

Efficient Video Object Segmentation via Network Modulation

Linjie Yang¹Yanran Wang^{2*}
¹Snap Inc.Xuehan Xiong³
²Northwestern UniversityJianchao Yang¹
³Google Inc.Aggelos K. Katsaggelos²

{linjie.yang, jianchao.yang}@snap.com {joycewang1026@u,aggk@eecs}.northwestern.edu xiong828@gmail.com

Abstract

Video object segmentation targets segmenting a specific object throughout a video sequence when given only an annotated first frame. Recent deep learning based approaches find it effective to fine-tune a general-purpose segmentation model on the annotated frame using hundreds of iterations of gradient descent. Despite the high accuracy that these methods achieve, the fine-tuning process is inefficient and fails to meet the requirements of real world applications. We propose a novel approach that uses a single forward pass to adapt the segmentation model to the appearance of a specific object. Specifically, a second meta neural network named modulator is trained to manipulate the intermediate layers of the segmentation network given limited visual and spatial information of the target object. The experiments show that our approach is 70× faster than fine-tuning approaches and achieves similar accuracy. Our model and code have been released at https://github.com/linjieyangsc/video_seg.

1. Introduction

Semantic segmentation plays an important role in understanding the visual content of an image as it assigns predefined objects or scene labels to each pixel and thus translates the image into a segmentation map. When dealing with video content, a human can easily segment an object in the whole video without knowing its semantic meaning, which inspired a research topic named semi-supervised video segmentation. In a typical scenario of semi-supervised video segmentation, one is given the first frame of a video along with an annotated object mask, and the task is to accurately locate the object in all following frames [24, 20]. The ability of performing accurate pixel-level video segmentation with minimum supervision (e.g., one annotated frame) can lead to a large amount of applications, such as accurate object tracking for video understanding, interactive video editing, augmented reality, and video-based advertisements. When

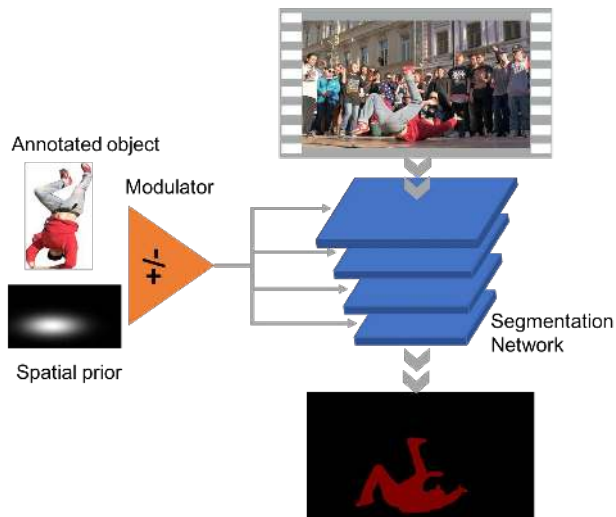


Figure 1. An overview of our approach. Our model is consisted of a modulator and a segmentation network. The modulator can adapt the segmentation model instantly to segment an arbitrary object through a video sequence.

the supervision is limited to only one annotated frame, researchers refer to this scenario as *one-shot learning*. In recent years, we have witnessed a rising amount of interest in developing one-shot learning techniques for video segmentation [2, 23, 35, 22, 32, 4]. Most of these works share a similar two-stage paradigm: first, train a general-purpose Fully Convolutional Network (FCN) [31] to segment the foreground object; second, fine-tune this network based on the first frame of the video for several hundred forward-backward iterations to adapt the model to the specific video sequence. Despite the high accuracies achieved by these approaches, the fine-tuning process is arguably time consuming, which hinders real-time applications. Some of these approaches [4] [23] also utilize optical flow information, which is computationally heavy for state-of-the-art algorithms [29] [15].

In order to reduce the computational cost of semi-supervised segmentation, we propose a novel approach to adapt the generic segmentation network to the appearance of a specific object instance in one single feed-forward pass.

*This work is done when Yanran was an intern at Snap Inc.

We propose to employ another meta neural network called *modulator* to learn to adjust the intermediate layers of the generic segmentation network given an arbitrary target object instance. Fig. 1 shows an illustration of our approach. By extracting information from the image of the annotated object and the spatial prior of the object, the modulator produces a list of parameters, which are injected into the segmentation model for layer-wise feature manipulation. Without one-shot fine-tuning, our model is able to change the behavior of the segmentation network with minimum extracted information from the target object. We name this process *network modulation*.

Our proposed model is efficient, requiring only one forward pass from the modulator to produce all parameters needed for the segmentation model to adapt to the specific object instance. Network modulation guided by the spatial prior facilitates the model to track the object even with the presence of multiple similar instances. The whole pipeline is differentiable and can be learned end-to-end using the standard stochastic gradient descent. The experiments show that our approach outperforms previous approaches without one-shot fine-tuning by a large margin, and achieves comparable performance with these approaches after one-shot fine-tuning with a $70\times$ speed up.

