available at <https://github.com/microsoft/transductive-vos.pytorch>.

1. Introduction

Video object segmentation addresses the problem of extracting object segments from a video sequence given the annotations in the starting frame. This semi-supervised setting is challenging as it requires the system to generalize to various objects, deformations, and occlusions. Nevertheless, video object segmentation has received considerable attention because of its broad practical applications in surveillance, self-driving cars, robotics, and video editing.

Despite the simplicity of the formulation, video object segmentation is closely related to many other visual problems, such as instance segmentation [19], object re-identification [13], optical flow estimation [15], and object

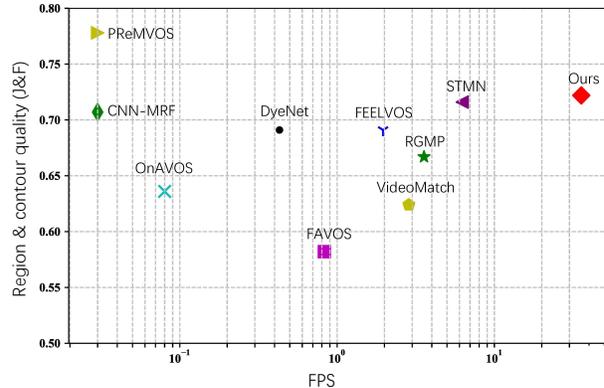


Figure 1: A comparison of performance and speed for semi-supervised video object segmentation methods on the DAVIS 2017 validation set. Ours performs comparably to the state-of-the-art methods, while running at an online speed (> 30 fps).

tracking [5]. As these tasks share similar challenges with video object segmentation, previous efforts [29, 30] attempt to transfer the modules trained for such tasks into video object segmentation pipeline. More specifically, optical flow and tracking encourage local dependencies by estimating displacements in nearby frames, while instance segmentation and object re-identification enforces global dependencies by learning invariances to large appearance changes. The integration of such modules allows a significant performance improvement in video object segmentation.

The idea of enforcing local and global dependencies has been a central topic in general semi-supervised learning [49, 14] (also known as transductive inference). The basic assumptions are: 1) nearby samples tend to have the same label and 2) samples that lie on the same manifold should have the same label. The local and global dependencies describe a sufficiently smooth affinity distribution, so that label propagation on the unlabeled data gives reliable estimates. Prior classical approaches that realize this idea include random walk [37], graph-cut [6] and spectral methods [4].

This inspires us to explore a unified approach for semi-supervised video object segmentation without the integration of the modules derived from other domains. We model

*Equal contribution. Work done when Yizhuo was an intern at MSRA.

the local dependency through a spatial prior and a motion prior. It is based on the assumption that spatially nearby pixels are likely to have same labels and that temporally distant frames weakens the spatial continuity. On the other hand, we model the global dependency through visual appearance, which is learned by convolutional neural networks on the training data.

The inference follows the regularization framework [49] which propagates labels in the constructed spatio-temporal dependency graph. While label propagation algorithms have been explored in the recent literature for video object segmentation [41, 9, 22, 36, 34], the manner in which they learn and propagate affinity is sparse and local, i.e., learning pixel affinities either between adjacent frames or between the first frame and a distant frame. We observe that there exists much smooth unlabeled structure in a temporal volume that these methods do not exploit. This may cause failures when handling deformations and occlusions. In contrast, our label propagation approach attempts to capture all frames which span the video sequence from the first frame to the frame preceding the current frame. To limit the computational overhead, sampling is performed densely within the recent history and sparsely in the more distant history, yielding a model that accounts for object appearance variation while reducing temporal redundancy.

In its implementation, our model does not rely on any other task modules, additional datasets, nor dedicated architectural designs beyond a pretrained ResNet-50 model from the ImageNet model zoo [20]. During inference, per-frame prediction involves only a feed-forward pass through the base network plus an inner product with the prediction history. Thus the inference is fast and also not affected by the number of objects. Experimentally, our model runs at a frame rate of 37 per second, achieving an overall score of 72.3% on Davis 2017 validation set, as well as 63.1% on Davis 2017 test set. Our model also achieves a competitive overall score of 67.8% on the recent Youtube-VOS validation set. Our method is competitive to current prevailing methods while being substantially simpler and faster. We hope the model can serve as a simple baseline for future works.

2. Related Work

We review related work on video object segmentation in the semi-supervised setting. For an overview of unsupervised and interactive video object segmentation, we refer readers to other papers [11, 12, 2, 39, 27, 28].

Single frame models. In the past few years, methods with leading performance have been based on finetuning the model on the single annotated frame and performing inference on individual test frames. These methods basically learn an objectness prior and spatial continuity without con-

sidering temporal information. The convolutional neural network architecture plays an important role for finetuning on a single frame to be effective. OSVOS [7] is the pioneering work in this direction. Lucid [25] seeks to augment the data specifically for each video from only a single frame of ground truth. OnAVOS [42] mines confident regions in the testing sequence to augment the training data. The later work of OSVOS-S [31] integrates semantic information from an instance segmentation model to boost performance. PReMVOS [30], CNN-MRF [3], and DyeNet [29] also build on top of single frame models.

