

Learning Video Object Segmentation from Static Images

*Federico Perazzi^{1,2} *Anna Khoreva³ Rodrigo Benenson³
Bernt Schiele³ Alexander Sorkine-Hornung¹

¹Disney Research ²ETH Zurich

³Max Planck Institute for Informatics, Saarbrücken, Germany

Abstract

Inspired by recent advances of deep learning in instance segmentation and object tracking, we introduce the concept of convnet-based guidance applied to video object segmentation. Our model proceeds on a per-frame basis, guided by the output of the previous frame towards the object of interest in the next frame. We demonstrate that highly accurate object segmentation in videos can be enabled by using a convolutional neural network (convnet) trained with static images only. The key component of our approach is a combination of offline and online learning strategies, where the former produces a refined mask from the previous frame estimate and the latter allows to capture the appearance of the specific object instance. Our method can handle different types of input annotations such as bounding boxes and segments while leveraging an arbitrary amount of annotated frames. Therefore our system is suitable for diverse applications with different requirements in terms of accuracy and efficiency. In our extensive evaluation, we obtain competitive results on three different datasets, independently from the type of input annotation.

1. Introduction

Convolutional neural networks (convnets) have shown outstanding performance in many fundamental areas in computer vision, enabled by the availability of large-scale annotated datasets (e.g., ImageNet classification [24, 43]). However, some important challenges in video processing can be difficult to approach using convnets, since creating a sufficiently large body of densely, pixel-wise annotated video data for training is usually prohibitive.

One example of such domain is video object segmentation. Given only one or a few frames annotated with segmentation masks of a particular object instance, the task of video object segmentation is to accurately segment the same

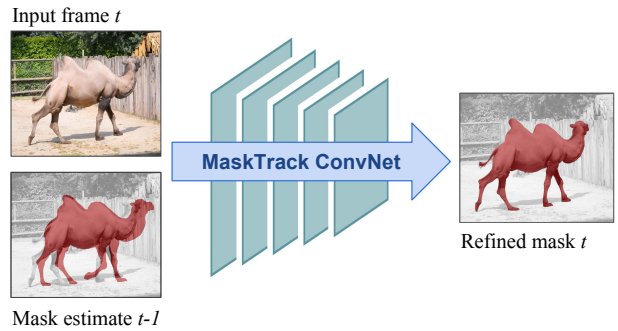


Figure 1: Given a rough mask estimate from the previous frame $t - 1$ we train a convnet to provide a refined mask output for the current frame t .

instance in all other frames of the video. Current top performing approaches either interleave box tracking and segmentation [53], or propagate the first frame mask annotation in space-time via CRF or GrabCut-like techniques [29, 49].

One of the key insights and contributions of this paper is that fully annotated video data is not necessary. We demonstrate that highly accurate video object segmentation can be enabled using a convnet trained with *static images* only. We show that a convnet designed for semantic image segmentation [8] can be utilized to perform per-frame instance segmentation, i.e., segmentation of generic objects while distinguishing different instances of the same class. For each new video frame the network is guided towards the object of interest by feeding in the previous frame mask estimate. We therefore refer to our approach as *guided instance segmentation*. To the best of our knowledge, it represents the first fully trained approach to video object segmentation.

Our system is efficient due to its feed-forward architecture and can generate high quality results in a single pass over the video, without the need for considering more than one frame at a time. This is in stark contrast to many other video segmentation approaches, which usually require global connections over multiple frames or even the whole

*The first two authors contributed equally.

video sequence in order to achieve coherent results. Furthermore, our method can handle different types of annotations and even simple bounding boxes as input are sufficient to obtain competitive results, making our method flexible with respect to various practical applications with different requirements in terms of human supervision.

Key to the video segmentation quality of our approach is the combination of offline and online learning strategies. In the offline phase, we use deformation and coarsening on the image masks in order to train the network to produce accurate output masks from their rough estimates. An online training phase extends ideas from previous works on object tracking [12, 32] to the task of video segmentation and enables the method to be easily optimized with respect to an object of interest in a novel input video.

The result is a single, homogeneous system that compares favourably to most classical approaches on three extremely heterogeneous video segmentation benchmarks, despite using the same model and parameters across all videos. We provide a detailed ablation study and explore the impact of varying number and types of annotations. Moreover, we discuss extensions of the proposed model, allowing to improve the quality even further.

2. Related Works

The idea of performing video object segmentation via tracking at the pixel level is at least a decade old [40]. Recent approaches interweave box tracking with box-driven segmentation (e.g. TRS [53]), or propagate the first frame segmentation via graph labeling approaches.