2. Related Work

Semi-supervised video segmentation. Semi-supervised video object segmentation aims at tracking an object mask from the first annotated frame throughout the rest of video. Many approaches have been proposed in recent literature, including those propagating superpixels [17] [35], patches [9], object proposals [25], or in bilateral space [22], and graphical model based optimization is usually performed to consider multiple frames simultaneously. With the success of FCN on static image segmentation [12], deep learning based methods [23, 2, 32, 34, 18, 4] have been recently proposed for video segmentation and promising results have been achieved. To model the temporal motion information, some works heavily rely on optical flow [34] [4], and use CNNs to learn mask refinement of an object from current frame to the next one [23], or combine the training of CNN with bilateral filtering between adjacent frames [18]. Chen *et al.* [4] use a CNN to jointly estimate the optical flow and provide the learned motion representation to generate motion consistent segmentation across time. Different from these approaches, Caelles *et al.* [2] combine offline and online training process on static images without using temporal information. While it saves the computation of optical flow and/or conditional random fields (CRF) [19] involved in some previous methods, online fine-tuning still requires many iterations of optimization, which poses a challenge for real-world applications that need rapid inference.

Meta-learning for low-shot learning. Current success of deep learning relies on the ability to learn from large-scale labeled datasets through gradient descent optimization. However, it is not optimal for the model to learn each task for each setting from scratch if we aim for the learning of many tasks adapted to many environments. Instead, we want our deep learning system to be able to learn new tasks very fast and from very limited quantities of data. In the extreme of “one-shot learning”, the algorithm needs to learn the new task with a single observation. One potential strategy for learning a versatile model is the notion of meta-learning, or learning to learn, which dates back to the late 1980s. Recently, meta-learning has become a hot research topic with publications on neural network optimization [3], finding good network architectures, fast reinforcement learning, and few-shot image recognition [36, 28, 13, 10, 30]. Ravi and Larochelle [28] proposed a LSTM meta-learner to learn the updating rules for few shot learning. The meta optimization over a large number of tasks in [10] targets learning a model that can quickly adapt to the new task with a limited number of updates. Hariharan and Girschick [13] trained a learner that generated new samples and used new samples for training new tasks. Our approach is similar to meta-learning because it learns to update the segmentation model rapidly with another meta learner, i.e. the modulator.

Network manipulation Several previous works try to incorporate modules to manipulate the behavior of a deep neural network, either to manipulate spatial arrangement of data [16] or connections of filter weights [5]. Our method is also heavily motivated by conditional batch normalization [8, 11, 14, 26], where the behavior of the deep model is manipulated by batch normalization parameters conditioned on a guidance input, e.g. a style image for image stylization or a language sentence for visual question answering.

3. Video Object Segmentation with Network Modulation

In our proposed framework, we utilize modulators to instantly adapt the segmentation network to a specific object, rather than performing hundreds of iterations of gradient descent. We can achieve similar accuracy by adjusting a limited number of parameters in the segmentation network, compared with the updating the whole network in one-shot learning approaches [23, 2]. There are two important cues for video object segmentation: visual appearance and continuous motion in space. To use information from both visual and spatial domains, we incorporate two network modulators, namely *visual modulator* and *spatial modulator*, to learn to adjust intermediate layers in the main segmentation

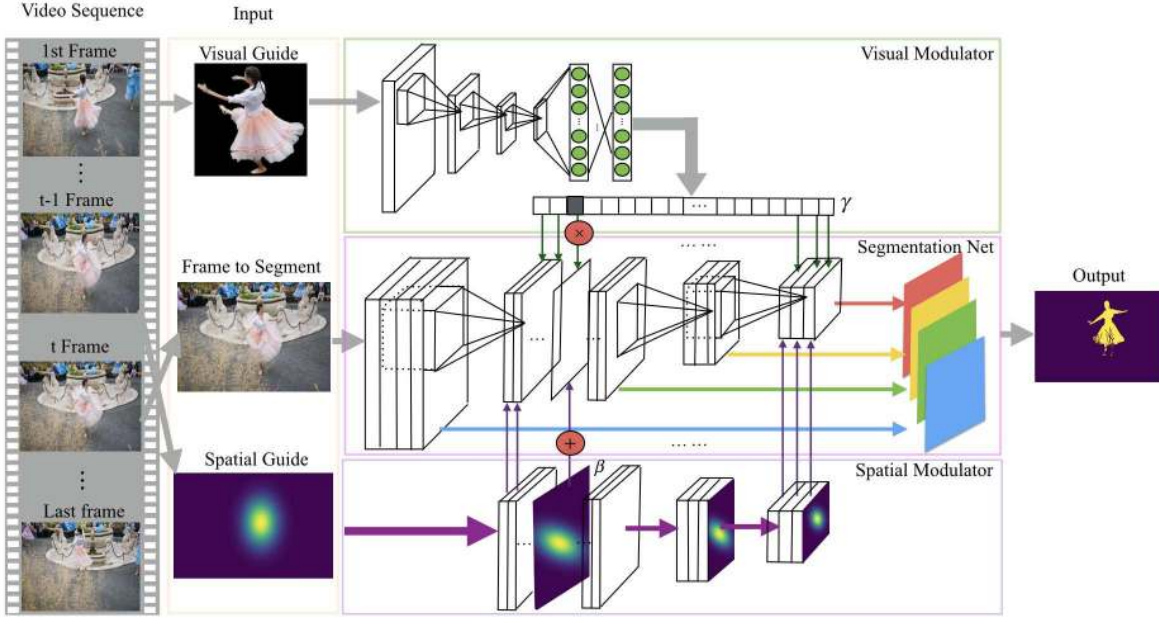


Figure 2. An illustration of our model with three components: a segmentation network, a visual modulator, and a spatial modulator. The two modulators produce a set of parameters that manipulates the intermediate feature maps of the segmentation network and adapts it to segment the specific object.