The effectiveness of single frame models demonstrates that optimizing a domain-specific spatial smoothness term greatly enhances performance. However, finetuning via gradient descent generally takes tens of seconds per video, which can make it impractical for many applications.

Propagation-based models. Propagation-based methods embed image pixels into a feature space and utilize pixel similarity in the feature space to guide label propagation. In methods such as VideoMatch [22, 9], only pixels in the first frame are used for reference in computing pixel similarity. Since no finetuning at run time is involved, propagation-based models run much faster than the aforementioned single frame models, but the lack of domain-specific finetuning leads to performance that is much worse. Later works [36, 45, 34, 41] explore adding the preceding frame to the first frame as reference, which significantly improves performance and leads to greater temporal smoothness. However, this local and sparse propagation scheme suffers from the drifting problem [16].

Long-range spatio-temporal models. There are two lines of work which attempt to optimize over a dense long-range spatio-temporal volume. The first [45, 21] builds a recurrent neural network which uses the estimate from the previous frame to predict the object segmentation in the current frame. The whole model is learned via backpropagation through time. However, such models are sensitive to estimation errors in the previous frame.

The second direction is based on graphical models [17, 38, 26, 48, 8, 32] (i.e., Markov Random Fields) defined over the spatio-temporal domain. These works were popular prior to deep learning, and employed edge potentials defined by handcrafted features such as SIFT. The models are computationally expensive and no longer competitive to learning-based methods.

Relation to other vision problems. As the above methods suggest, video object segmentation is closely related to a variety of computer vision problems such as instance segmentation, object re-identification, and optical flow estimation and tracking. Many recent methods integrate components for these other tasks into the video object segmentation pipeline. For example, OSVOS-S [31] includes an instance

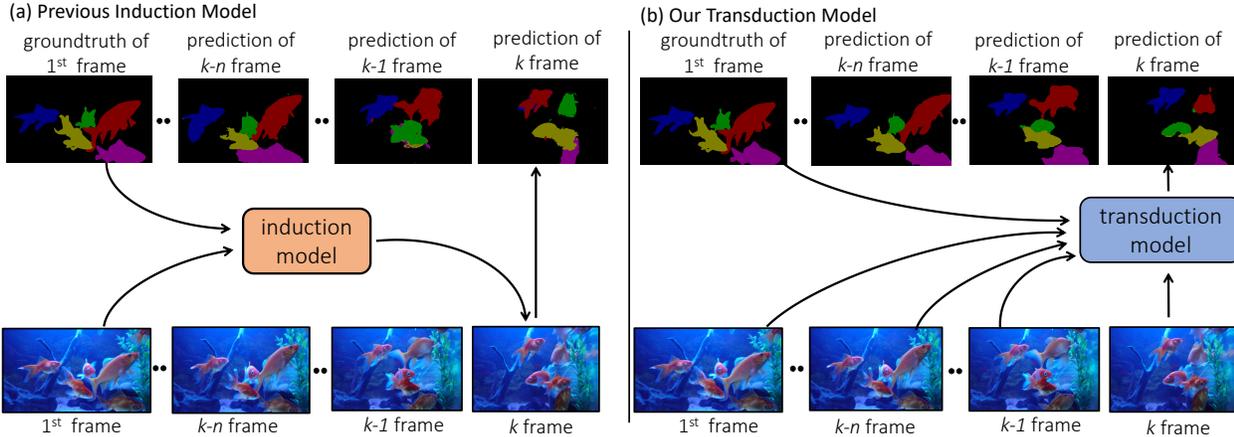


Figure 2: We pose video object segmentation from a transductive inference perspective, where dense long-term similarity dependencies are constructed to discover structures in the spatio-temporal volume. a) Previous induction model transfers knowledge from the first frame to other frames. b) Our transduction model considers holistic dependencies in the unlabeled spatio-temporal volume for joint inference.

segmentation module; PRemVOS [30] and DyeNet [29] incorporate a object re-identification module; CNN-MRF [3], MaskTrack [34] and MaskRNN [21] rely on optical flow estimation. The integration of other modules heavily depends on transfer learning from other datasets. Though performance improvement is observed, it usually involves further complications. For example, instance segmentation becomes less useful when the video encounters a new object category which is not present in instance segmentation model. Optical flow [15] suffers from occlusions which would mislead label propagation.

Most relevant works. Space-time memory network (STM) [33] is a significant work and most similar with ours. Ours is developed independently from STM, while STM is published earlier than ours. The insight that exploiting dense long-term information is similar. However, the transductive framework in the proposed approach, which stems from the classical semi-supervised learning, brings theoretical foundations to video object segmentation. Moreover, in the implementation, ours is much simpler and more efficient which does not require additional datasets, and infers all objects simultaneously.