Local propagation. JOTS [52] builds a graph over neighboring frames connecting superpixels and (generic) object parts to solve the video labeling task. ObjFlow [49] builds a graph over pixels and superpixels, uses convnet based appearance terms, and interleaves labeling with optical flow estimation. Instead of using superpixels or proposals, BVS [29] formulates a fully-connected pixel-level graph between frames and efficiently infer the labeling over the vertices of a spatio-temporal bilateral grid [7]. Because these methods propagate information only across neighbor frames they have difficulties capturing long range relationships and ensuring globally consistent segmentation.

Global propagation. In order to overcome these limitations, some methods have proposed to use long-range connections between video frames [15, 25, 48, 55]. In particular, we compare to FCP [35], Z15 [56], NLC [15] and W16 [50] which build a global graph structure over object proposal segments, and then infer a consistent segmentation. A limitation of methods utilizing long-range connections is that they have to operate on larger image regions such as superpixels or object proposals for acceptable speed and memory usage, compromising on their ability to

handle fine details.

Unsupervised segmentation. Another family of works perform moving object segmentation (over all parts of the image), and selects post-hoc the space-time tube that best match the annotation [18, 26, 33, 53]. In contrast, our approach side-steps the use of any intermediate tracked boxes, superpixels or object proposals and proceeds on a per-frame basis, therefore efficiently handling even long sequences at full detail. We focus on propagating the first frame segmentation forward onto future frames, using an online fine-tuned convnet as appearance model for segmenting the object of interest in the next frames.

Box tracking. Some previous works have investigated approaches that improve segmentation quality by leveraging object tracking and vice versa [10, 13, 17, 40, 53]. More recent, state-of-the-art tracking methods are based on discriminative correlation filters over handcrafted features (e.g. HOG) and over frozen deep learned features [11, 12], or are convnet based trackers on their own right [20, 32]. Our approach is most closely related to the latter group. GOTURN [20] proposes to train offline a convnet so as to directly regress the bounding box in the current frame based on the object position and appearance in the previous frame. MDNet [32] proposes to use online fine-tuning of a convnet to model the object appearance. Our training strategy is inspired by GOTURN for the offline part, and MDNet for the online stage. Compared to the aforementioned methods our approach operates at pixel level masks instead of boxes. Differently from MDNet, we do not replace the domain-specific layers, instead fine-tuning all the layers on the available annotations for each individual video sequence.

Instance segmentation. At each frame, video object segmentation outputs a single instance segmentation. Given an estimate of the object location and size, bottom-up segment proposals [38] or GrabCut [42] variants can be used as shape guesses. Also specific convnet architectures have been proposed for instance segmentation [19, 36, 37, 54]. Our approach outputs per-frame instance segmentation using a convnet architecture, inspired by works from other domains like [6, 44, 54]. A concurrent work [5] also exploits convnets for video object segmentation. Differently from our approach their segmentation is not guided, and therefore it cannot distinguish multiple instances of the same object.

Interactive video segmentation. Applications such as video editing for movie production often require a level of accuracy beyond the current state-of-the-art. Thus several works have also considered video segmentation with variable annotation effort, enabling human interaction using clicks [22, 47, 51] or strokes [1, 16, 57]. In this work we consider instead box or segment annotations on multiple frames. In §5 we report results when varying the amount of annotation effort, from one frame per video to all frames.

3. Method

We approach the video object segmentation problem from a different perspective we refer as: convnet-based *guided instance segmentation*. For each new frame we wish to label pixels as object/non-object of interest, for this we build upon the architecture of the existing pixel labelling convnet and train it to generate per-frame instance segments. We pick DeepLabv2 [8], but our approach is agnostic of the specific architecture selected. The challenge is then: how to inform the network which instance to segment? We solve this by using two complementary strategies. First we guide the network towards the instance of interest by feeding in the previous frame mask estimate during offline training (§3.1). Second, we employ online training to fine-tune the model to incorporate specific knowledge of the object instance (§3.2).

3.1. Offline Training

In order to guide the pixel labeling network to segment the object of interest, we begin by expanding the convnet input from RGB to RGB+mask channels. The extra mask channel is meant to provide an estimate of the visible area of the object in the current frame, its approximate location and shape. We can then train the labelling convnet to output an accurate segmentation of the object, given as input the current image and a rough estimate of the object mask. Our tracking network is de-facto a "mask refinement" network.