network, based on the annotated first frame and spatial location of the object, respectively.

3.1. Conditional batch normalization

Our approach is inspired by recent works using Conditional Batch Normalization (CBN) [6, 14, 26], where the scale and bias parameters of each batch-normalization layer are produced by a second controller network. These parameters are used to control the behavior of the main network for tasks such as image stylization and question answering. Mathematically, each CBN layer can be formulated as follows:

$$\mathbf{y}_c = \gamma_c \mathbf{x}_c + \beta_c, \quad (1)$$

where \mathbf{x}_c and \mathbf{y}_c are the input and output feature maps in the c_{th} channel, and γ_c and β_c are the scale and bias parameters produced by the controller network, respectively. The mean and variance parameters are omitted for clarity.

3.2. Visual and spatial modulation

The CBN layer is a special case of the more general scale-and-shift operation on feature maps. Following each convolution layer, we define a new *modulation layer* with parameters generated by both visual and spatial modulators that are jointly trained. We design the two modulators such that the visual modulator produces channel-wise scale parameters to adjust the weights of different channels in the feature maps, while the spatial modulator generates element-wise bias parameters to inject spatial prior to the

modulated features. Specifically, our modulation layer can be formulated as follows:

$$\mathbf{y}_c = \gamma_c \mathbf{x}_c + \beta_c, \quad (2)$$

where γ_c and β_c are modulation parameters from the visual and spatial modulators, respectively. γ_c is a scalar for channel-wise weighting, while β_c is a two-dimensional matrix to apply point-wise bias values.

Fig. 2 shows an illustration of the proposed approach, which consists of three networks: a fully-convolutional main segmentation network, a visual modulator network, and a spatial modulator network. The visual modulator network is a CNN that takes the annotated visual object image as input and produces a vector of scale parameters for all modulation layers, while the spatial modulator network is a very efficient network that produces bias parameters based on the spatial prior input. We will discuss the two modulators in more detail in the following sections.

3.3. Visual modulator

The visual modulator is used to adapt the segmentation network to focus on a specific object instance, which is the annotated object in the first frame. The annotated object is referred to as *visual guide* hereafter for convenience. The visual modulator extracts semantic information such as category, color, shape, and texture from the visual guide and generates corresponding channel-wise weights in order to re-target the segmentation network to segment the object. We use VGG16 [33] neural network as the model for the

visual modulator. We modify its last layer trained for ImageNet classification to match the number of parameters in the modulation layers for the segmentation network.

The visual modulator implicitly learns an embedding of different types of objects. It should produce similar parameters to adjust the segmentation network for similar objects and different parameters for different objects. It is demonstrated in Sec. 4.2 that the embedding of the modulator outputs closely correlates with object appearance. One significant advantage of using such a visual modulator is that we can potentially transfer the knowledge learned with a large number of object classes, e.g., ImageNet, in order to learn a good embedding.

3.4. Spatial modulator

Our spatial modulator takes a prior location of the object in the image as input. Since objects move continuously in a video, we set the prior as the predicted location of the object mask in the previous frame. Specifically, we encode the location information as a heatmap with a two-dimensional Gaussian distribution on the image plane. The center and standard deviations of the Gaussian distribution are computed from the predicted mask of the previous frame. This heatmap is referred to as a *spatial guide* hereafter for convenience. The spatial modulator downsamples the spatial guide into different scales, to match the resolution of different feature maps in the segmentation network, and then applies a scale-and-shift operation on each downsampled heatmap to generate the bias parameters of the corresponding modulation layer. Mathematically,

$$\beta_c = \tilde{\gamma}_c \mathbf{m} + \tilde{\beta}_c \quad (3)$$

where \mathbf{m} is a down-sampled Gaussian heat map for the corresponding modulation layer, $\tilde{\gamma}_c$ and $\tilde{\beta}_c$ are the scale-and-shift parameters for the c -th channel, respectively. This is implemented with a computationally efficient 1×1 convolution. At the bottom of Fig. 2, we illustrate the structure of the spatial modulator.

Our method shares some similarities with the previous work MaskTrack [23] in utilizing information from the previous mask. Compared to their approach that uses the exact foreground mask of the previous frame, we use only a very coarse location prior. At first, it may seem that our method discards more information from the previous frame. However, we argue that the rough position and size in the previous frame possess enough information to infer the object mask with the RGB image. This prevents the model from relying too much on the mask and as a result the error propagation, which can be catastrophic when the object has large movements in the video. One drawback of such over-utilization of the mask is that MaskTrack has to apply plenty of well-engineered data augmentation to prevent over-fitting, while we only apply simple shift and scaling.