3. Approach

In contrast to much prior work on finetuning a model on a single annotated frame or transferring knowledge from other related tasks, our approach focuses on fully exploiting the unlabeled structure in a video sequence. This enables us to build a simple model that is both strong in performance and fast in inference.

We first describe a generic semi-supervised classification

framework [49] and then adapt it to online video object segmentation in a manner that follows our ideas.

3.1. A Transductive Inference Framework

Let us first consider a general semi-supervised classification problem. Suppose that we have a dataset $\mathcal{D} = \{(x_1, y_1), (x_2, y_2), \dots, (x_l, y_l), x_{l+1}, \dots, x_n\}$, which contains l labeled data pairs and $n - l$ unlabeled data points. The task is to infer the labels $\{\hat{y}_i\}_{i=l+1}^n$ for the unlabeled data $\{x_{l+1}, \dots, x_n\}$ based on all the observation \mathcal{D} . Inference of the unlabeled data is formulated in prior work [49] as a transductive regularization framework,

$$\mathcal{Q}(\hat{\mathbf{y}}) = \sum_{i,j} w_{ij} \left\| \frac{\hat{y}_i}{\sqrt{d_i}} - \frac{\hat{y}_j}{\sqrt{d_j}} \right\|^2 + \mu \sum_{i=1}^l \|\hat{y}_i - y_i\|^2, \quad (1)$$

where w_{ij} encodes the similarity between data points (x_i, x_j) , and d_i denotes the degree $d_i = \sum_j w_{ij}$ for pixel i . The first term is a smoothness constraint that enforces similar points to have identical labels. The second term is a fitting constraint, which penalizes solutions that deviate from the initial observations. The parameter μ balances these two terms. Semi-supervised classification amounts to solve the following optimization problem,

$$\hat{\mathbf{y}} = \operatorname{argmin} \mathcal{Q}(\mathbf{y}). \quad (2)$$

It is shown [49] that the above energy minimization problem can be solved by iterative algorithm as follows. Let $\mathbf{S} = \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$ be the normalized similarity matrix constructed from w_{ij} . Iteratively solve for $\hat{\mathbf{y}}(k)$ until

convergence, as ¹

$$\hat{\mathbf{y}}(k+1) = \alpha \mathbf{S} \hat{\mathbf{y}}(k) + (1-\alpha) \mathbf{y}(0), \quad (3)$$

where $\alpha = \mu/(\mu+1)$, and $\mathbf{y}(0) = [y_1, y_2, \dots, y_n]^T$ is the initial observation of the label clamped with supervised labels. The typical value of α is 0.99. The power of this transduction model comes from the globalized model it builds over the dense structures in the unlabeled data.

3.2. Online Video Object Segmentation

Based on this general framework [49], we build a transductive model for semi-supervised video object segmentation that accounts for dense long-range interactions.

This gives rise to three challenges. First, video frames stream sequentially, so the model must work in an online fashion, where the inference of one frame should not depend on future frames. Second, the number of pixels in one video can scale into the tens of millions. A similarity matrix over all the pixels would thus be intractable to compute. Third, an effective similarity measure W needs to be learned between pixels in a video sequence.

For the algorithm to run online, it is assumed that the predictions on all prior frames have been determined when the current frame t arrives. We therefore approximate the Eqn 3 by expanding the inference procedure through time,

$$\hat{\mathbf{y}}(t+1) = \mathbf{S}_{1:t \rightarrow t+1} \hat{\mathbf{y}}(t). \quad (4)$$

$\mathbf{S}_{1:t \rightarrow t+1}$ represents the similarity matrix \mathbf{S} that is only constructed between pixels up to the t -th frame and the pixels in the $t+1$ -th frame. Since no labels are provided beyond the first frame, the prior term $\mathbf{y}(0)$ is omitted for the frame $t+1$.

For time $t+1$, the above propagation procedure is equivalently minimizing a set of smoothness terms in the spatio-temporal volume,

$$\mathcal{Q}^{t+1}(\hat{\mathbf{y}}) = \sum_i \sum_j w_{ij} \left\| \frac{\hat{y}_i}{\sqrt{d_i}} - \frac{\hat{y}_j}{\sqrt{d_j}} \right\|^2, \quad (5)$$

where i indexes the pixels at the target time $t+1$, j indexes the pixels in all frames prior to and including time t .

3.3. Label Propagation

Given the annotations on the starting frame of a video, we process the remaining frames sequentially, propagating labels to each frame based on Eqn. 4. The quality of video object segmentation heavily depends on the similarity metric \mathbf{S} , whose core component is the the affinity matrix \mathbf{W} .

¹Note that there exists a closed-form solution to the Eqn 2 shown in [49]. However, this requires the inverse of matrix S which is often computationally demanding when S is a large matrix.



Figure 3: Sampling strategy for label propagation. We sample densely in the recent history, and more sparsely in the distant history.