There are two key observations that make this approach practical. First, very rough input masks are enough for our trained network to provide sensible output segments. Even a large bounding box as input will result in a reasonable output (see §5.2). The main role of the input mask is to point the convnet towards the correct object instance to segment. Second, this particular approach does not require us to use video as training data, such as done in [3, 5, 20, 32]. Because we only use a mask as additional input, instead of an image crop as in [3, 20], we can synthesize training samples from single frame instance segmentation annotations. This allows us to train from a large set of diverse images, instead of having to rely on scarce video annotations.

Figure 1 shows our simplified model. To simulate the noise of the previous frame output, during offline training, we generate input masks by deforming the annotations via affine transformation as well as non-rigid deformations via thin-plate splines [4], followed by a coarsening step (dilation morphological operation) to remove details of the object contour. We apply this data generation procedure over a dataset of $\sim 10^4$ images containing diverse object instances, see examples in the supplementary material. At test time, given the mask estimate at time $t - 1$, we apply the dilation operation and use the resulting rough mask as input for object segmentation in frame t .

The affine transformations and non-rigid deformations

aim at modelling the expected motion of an object between two frames. The coarsening permits us to generate training samples that resembles the test time data, simulating the blobby shape of the output mask given from the previous frame by the convnet. These two ingredients make the estimation more robust to noisy segmentation estimates while helping to avoid accumulation of errors from the preceding frames. The trained convnet has learnt to do guided instance segmentation similar to networks like SharpMask [37], DeepMask [36] and Hypercolumns [19], but instead of taking a bounding box as guidance, we can use an arbitrary input mask. The training details are described in §4.

When using offline training only, the segmentation procedure consists of two steps: the previous frame mask is coarsened and then fed into the trained network to estimate the current frame mask. Since objects have a tendency to move smoothly through space, the object mask in the preceding frame will provide a good guess in the current frame and simply copying the coarse mask from the previous frame is enough. This approach is fast and already provides good results. We also experimented using optical flow to propagate the mask from one frame to the next, but found the optical flow errors to offset the gains.

With only the offline trained network, the proposed approach allows us to achieve competitive performance compared to previously reported results (see §5.2). However the performance can be further improved by integrating online training strategy as described in the next section.

3.2. Online Training

For further boosting the video segmentation quality, we borrow and extend ideas that were originally proposed for object tracking. Current top performing tracking techniques [12, 32] use some form of online training. We thus consider improving results by adding online fine-tuning as a second strategy.

The idea is to use, at test time, the segment annotation of the first video frame as additional training data. Using augmented versions of this single frame annotation, we proceed to fine-tune the model to become more specialized for the specific object instance at hand. We use a similar data augmentation as for offline training. On top of affine and non-rigid deformations for the input mask, we also add image flipping and rotations. We generate $\sim 10^3$ training samples from this single annotation, and proceed to fine-tune the model previously trained offline.

With online fine-tuning, the network weights partially capture the appearance of the specific object being tracked. The model aims to strike a balance between general instance segmentation (so as to generalize to the object changes), and specific instance segmentation (so as to leverage the common appearance across video frames). The details of the online fine-tuning are provided in §4. In our experiments




Figure 2: Examples of optical flow magnitude images. Top: RGB images. Bottom: corresponding motion magnitude estimates encoded into as gray-scale images.

we only perform fine-tuning using the annotated frame(s).

To the best of our knowledge our approach is the first to use a pixel labelling network (like DeepLabv2 [8]) for the task of video object segmentation. We name our full approach, using both offline and online training, MaskTrack.

3.3. Variants

Additionally we consider variations of the proposed model. First, we demonstrate that our approach is flexible and could handle different types of input annotations, using less supervision in the first frame annotation. Second, we describe how motion information could be easily integrated in the system, improving the quality of the object segments.

Box annotation. In this paragraph, we discuss a variant named MaskTrack_{Box}, that takes a bounding box annotation in the first frame as an input supervision instead of a segmentation mask. To this end, we train a similar convnet that fed with a bounding-box annotation as an input outputs a segment. Once the first frame bounding box is converted to a segment, we switch back to the MaskTrack model that uses as guidance the output mask from the previous frame.

Optical flow. On top of MaskTrack, we consider to employ optical flow as a source of additional information to guide the segmentation. Given a video sequence, we compute the optical flow using EpicFlow [41] with Flow Fields matches [2] and convolutional boundaries [30]. In parallel to the vanilla MaskTrack, we proceed to compute a second output mask using the magnitude of the optical flow as input image (replicated into a three channel image). The model is used as-is, without retraining. Although it has been trained on RGB images, this strategy works as object flow magnitude roughly looks like a gray-scale object, and still captures useful object shape information, see examples in Figure 2. Using the RGB model allows to avoid training the convnet on video datasets annotated with masks. We then fuse by averaging the output scores given by the two parallel networks, respectively fed with RGB images and optical flow magnitude as input. As shown in Table 1, optical flow provides complementary information to MaskTrack with RGB images, improving the overall performance.