3.5. Implementation details

Our FCN structure follows the one used by [2], which is a VGG16 [33] model with a hyper-column structure [12]. Intuitively, we should add modulation layers after each convolution layer in the FCN. However, we found that adding modulation layers in-between the early convolution layers actually worsens the model’s performance. One possible reason is that early layers extract low-level features that are very sensitive to the scale-and-shift operations introduced by the modulator. In our implementation, we add modulation operations to all convolution layers in VGG16 except the first four layers, which results in nine modulation layers.

Similar to MaskTrack [23], we also utilize static images for training our model. Ideally, the visual modulator should learn a mapping from any object to modulation weights of different layers in a FCN, which requires the model to see all possible different objects. However, most video semantic segmentation datasets only contain a very limited number of categories. We tackle this challenge by using the largest public semantic segmentation dataset MSCOCO [21], which has 80 object categories. We select objects that are larger than 3% of the image size for training, resulting in a total number of 217,516 objects. For preprocessing the input for the visual modulator, we first crop the object using the annotated mask, then set the background pixels to mean image values, and then resize the cropped image to a constant resolution of 224×224 . The object is also augmented with up to 10% random scaling and 10° random rotation. For preprocessing the spatial guide as input to the spatial modulator, we first compute the mean and standard deviation of the mask, and then augment the mask with up to 20% random shift and 40% random scaling. For the whole image fed into the FCN, we use a random size from 320, 400, and 480 with a square shape.

The visual modulator and segmentation network are both initialized with VGG16 model pretrained on the ImageNet [7] classification task. The modulation parameters $\{\gamma_c\}$ are initialized to ones by setting the weights and biases of the last fully-connected layer of the visual modulator to zeros and ones, respectively. The weights of spatial modulator are initialized randomly. We used the same balanced cross-entropy loss as in [2]. A mini-batch size of 8 is used. We use Adam optimizer with default momentum 0.9 and 0.999 for β_1 and β_2 , respectively. The model is first trained for 10 epochs with learning rate 10^{-5} and then trained for another 5 epochs with learning rate 10^{-6} .

Further, in order to model appearance variations of moving objects in videos, the model can be finetuned on video segmentation dataset such as DAVIS 2017 [27]. To be more robust to appearance variations, we randomly pick a foreground object from the whole video sequence as the visual guide for each frame. The spatial guide is obtained from the ground truth mask of the object in the previous frame.

The same data augmentations are applied as training on MSCOCO. The model is fine-tuned for 20 epochs with learning rate 10^{-6} .

4. Experiments

In this section, we will introduce three parts of experiment: the comparison of our approach to previous methods, the visualization of modulation parameters, and an ablation study. Our model is tested on several popular video segmentation datasets, including DAVIS [24] [27] and YoutubeObjects [17].

4.1. Semi-supervised Video Segmentation

In this section, we compare with traditional approaches including OFL [35], BVS[22], and deep learning-based approaches including PLM [32], MaskTrack [23], OSVOS [2], VPN [18], SFL [4], and ConvGRU [34].

4.1.1 DAVIS 2016 & YoutubeObjects

First, we compare our approach to previous approaches on DAVIS 2016 and YoutubeObjects. Some approaches (MaskTrack[23], SFL [4] and OSVOS[2]) reported results both with and without model fine-tuning on the target sequences. We include both of them and denote the variants without fine-tuning as MaskTrack-B, SFL-B, and OSVOS-B, respectively. Our model has two variants, with the first only trained on static images (Stage 1) and the second fine-tuned on DAVIS 2017 (Stage 1&2). Since there are several popular add-ons for this line of research, such as optical flow and CRF [19], which both have a lot of variants and make a fair comparison hard, we only include the performances without optical flow and CRF if possible, and mark those with add-ons in Table 1.

In Table 1, by comparing our method with OFL [35], an expensive graphical model based approach, we achieve better accuracy on both DAVIS 2016 and YoutubeObjects. Compared to deep learning approaches without model fine-tuning, and therefore, similar speed as ours, our method achieves the best accuracy on both DAVIS 2016 and YoutubeObjects. Compared to the four approaches using model fine-tuning on target videos (PLM, MaskTrack, SFL, and OSVOS), our approach achieves better performance than PLM and MaskTrack, and is on-par with SFL. OSVOS achieves higher accuracy but also utilizes a boundary snapping approach which contributes 2.4% in mean IU. Our method is $70\times$ faster than MaskTrack and OSVOS, $50\times$ faster than SFL. We measure the running time of MaskTrack-B, OSVOS-B, and our method on a NVIDIA Quadro M6000 GPU using Tensorflow [1]. The speed of other methods is derived from the corresponding papers ¹.

¹Speed of ConvGRU is estimated with the expensive optical flow they use, speed of PLM is derived through communication with the authors.