Similarity metric. In order to build a smooth classification function, the similarity metric should account for global high-level semantics and local low-level spatial continuity. Our similarity measure w_{ij} includes an appearance term and a spatial term,

$$w_{ij} = \exp(f_i^T f_j) \cdot \exp\left(-\frac{\|\text{loc}(i) - \text{loc}(j)\|^2}{\sigma^2}\right), \quad (6)$$

where f_i, f_j are the feature embeddings for pixels p_i, p_j through a convolutional neural network. $\text{loc}(i)$ is the spatial location of pixel i . The spatial term is controlled by a locality parameter σ . Learning of the appearance model is described in the next section.

Frame sampling. Computing a similarity matrix \mathbf{S} over all the previous frames is computationally infeasible, as long videos can span hundreds of frames or more. Inspired by Temporal Segment Networks [43], we sample a small number of frames in the observance of the temporal redundancy in videos.

Specifically, as in Figure 3, we sample a total of 9 frames from the preceding 40 frames: the 4 consecutive frames before the target frame to model the local motion, and 5 more frames sparsely sampled from the remaining 36 frames to model long-term interactions. We find this sampling strategy to strike a good balance between efficiency and effectiveness. Detailed ablations about the choice of frame sampling are presented in the experiments.

A simple motion prior. Pixels that are more distant in the temporal domain have weaker spatial dependencies. To integrate this knowledge, we use a simple motion prior where a smaller $\sigma = 8$ is used when the temporal references are sampled locally and densely, and a larger $\sigma = 21$ is employed when the reference frames are distant. We find this simple motion model to be effective for finding long-term dependencies.

Methods	Architecture	Optical	Proposal	Tracking	Re-ID
DyeNet [29]	ResNet 101	✓	✓	✓	✓
CNN-MRF [3]	Deeplab	✓	✗	✗	✗
PRemVOS [30]	Deeplab-V3+	✓	✓	✓	✓
FEELVOS [41]	Deeplab-V3+	✗	✗	✗	✗
STM [33]	2×ResNet-50	✗	✗	✗	✗
TVOS (ours)	ResNet-50	✗	✗	✗	✗

Table 1: A brief overview of leading VOS methods with dependent modules for other related vision tasks.

3.4. Learning the appearance embedding

We learn the appearance embedding in a data-driven fashion using a 2D convolutional neural network. The embedding aims to capture both short-term and long-term variations due to motion, scale and deformations. The embedding is learned from the training data in which each frame from the video is annotated with the segmented object and the object identity.

Given a target pixel x_i and we consider all pixels in the prior frames as references. Denote f_i and f_j the feature embeddings for pixel x_i and a reference pixel x_j . Then the predicted label \hat{y}_i of x_i is given by

$$\hat{y}_i = \sum_j \frac{\exp(f_i^T f_j)}{\sum_k \exp(f_i^T f_k)} \cdot y_j, \quad (7)$$

where the reference indexes j, k span the temporal history before the current frame. We show detailed ablations on how sampling the historical frames affects the learning quality.

We optimize the embedding via a standard cross-entropy loss on all pixels in the target frame,

$$\mathcal{L} = - \sum_i \log P(\hat{y}_i = y_i | x_i). \quad (8)$$

3.5. Implementation Details

We use a ResNet-50 to train the embedding model. The convolution stride of the third and the fourth residual blocks is set to 1 to maintain a high-resolution output. We add one additional 1×1 convolutional layer to project the feature to a final embedding of 256 dimensions. The embedding model produces a feature with a total stride of 8.

During training, we take the pretrained weights from the ImageNet model zoo, and finetune the model on the Davis 2017 [35] training set for 240 epochs and Youtube-VOS [46] for 30 epochs. We apply the standard augmentations of random flipping and random cropping of size 256×256 on the input images. We use a SGD solver with an initial learning rate of 0.02 and a cosine annealing scheduler. The optimization takes 16 hours on 4 Tesla P100 GPUs, with a batch size of 16, each containing 10 snippets from a video sequence.

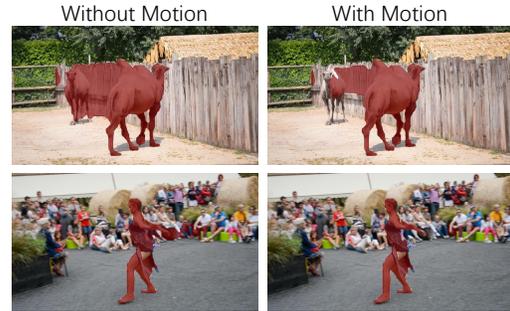


Figure 4: The effect of our simple motion model. Distant frames have weaker spatial priors on the location of objects, thus reducing the drifting problem.

During tracking, we extract features at the original image resolution of 480p. The results of each video frame are predicted sequentially online.

4. Results

In this section, we first describe our experimental settings and datasets. Then we show detailed ablations on how the transductive approach takes advantage of unlabeled structures in the temporal sequence to significantly improve the performance. Results are conducted on various datasets to compare with the state-of-the-art. Finally, we discuss temporal stability and the relationship to optical flow. Our method is abbreviated as TVOS in the result tables.