4. Network Implementation and Training

Following, we describe the implementation details of our approach. Specifically, we provide additional information regarding the network initialization, the offline and online training strategies and the data augmentation.

Network. For all our experiments we use the training and test parameters of DeepLabv2-VGG network [8]. The model is initialized from a VGG16 network pre-trained on ImageNet [46]. For the extra mask channel of filters in the first convolutional layer we use gaussian initialization. We also tried zero initialization, but observed no difference.

Offline training. The advantage of our method is that it does not require expensive pixel-wise video annotations for training. Thus we can employ existing image datasets. However, in order for our model to generalize well across different videos, we avoid training on datasets that are biased towards certain semantic classes, such as COCO [28] or Pascal [14]. Instead we combine images and annotations from several saliency segmentation datasets (ECSSD [45], MSRA10K [9], SOD [31], and PASCAL-S [27]), resulting in an aggregated set of 11 282 training images.

The input masks for the extra channel are generated by deforming the binary segmentation masks via affine transformation and non-rigid deformations, as discussed in §3.1. For affine transformation we consider random scaling ($\pm 5\%$ of object size) and translation ($\pm 10\%$ shift). Non-rigid deformations are done via thin-plate splines [4] using 5 control points and randomly shifting the points in x and y directions within $\pm 10\%$ margin of the original segmentation mask width and height. Next, the mask is coarsened using dilation operation with 5 pixel radius. This mask deformation procedure is applied over all object instances in the training set. For each image two different masks are generated. We refer the reader to the supplementary material for visual examples of deformed masks.

The convnet training parameters are identical to those proposed in [8]. Therefore we use stochastic gradient descent (SGD) with mini-batches of 10 images and a polynomial learning policy with initial learning rate of 0.001. The momentum and weight decay are set to 0.9 and 0.0005, respectively. The network is trained for 20k iterations.

Online training. For online adaptation we fine-tune the model previously trained offline on the first frame for 200 iterations with training samples generated from the first frame annotation. We augment the first frame by image flipping and rotations as well as by deforming the annotated masks for an extra channel via affine and non-rigid deformations with the same parameters as for the offline training. This results in an augmented set of $\sim 10^3$ training images.

The network is trained with the same learning parameters as for offline training, fine-tuning all convolutional layers.

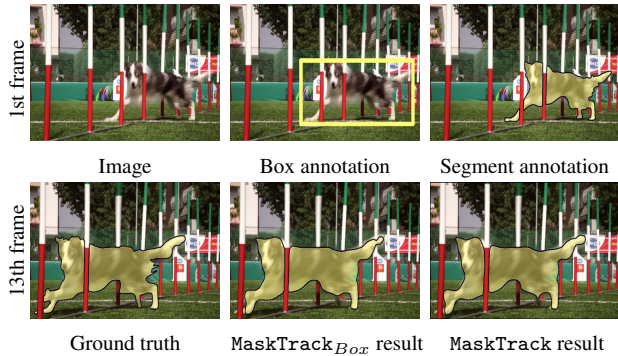


Figure 3: By propagating annotation from the 1st frame, either from segment or just bounding box annotations, our system generates results comparable to ground truth.

At test time our base MaskTrack system runs at about 12 seconds per frame (averaged over DAVIS, amortizing the online fine-tuning time over all video frames), which is a magnitude faster compared to ObjFlow [49] (takes 2 minutes per frame, averaged over DAVIS).

5. Results

In this section we describe our evaluation protocol (§5.1), study the importance of the different components of our system (§5.2), and report results comparing to state-of-the-art techniques over three datasets (§5.3), as well as comparing the effects of different amounts of annotation on the resulting segmentation quality (§5.4). Additional results are provided in the supplementary material.

5.1. Experimental setup

Datasets. We evaluate the proposed approach on three different video object segmentation datasets: DAVIS [34], YoutubeObjects [39], and SegTrack-v2 [26]. These datasets include assorted challenges such as appearance change, occlusion, motion blur and shape deformation.

DAVIS [34] consists of 50 high quality videos, totaling 3 455 frames. Pixel-level segmentation annotations are provided for each frame, where one single object or two connected objects are separated from the background.

YoutubeObjects [39] includes videos with 10 object categories. We consider the subset of 126 videos with more than 20 000 frames, for which the pixel-level ground truth segmentation masks are provided by [21].