In our method, one forward pass for the visual modulator allows the segmentation model to adapt, which is more efficient than existing approaches with model fine-tuning on target videos. The visual modulator only needs to be computed once for the whole video. Meanwhile, the spatial modulator needs to be computed for every frame but the overhead is negligible, i.e., the average speed of our model on a video sequence is about the same as FCN itself. Our method is the second fastest of all compared methods, with only MaskTrack-B and OSVOS-B achieving similar speed but with lesser accuracies.

4.1.2 DAVIS 2017

To further investigate the capability of our model, we conduct more experiments on DAVIS 2017 [27], which is the largest video segmentation dataset to date. DAVIS 2017 is more challenging than DAVIS 2016 and YoutubeObjects in that it has multiple objects for each video sequence and some of the objects are very similar. We compare our method with two most related approaches, MaskTrack [23] and OSVOS [2]. For fair comparison, we only use their single network and adds-on free versions. We directly use open source code of OSVOS and adapt MaskTrack model to Tensorflow [1]. For each video sequence, OSVOS and MaskTrack are fine-tuned with 1000 iterations. To show that the network modulation is capable of adapting different model structures to specific object instances, we also experiment with modified OSVOS and MaskTrack models by adding a visual modulator to each of them, which are named OSVOS-M and MaskTrack-M respectively. For these two models, we only update the weights of the visual modulators and keep the weights of the segmentation model fixed in training.

Table 2 shows the results of different approaches on DAVIS 2017. We utilize the official evaluation metrics of DAVIS dataset: mean, recall, and decay of region similarity \mathcal{J} and contour accuracy \mathcal{F} , respectively. Note \mathcal{J} mean is equivalent to mean IU we used above. Again, our model outperforms OSVOS-B and MaskTrack-B by a large margin, while obtaining comparable performance with the two methods with model fine-tuning. OSVOS-M and MaskTrack-M are both better than their baseline implementations with a 18% and 9.3% gain in \mathcal{J} mean, respectively. Since the weights of the segmentation model are fixed, the accuracy gain comes solely from the modulator, which proves that the visual modulator is capable of improving different model structures by manipulating the scales of the intermediate feature maps. Noticeably, our method obtains a much lower decay rate for both region similarity and contour accuracy compared to OSVOS and MaskTrack. The accuracy changes of the different methods over time are illustrated in Fig. 4. In the beginning

Table 1. Performance comparison of our approach with recent approaches on DAVIS 2016 and YoutubeObjects. Performance measured in mean IU.

Method	DAVIS 16	YoutubeObjs	with FT	OptFlow	CRF	Speed (s)
OFL [35]	68.0	67.5	-	✓	✓	42.2
BVS [22]	60.0	58.4	-	✗	✗	0.37
ConvGRU[34]	70.1	-	✗	✓	✗	20
VPN[18]	70.2	-	✗	✗	✗	0.63
MaskTrack-B [23]	63.2	66.5	✗	✗	✗	0.24
SFL-B [4]	67.4	-	✗	✓	✗	0.30
OSVOS-B [2]	52.5	44.7	✗	✗	✗	0.14
Ours (Stage 1)	72.2	66.4	✗	✗	✗	0.14
Ours (Stage 1&2)	74.0	69.0	✗	✗	✗	0.14
PLM [32]	70.0	-	✓	✗	✗	0.50
MaskTrack [23]	69.8	71.7	✓	✗	✗	12
SFL [4]	74.8	-	✓	✓	✗	7.9
OSVOS [2]	79.8	74.1	✓	✗	✗	10