4.1. Experimental Setup

Datasets. We evaluate our method on the Davis 2017 [35], and Youtube-VOS [46] datasets. Our model is trained on the respective training set and evaluated on the validation set. For Davis 2017, we also train our model on the combined train-val set, and submit the results on the testing set on the evaluation server.

Davis 2017 contains 150 video sequences and it involves multiple objects with drastic deformations, heavy and prolonged occlusions, and very fast motions. High-definition annotations are available for all frames in the training sequences. Youtube-VOS is the largest dataset for this task to-date, containing 4453 training sequences and 474 validation sequences. It captures a comprehensive collection of 94 daily object categories. However, the frame rate of the video is much lower than videos in Davis (5 fps compared to 24 fps).

Evaluation metrics. We use the standard evaluation metric of mean intersection over union (mIoU), averaged across objects and summed over all frames. The mIoU is evaluated on both the full objects (J measure) and only on the object boundaries (F measure). The global metric (G measure) is the average of the J and F measures. Youtube-



Figure 5: Using dense long-range dependencies improves the tracking performance. The spatial term smooths the object boundaries, while long-term dependencies up to 40 frames help to re-detect objects.

train / tracking	1 frame	3 frames	9 frames	uniform sample	sparse sample	sparse + motion
1 frame	55.8	60.4	63.4	63.8	64.0	64.3
3 frames	56.0	61.4	65.4	65.5	66.1	66.7
9 frames	60.7	63.4	68.6	68.6	69.0	69.9
uniform sample	55.8	60.2	64.4	65.0	65.1	65.3
sparse sample	59.9	62.9	66.2	67.2	68.5	68.6
Supervised	47.5	52.2	53.8	54.0	54.5	54.8
InstDisc[44]	42.4	47.3	51.3	51.3	52.1	52.2
MoCo[18]	43.5	48.7	53.0	53.2	53.8	54.0

Table 2: Ablation study on the range of temporal dependencies and the simple motion component. The mean J measure on the Davis 2017 validation set is reported. See text for details.

VOS also includes a separate measure of seen objects and unseen objects to measure the generalization ability. In Section 4.4, we provide a discussion of temporal stability.

4.2. Ablation Study

Dense local and global dependencies. While most prior works focus on optimizing single frame models, the key idea of this paper is to build dense long-term models over the spatio-temporal volume. In Table 2, we summarize the effect of such long-term potentials which capture both local and global dependency. Each row is an appearance embedding model trained with different reference frame sampling strategies. Each column corresponds to a tracking sampling strategy. We study the following settings: one reference frame preceding the target frame, 3

consecutive frames preceding the target frame, 9 consecutive frames preceding the target frame, uniform sampling of 9 frames in the preceding 40 frames, and our sparse sampling of 9 frames in the preceding 40 frames as in Figure 3. We find that tracking over a longer term generally improves the performance, and denser sampling near the target frame is helpful. For learning the appearance embedding, training with 9 consecutive frames produces the best results, while longer ranges do not always lead to improvements. This may be due to very long ranges covering almost the entire video reduces the variations in the dataset, which leads to worse generalization for training.

In Figure 5, we show some qualitative examples for long range tracking. Using 9 consecutive frames yields more stable predictions than using only the previous frame. Adding the spatial term smooths the object boundaries. A long range of 40 frames enables the model to re-detect objects after heavy occlusions.

Transferred representations. In the last rows of Table 2, we also test the tracking performance for models pretrained on ImageNet but without further training on the DAVIS dataset. The transferred ImageNet model obtains a mean J measure of 54.8%, which is actually better than some prior methods [47, 10] trained with additional Davis data. Also, even an unsupervised pretrained model on images obtains performance competitive to network modulation [47] using our transductive inference algorithm. Two recent unsupervised pretrained models on ImageNet are investigated [44, 18]. Since no domain-specific training is involved for the appearance embedding, the evaluation of transferred representations clearly validates the effective-

Methods	FT	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	Speed
OnAVOS [42]	✓	61.0	66.1	63.6	0.08
DyeNet [29]	✓	67.3	71.0	69.1	0.43
CNN-MRF [3]	✓	67.2	74.2	70.7	0.03
PReMVOS [30]	✓	73.9	81.7	77.8	0.03
Modulation [47]	✗	52.5	57.1	54.8	3.57
FAVOS [10]	✗	54.6	61.8	58.2	0.83
VideoMatch [22]	✗	56.5	68.2	62.4	2.86
RGMP [45]	✗	64.8	68.8	66.7	3.57
FEELVOS [41]	✗	65.9	72.3	69.1	1.96
STM [33]	✗	69.2	74.0	71.6	6.25
STM [33]+Pretrain	✗	81.7	79.2	84.3	6.25
TVOS	✗	69.9	74.7	72.3	37

Table 3: Quantitative evaluation on the Davis 2017 validation set. FT denotes methods that perform online training.

ness of dense long-term modeling.

The simple motion prior. As a weak spatial prior for modeling the dependency between distant frames, our simple motion model reduces noise from the model predictions and leads to about 1% improvement. Figure 4 displays two concrete examples. More complicated motion models, such as a linear motion model [1], may be even more effective.

4.3. Quantitative Results

In Table 1, we first give a brief overview of the current leading methods, including those that use first-frame finetuning (CNN-MRF [3], DyeNet [29], PReMVOS [30]) and those that do not (FEELVOS [41], STM [33] and our TVOS). For DyeNet and PReMVOS, their sub-modules are learned on dedicated datasets such as optical flow on Flying Chairs, object proposal on MSCOCO, and object segmentation on PASCAL VOC. Since Davis is much smaller than the large-scale datasets, it remains unknown how much of the gains can be attributed to knowledge transfer or to the methods themselves. Therefore, the mentioned methods are not directly comparable with our method. FEELVOS, STM and ours are much simpler, as they do not rely on additional modules for this problem. STM additionally requires heavy pretraining on large-scale image datasets.

It is also important to note that for PReMVOS, DyeNet, CNN-MRF, they are not able to run tracking **in an online fashion**. They use information from future frames to stabilize prediction for the target frame. Also, instead of using the first frame from the given video for training, they use the first frames from the entire test-dev set for training. Propagation-based methods are able to track objects sequentially online.

DAVIS 2017. We summarize our results on the Davis 2017 validation set in Table 3, and on the Davis 2017 test-dev set in Table 4. On the validation set, our method performs slightly better than STM [33] under the same amount

Methods	FT	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	Speed
OnAVOS [42]	✓	53.4	59.6	56.5	0.08
DyeNet [29]	✓	65.8	70.5	68.2	0.43
CNN-MRF [3]	✓	64.5	70.5	67.5	0.02
PReMVOS [30]	✓	67.5	75.7	71.6	0.02
RGMP [45]	✗	51.4	54.4	52.9	2.38
FEELVOS [41]	✗	51.2	57.5	54.4	1.96
TVOS	✗	58.8	67.4	63.1	37

Table 4: Quantitative evaluation on the Davis 2017 test-dev set. FT denotes methods that perform online training.

Methods	Overall	Seen		Unseen	
		\mathcal{J}	\mathcal{F}	\mathcal{J}	\mathcal{F}
RGMP [45]	53.8	59.5	-	45.2	-
OnAVOS [42]	55.2	60.1	62.7	46.6	51.4
RVOS [40]	56.8	63.6	67.2	45.5	51.0
OSVOS [7]	58.8	59.8	60.5	54.2	60.7
S2S [46]	64.4	71.0	70.0	55.5	61.2
PReMVOS [30]	66.9	71.4	75.9	56.5	63.7
STM [33]+Pretrain	79.4	79.7	84.2	72.8	80.9
TVOS	67.8	67.1	69.4	63.0	71.6
TVOS (from DAVIS)	67.4	66.7	69.8	62.5	70.6

Table 5: Quantitative evaluation on the Youtube-VOS validation set.

of training data, while surpassing other propagation-based methods which do not need fine-tuning, by 4% for mean \mathcal{J} and 3% for mean $\mathcal{J}\&\mathcal{F}$. In comparison to finetuning based methods, our TVOS also outperforms DyeNet and CNN-MRF by 2% while being significantly simpler and faster.

We train our model on the combined training and validation set for evaluating on the test-dev set. We find that there is a large gap of distribution between the Davis 2017 test-dev and validation sets. Heavy and prolonged occlusions among objects belonging to the same category are more frequent in the test-dev, which favors methods with re-identification modules. As a result, we are 4 – 5% lower than DyeNet and CNN-MRF on the test-dev set. FEELVOS is even more negatively affected, performing 8% lower than ours in terms of mean $\mathcal{J}\&\mathcal{F}$. STM [33] does not provide an evaluation on the test set.

Youtube-VOS. We summarize the results on youtube-VOS validation set in Table 5. Ours surpasses all prior works except STM [33], which relies on heavy pretraining on a variety of segmentation datasets such as saliency detection and instance segmentation. Without pretraining, STM obtains a comparable result of 68.1%. We also test the generalization ability of our model trained on DAVIS train-val and test on Youtube-VOS val. The transferred model shows great generalization ability with an overall score of 67.4%

Speed analysis. During tracking, we cache the appearance embeddings for a history up to 40 frames. Inference



Figure 6: Per-frame IoU over time for PreMVOS and our method on an example video sequences from the DAVIS validation set. PreMVOS switches object identities frequently, while our predictions are temporally smooth. The color of each IoU curve matches its corresponding object segment.