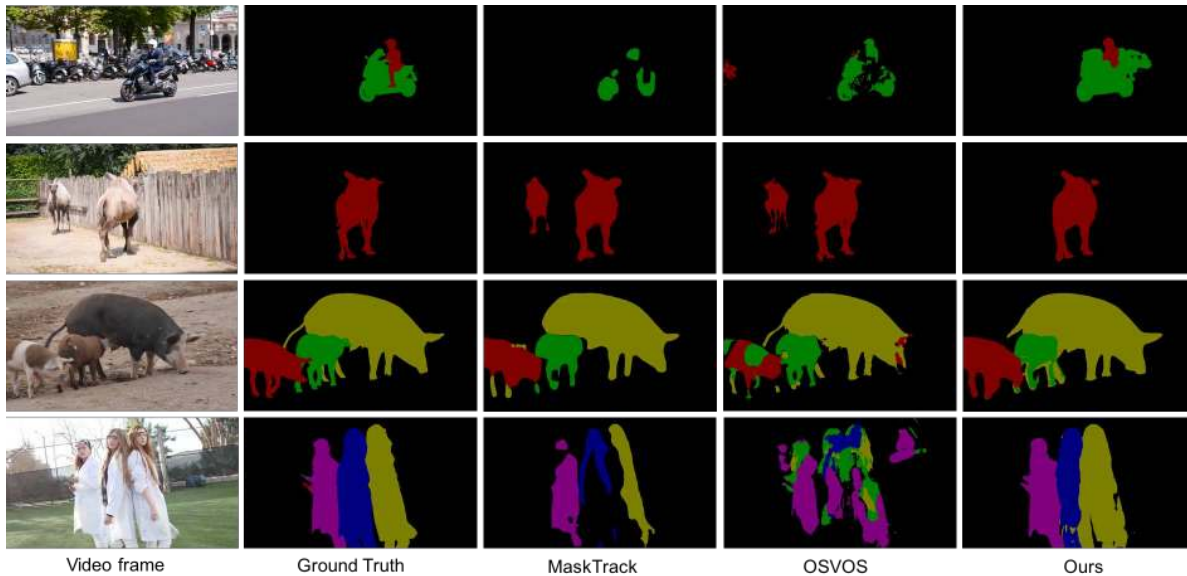


Figure 3. Some qualitative results of our approach compared with two recent state-of-the-art approaches on DAVIS 2017.

of the video, our method lags behind OSVOS and MaskTrack. However, when it proceeds to around 40% of the video, our method is on par with OSVOS and outperforms MaskTrack towards the end of the video. With one-shot fine-tuning, OSVOS and MaskTrack fit to the first frame very well. They are able to obtain high accuracy in the beginning of the video since these frames are all similar to the first one. But as time goes on and the object assumes different poses and appearances, it becomes harder for the fine-tuned model to generalize to new object appearances. Our model is more robust to the appearance changes since it learns a feature embedding (see Section 4.2) for the annotated object which is more tolerant to pose and appearance changes compared to one-shot fine-tuning.

Some qualitative results of our methods compared with

the two previous approaches are shown in Fig. 3. Compared with MaskTrack, our method generally obtains more accurate boundaries, partially due to that the coarse spatial prior forces the model to explore more cues on the image rather than the mask in the previous frame. Compared with OSVOS, our method shows better results when there are multiple similar objects in the image, thanks to the tracking capability provided by the spatial modulator. On the other hand, our method is also shown to work well on unseen object categories in training data. In Fig. 3, the camel and the pigs are unseen object categories in MS-COCO dataset.

4.2. Visualization of the modulation parameters

Our model implicitly learns an embedding with the modulation parameters from the visual modulator for the anno-

Table 2. Comparisons of our approach and two state-of-the-art algorithm on DAVIS 2017 validation set.

Method	with FT	\mathcal{J} mean \uparrow	\mathcal{J} recall \uparrow	\mathcal{J} decay \downarrow	\mathcal{F} mean \uparrow	\mathcal{F} recall \uparrow	\mathcal{F} decay \downarrow
OSVOS-B [2]	\times	18.5	15.9	-0.8	30.0	20.0	0.1
MaskTrack-B [23]	\times	35.3	37.8	39.3	36.4	36.0	42.0
OSVOS-M	\times	36.4	34.8	14.8	39.5	35.3	9.1
MaskTrack-M	\times	44.6	48.7	27.1	47.6	49.3	27.9
OSVOS [2]	\checkmark	55.1	60.2	28.2	62.1	71.3	29.3
MaskTrack [23]	\checkmark	51.2	59.7	28.3	57.3	65.5	29.1
Ours	\times	52.5	60.9	21.5	57.1	66.1	24.3

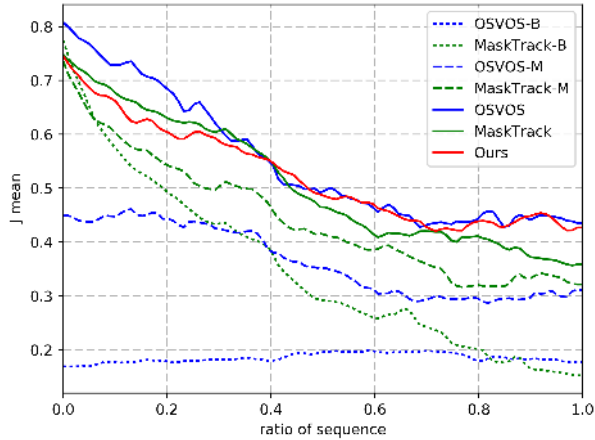


Figure 4. The \mathcal{J} mean performance of different methods over time on DAVIS 2017. Best viewed in color.

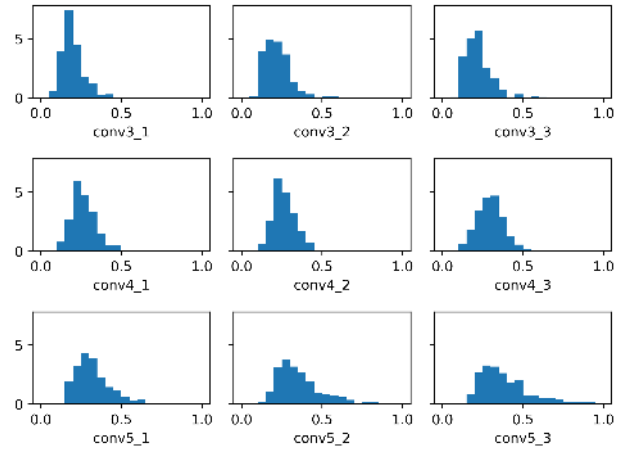


Figure 6. Histograms of standard deviations of γ_c from the visual modulator in different modulation layers. The annotated names are the corresponding convolution layers in VGG16.