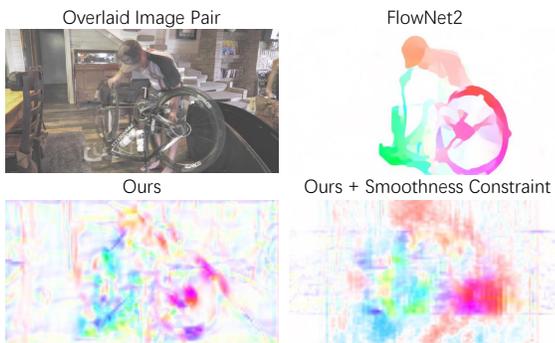


Figure 7: An example of optical flow computed from our model compared to FlowNet2.

per frame thus only involves a feed-forward pass of the target frame through the base network, and an additional dot product of target embeddings to prior embeddings. Computation is also constant of any number of objects. This makes our algorithm extremely fast, with a runtime of 37 frames per second on a single Titan Xp GPU. Figure 1 compares current algorithms on their trade-off between speed and performance. Ours is an order of magnitude faster than prior methods, while achieving the results comparable to state-of-the-art methods.

4.4. Discussions

Temporal stability. Temporal stability is often a desirable property in video object segmentation, as sharp inconsistencies may be disruptive to downstream video analysis. However, temporal stability is typically not included as an evaluation criterion. Here, we give qualitative examples showing the difference in temporal stability between our model and the state-of-the-art PreMVOS [30].

In Figure 6, we show examples of per-frame evaluation along video sequences. Although the state-of-the-art integrates various temporal smoothing modules, such as opti-

cal flow, merging and tracking, we observe the detection-based method to be prone to noise. For example, objects are lost suddenly, or being tagged with a different identity. Our method, on the other hand, makes temporally consistent predictions.

Does our model learn optical flow? Our method learns a soft mechanism for associating pixels in the target frame with pixels in the history frames. This is similar to optical flow where hard correspondences are computed between pixels. We examine how much our learned model aligns with optical flow.

We take two adjacent two frames and calculate the optical flow from our model as $\Delta d_i = \sum_j s_{ij} \Delta d_{ij}$, where s_{ij} is the normalized similarity, and Δd_{ij} is the displacement between i, j . Figure 7 shows an example visualization of the flow. Compared to the optical flow computed by FlowNet2 [23], our flow makes sense on objects that would be segmented, but is much more noisy on the background. We have further added a spatial smoothness constraint on the computed optical flow for jointly learning the embeddings, as widely used for optical flow estimation [15, 24]. We observe that the constraint smooths the optical flow on the background, but fails to regularize the model for tracking. Adding the term consistently hurts the performance of video object segmentation.

5. Conclusion

We present a simple approach to semi-supervised video object segmentation. Our main insight is that much more unlabeled structure in the spatio-temporal volume can be exploited for video object segmentation. Our model finds such structure via transductive inference. The approach is learned end-to-end, without the need of additional modules, additional datasets, or dedicated architectural designs. Our vanilla ResNet50 model achieves competitive performance with a compelling speed of 37 frames per second. We hope our model can serve as a solid baseline for future research.