Figure 5. Visualization of learned modulation parameters for 100 objects from 10 categories: bicycle, motorcycle, car, bus, truck, dog, cat, horse, cow, person. Zoom in to see details.

tated objects. Intuitively, similar objects should have similar modulation parameters, while different objects should have dramatically different modulation parameters. To visualize this embedding, we extract modulation parameters from 100 object instances in 10 object classes in MS-COCO, and visualize the parameters in a two-dimensional embedding

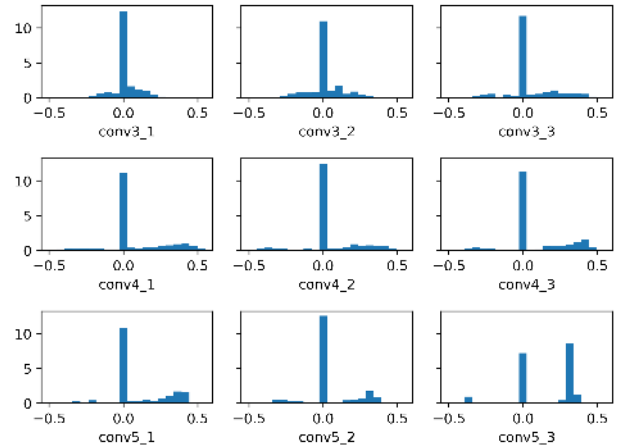


Figure 7. Histograms of magnitude of $\tilde{\gamma}_c$ from the spatial modulator in different modulation layers. The annotated names are the corresponding convolution layers in VGG16.

space using multi-dimensional scaling in Fig. 5. We can see that objects in the same category are mostly clustered together, and similar categories are closer to each other than dissimilar categories. For example, cats and dogs, cars and buses are mixed up due to their similar appearance, while

Table 3. Ablation study of our method on DAVIS 2017.

	Variants	mIU	Δ mIU
Add-on	Ours + Online finetuning	60.8	+8.3
	Ours + CRF	54.4	+1.9
	Ours	52.5	
Model	no visual modulator	33.0	-19.5
	no spatial modulator	40.1	-12.4
Data	- random crop	50.6	-1.9
	- visual guide augmentation	49.5	-1.1
	- spatial guide augmentation	35.6	-13.9

bicycles and dogs, buses and horses are far from each other due to the big visual difference. Mammal classes (cats, dogs, cows, horses, human) are generally clustered together, and man-made objects (cars, buses, bicycles, motorcycles, trucks) are clustered together.

We also investigate the magnitude of the modulation parameters in different layers. The modulation parameters $\{\gamma_c\}$ change according to the visual guide. Therefore, we compute the standard deviations of modulation parameters $\{\gamma_c\}$ in each modulation layer for images in MS-COCO validation set and illustrate them in Fig. 6. An interesting observation is that towards deeper level of the network, the variations of modulation parameters get larger. This shows that the manipulation of feature maps is more dramatic in the last few layers than in early layers of the network. The last few layers of a deep neural network usually learn high-level semantic meanings [37], which could be used to adjust the segmentation model to a specific object more effectively.

We also look into the spatial modulator by extracting the scale parameters $\{\tilde{\gamma}_c\}$ in each layer of the spatial modulator and visualize them in Fig. 7. The magnitudes of $\{\tilde{\gamma}_c\}$ are the relative scales of the spatial guide added to the feature maps in the FCN. The scale of $\{\tilde{\gamma}_c\}$ is proportional to the impact of spatial prior on the intermediate feature maps. Interestingly, we observe sparsity in the values of $\{\tilde{\gamma}_c\}$. Except the last convolution layer `conv5_3`, around 60% of the parameters have zero values, which means only 40% of the feature maps are affected by the spatial prior in these layers. In the layer `conv5_3`, around 70% of the feature maps interact with the spatial guide and most of them are added with a similar scale (note the peak around 0.4) of the spatial guide. This shows that the spatial prior is fused into the feature maps gradually, rather than being effective at the beginning of the network. After all feature extractions are done, the spatial modulator makes a large adjustment to the feature maps, which provides a strong prior of the location of the target object.

4.3. Ablation Study

We study the impact of different ingredients in our method. We conduct experiments on DAVIS 2017 and measure the performance using mean IU. For variants of model structures, we experiment with only using spatial or visual modulators. For data augmentation methods, we experiment with no random crop augmentation for the FCN input, and no affine transformation for the visual guide and the spatial guide. We experiment with CRF as a post-processing step. To investigate the effect of one-shot fine-tuning on our model, we also experiment with standard one-shot fine-tuning using a small number of iterations. Results are shown in Table 3.