References

- [1] A. Andriyenko and K. Schindler. Multi-target tracking by continuous energy minimization. In *CVPR 2011*, pages 1265–1272. IEEE, 2011. [7](#)
- [2] X. Bai, J. Wang, D. Simons, and G. Sapiro. Video snapshot: robust video object cutout using localized classifiers. In *ACM Transactions on Graphics (ToG)*, volume 28, page 70. ACM, 2009. [2](#)
- [3] L. Bao, B. Wu, and W. Liu. Cnn in mrf: Video object segmentation via inference in a cnn-based higher-order spatio-temporal mrf. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5977–5986, 2018. [2](#), [3](#), [5](#), [7](#)
- [4] M. Belkin and P. Niyogi. Semi-supervised learning on manifolds. *Machine Learning Journal*, 2002. [1](#)
- [5] K. Bernardin and R. Stiefelwagen. Evaluating multiple object tracking performance: the clear mot metrics. *Journal on Image and Video Processing*, 2008:1, 2008. [1](#)
- [6] A. Blum and S. Chawla. Learning from labeled and unlabeled data using graph mincuts. 2001. [1](#)
- [7] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-shot video object segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5320–5329. IEEE, 2017. [2](#), [7](#)
- [8] S. Chandra, C. Couprie, and I. Kokkinos. Deep spatio-temporal random fields for efficient video segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8915–8924, 2018. [2](#)
- [9] Y. Chen, J. Pont-Tuset, A. Montes, and L. Van Gool. Blazingly fast video object segmentation with pixel-wise metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1189–1198, 2018. [2](#)
- [10] J. Cheng, Y.-H. Tsai, W.-C. Hung, S. Wang, and M.-H. Yang. Fast and accurate online video object segmentation via tracking parts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7415–7424, 2018. [6](#), [7](#)
- [11] A. Faktor and M. Irani. Video segmentation by non-local consensus voting. In *BMVC*, volume 2, page 8, 2014. [2](#)
- [12] Q. Fan, F. Zhong, D. Lischinski, D. Cohen-Or, and B. Chen. Jumpcut: non-successive mask transfer and interpolation for video cutout. *ACM Trans. Graph.*, 34(6):195–1, 2015. [2](#)
- [13] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2360–2367. IEEE, 2010. [1](#)
- [14] R. Fergus, Y. Weiss, and A. Torralba. Semi-supervised learning in gigantic image collections. In *Advances in neural information processing systems*, pages 522–530, 2009. [1](#)
- [15] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. Van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. *arXiv preprint arXiv:1504.06852*, 2015. [1](#), [3](#), [8](#)
- [16] H. Grabner, C. Leistner, and H. Bischof. Semi-supervised on-line boosting for robust tracking. In *European conference on computer vision*, pages 234–247. Springer, 2008. [2](#)
- [17] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph-based video segmentation. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 2141–2148. IEEE, 2010. [2](#)
- [18] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019. [6](#)
- [19] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [1](#)
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#)
- [21] Y.-T. Hu, J.-B. Huang, and A. Schwing. Maskrnn: Instance level video object segmentation. In *Advances in Neural Information Processing Systems*, pages 325–334, 2017. [2](#), [3](#)
- [22] Y.-T. Hu, J.-B. Huang, and A. G. Schwing. Videomatch: Matching based video object segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 54–70, 2018. [2](#), [7](#)
- [23] 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 *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, pages 1647–1655. IEEE, 2017. [8](#)
- [24] J. Janai, F. Güney, A. Ranjan, M. Black, and A. Geiger. Unsupervised learning of multi-frame optical flow with occlusions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 690–706, 2018. [8](#)
- [25] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele. Lucid data dreaming for video object segmentation, 2018. [2](#)
- [26] Y. J. Lee, J. Kim, and K. Grauman. Key-segments for video object segmentation. In *2011 International conference on computer vision*, pages 1995–2002. IEEE, 2011. [2](#)
- [27] S. Li, B. Seybold, A. Vorobyov, A. Fathi, Q. Huang, and C.-C. Jay Kuo. Instance embedding transfer to unsupervised video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6526–6535, 2018. [2](#)
- [28] S. Li, B. Seybold, A. Vorobyov, X. Lei, and C.-C. Jay Kuo. Unsupervised video object segmentation with motion-based bilateral networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 207–223, 2018. [2](#)
- [29] X. Li and C. Change Loy. Video object segmentation with joint re-identification and attention-aware mask propagation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 90–105, 2018. [1](#), [2](#), [3](#), [5](#), [7](#)
- [30] J. Luiten, P. Voigtlaender, and B. Leibe. Premvos: Proposal-generation, refinement and merging for video object segmentation. *arXiv preprint arXiv:1807.09190*, 2018. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#)
- [31] 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. *arXiv preprint arXiv:1709.06031*, 2017. [2](#)
- [32] N. Märki, F. Perazzi, O. Wang, and A. Sorkine-Hornung. Bilateral space video segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*

- tion, pages 743–751, 2016. [2](#)
- [33] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim. Video object segmentation using space-time memory networks. *arXiv preprint arXiv:1904.00607*, 2019. [3](#), [5](#), [7](#)
- [34] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2663–2672, 2017. [2](#), [3](#)
- [35] 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. [5](#)
- [36] J. Shin Yoon, F. Rameau, J. Kim, S. Lee, S. Shin, and I. So Kweon. Pixel-level matching for video object segmentation using convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2167–2176, 2017. [2](#)
- [37] M. Szummer and T. Jaakkola. Partially labeled classification with markov random walks. In *Advances in neural information processing systems*, pages 945–952, 2002. [1](#)
- [38] D. Tsai, M. Flagg, A. Nakazawa, and J. M. Rehg. Motion coherent tracking using multi-label mrf optimization. *International journal of computer vision*, 100(2):190–202, 2012. [2](#)
- [39] Y.-H. Tsai, M.-H. Yang, and M. J. Black. Video segmentation via object flow. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3899–3908, 2016. [2](#)
- [40] C. Ventura, M. Bellver, A. Girbau, A. Salvador, F. Marques, and X. Giro-i Nieto. Rvos: End-to-end recurrent network for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5277–5286, 2019. [7](#)
- [41] 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. *arXiv preprint arXiv:1902.09513*, 2019. [2](#), [5](#), [7](#)
- [42] P. Voigtlaender and B. Leibe. Online adaptation of convolutional neural networks for video object segmentation. *arXiv preprint arXiv:1706.09364*, 2017. [2](#), [7](#)
- [43] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, pages 20–36. Springer, 2016. [4](#)
- [44] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018. [6](#)
- [45] S. Wug Oh, J.-Y. Lee, K. Sunkavalli, and S. Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7376–7385, 2018. [2](#), [7](#)
- [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 preprint arXiv:1809.03327*, 2018. [5](#), [7](#)
- [47] L. Yang, Y. Wang, X. Xiong, J. Yang, and A. K. Katsaggelos. Efficient video object segmentation via network modulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6499–6507, 2018. [6](#), [7](#)
- [48] D. Zhang, O. Javed, and M. Shah. Video object segmentation through spatially accurate and temporally dense extraction of primary object regions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 628–635, 2013. [2](#)
- [49] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. In *Advances in neural information processing systems*, pages 321–328, 2004. [1](#), [2](#), [3](#), [4](#)