By adding a CRF post-processing, our method achieves mIU (mean IU) of 54.4. By one-shot fine-tuning with only 100 iterations for each sequence, our method achieves mIU of 60.8, which is 5.7 better than OSVOS with 1000 iterations. With fine-tuning, our method is still relatively efficient with average running time around 1 s/frame. Without a visual modulator, our model deteriorates to 33.0, while without spatial modulator, our model obtains mIU of 40.1, which shows that the visual guide is more important than the spatial guide. For data augmentation, without random crop, the accuracy drops by 1.9. Without affine data augmentation on the visual guide, the accuracy further decreases by 1.1. Without augmentation on the spatial guide, our model only obtains mIU of 35.6, which is a dramatic drop from 49.5. The results indicate that the spatial guide augmentation contributes most significantly to the performance. Without perturbation, the model might rely on the location of the spatial prior too much and therefore cannot deal with moving objects in real video sequences.

5. Conclusions

In this work, we propose a novel framework to process one-shot video segmentation efficiently. To alleviate the slow speed of one-shot fine-tuning developed by previous FCN-based methods, we propose to use a network modulation approach that mimics the fine-tuning process with one forward pass of the modulator network. We show in experiments that by injecting a limited number of parameters computed by the modulators, the segmentation model can be repurposed to segment an arbitrary object. The proposed network modulation method is a general learning method for few-shot learning problems, which could be applied to other tasks such as visual tracking and image stylization. Our approach falls into the general category of meta-learning, and it would also be worthwhile to investigate other meta-learning approaches for video segmentation. Another piece of future work would be to learn a recurrent representation of the modulation parameters to manipulate the FCN based on temporal information.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. 5
- [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. 1, 2, 4, 5, 6, 7
- [3] Y. Chen, M. W. Hoffman, S. G. Colmenarejo, M. Denil, T. P. Lillicrap, M. Botvinick, and N. de Freitas. Learning to learn without gradient descent by gradient descent. In *ICML*, 2016. 2
- [4] J. Cheng, Y.-H. Tsai, S. Wang, and M.-H. Yang. Segflow: Joint learning for video object segmentation and optical flow. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2, 5, 6
- [5] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei. Deformable convolutional networks. *ICCV*, 2017. 2
- [6] H. de Vries, F. Strub, J. Mary, H. Larochelle, O. Pietquin, and A. C. Courville. Modulating early visual processing by language. *CoRR*, abs/1707.00683, 2017. 3
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 4
- [8] V. Dumoulin, J. Shlens, and M. Kudlur. A learned representation for artistic style, 2017. 2
- [9] Q. Fan, F. Zhong, D. Lischinski, D. Cohen-Or, and B. Chen. Jumpcut: non-successive mask transfer and interpolation for video cutout. In *ACM Trans. Graph.*, 34(6), 2015. 2
- [10] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 2
- [11] G. Ghiasi, H. Lee, M. Kudlur, V. Dumoulin, and J. Shlens. Exploring the structure of a real-time, arbitrary neural artistic stylization network. In *BMVC*, 2017. 2
- [12] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *CVPR*, 2015. 2, 4
- [13] B. Hariharan and R. Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *ICCV*, 2017. 2
- [14] X. Huang and S. J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *CoRR*, abs/1703.06868, 2017. 2, 3
- [15] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. *CVPR*, 2017. 1
- [16] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *NIPS*, pages 2017–2025, 2015. 2
- [17] S. D. Jain and K. Grauman. Supervoxel-consistent foreground propagation in video. In *ECCV*, 2014. 2, 5
- [18] V. Jampani, R. Gadde, and P. V. Gehler. Video propagation networks. In *CVPR*, 2017. 2, 5, 6
- [19] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, pages 109–117, 2011. 2, 5
- [20] F. Li, T. Kim, A. Humayun, D. Tsai, and J. M. Rehg. Video segmentation by tracking many figure-ground segments. In *ICCV*, 2013. 1
- [21] 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. 4
- [22] N. Märki, F. Perazzi, O. Wang, and A. Sorkine-Hornung. Bilateral space video segmentation. In *CVPR*, 2016. 1, 2, 5, 6
- [23] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *CVPR*, 2017. 1, 2, 4, 5, 6, 7
- [24] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 1, 5
- [25] F. Perazzi, O. Wang, M. Gross, and A. Sorkine-Hornung. Fully connected object proposals for video segmentation. In *ICCV*, 2015. 2
- [26] E. Perez, H. de Vries, F. Strub, V. Dumoulin, and A. C. Courville. Learning visual reasoning without strong priors. *CoRR*, abs/1707.03017, 2017. 2, 3
- [27] 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. 4, 5
- [28] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. *ICLR*, 2017. 2
- [29] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *CVPR*, 2015. 1
- [30] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, 2016. 2
- [31] E. Shelhamer, J. Long, and T. Darrell. Fully convolutional networks for semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):640–651, 2017. 1
- [32] 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 *CVPR*, 2017. 1, 2, 5, 6
- [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3, 4
- [34] P. Tokmakov, K. Alahari, and C. Schmid. Learning video object segmentation with visual memory. In *ICCV*, 2017. 2, 5, 6
- [35] Y.-H. Tsai, M.-H. Yang, and M. J. Black. Video segmentation via object flow. In *CVPR*, 2016. 1, 2, 5, 6
- [36] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *NIPS*, 2016. 2
- [37] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014. 8