SegTrack-v2 [26] contains 14 video sequences with 24 objects and 947 frames. Every frame is annotated with a pixel-level object mask. As instance-level annotations are provided for sequences with multiple objects, each specific instance segmentation is treated as separate problem.

Evaluation. We evaluate using the standard mIoU metric: intersection-over-union of the estimated segmentation and the ground truth binary mask, also known as *Jaccard Index*, averaged across videos. For DAVIS we use the pro-

vided benchmark code [34], which excludes the first and the last frames from the evaluation. For YoutubeObjects and SegTrack-v2 only the first frame is excluded.

Previous works used different evaluation procedures. To ensure a consistent comparison between methods, when needed, we re-computed scores from the publicly available output masks, or reproduced the results using the available open source code. In particular, we collected new results for ObjFlow [49] and BVS [29] in order to present other methods with results across the three datasets.

5.2. Ablation study

We first study different ingredients of our method. We experiment on the DAVIS dataset and measure the performance using the mean intersection-over-union metric (mIoU). Table 1 shows the importance of each of the ingredients described in §3 and reports the improvement of adding extra components to the MaskTrack model.

Add-ons. We first study the effect of adding a couple of ingredients on top of our base MaskTrack system, which are specifically fine-tuned for DAVIS. We see that optical flow provides complementary information to the appearance, boosting further the results (74.8 \rightarrow 78.4). Adding on top a well-tuned post-processing CRF [23] can gain a couple of mIoU points, reaching 80.3% mIoU on DAVIS, the best known result on this dataset.

Albeit optical flow can provide interesting gains, we found it to be brittle when going across different datasets. Different strategies to handle optical flow provide 1 \sim 4% on each dataset, but none provide consistent gains across all datasets; mainly due to failure modes of the optical flow algorithms. For the sake of presenting a single model with fix parameters across all datasets, we refrain from using a per-dataset tuned optical flow in the results of §5.3.

Training. We next study the effect of offline/online training of the network. By disabling online fine-tuning, and only relying on offline training we see a \sim 5 IoU percent points drop, showing that online fine-tuning indeed expand the tracking capabilities. If instead we skip offline training and only rely on online fine-tuning performance drop drastically, albeit the absolute quality (57.6 mIoU) is surprisingly high for a system trained on ImageNet+single frame.

By reducing the amount of training data from 11k to 5k we only see a minor decrease in mIoU; this indicates that even with the small amount of training data we can achieve reasonable performance. That being said, further increase of the training data volume would lead to improved results.

Additionally, we explore the effect of the offline training on video data instead of using static images. We train the model on the annotated frames of two combined datasets, SegTrack-v2 and YoutubeObjects. By switching to train on video data we observe a minor decrease in mIoU; this could be explained by lack of diversity in the video training data

Aspect	System variant	mIoU	Δ mIoU
Add-ons	MaskTrack+Flow+CRF	80.3	+1.9
	MaskTrack+Flow	78.4	+3.6
	MaskTrack	74.8	-
Training	No online fine-tuning	69.9	-4.9
	No offline training	57.6	-17.2
	Reduced offline training	73.2	-1.6
	Training on video	72.0	-2.8
Mask deformation	No dilation	72.4	-2.4
	No deformation	17.1	-57.7
	No non-rigid deformation	73.3	-1.5
Input channel	Boxes	69.6	-5.2
	No input	72.5	-2.3

Table 1: Ablation study of our MaskTrack method on DAVIS. Given our full system, we change one ingredient at a time, to see each individual contribution. See §5.2.

due to the small scale of the existing datasets, as well as the effect of the domain shift between different benchmarks. This shows that employing static images in our approach does not result in any performance drop.

Mask deformation. We also study the influence of mask deformations. We see that coarsening the mask via dilation provides a small gain, as well as adding non-rigid deformations. All-and-all, Table 1 shows that the main factor affecting the quality is using any form of mask deformations when creating the training samples (both for offline and online training). This ingredient is critical for our overall approach, making the segmentation estimation more robust at test time to the noise in the input mask.

Input channel. Next we experiment with different variants of the extra channel input. Even by changing the input from segments to boxes, a model trained for this modality still provides reasonable results. A failure mode of this approach is the generation of small blobs outside of the object. As a result the guidance from the previous frame produces a noisy box, yielding inaccurate segmentation. The error accumulates forward over the entire video sequence.

We also evaluated a model that does not use any mask input. Without the additional input channel, this pixel labelling convnet was trained offline as a salient object segmenter and fine-tuned online to capture the appearance of the object of interest. This model obtains competitive results (72.5 mIoU) on DAVIS, since the object to segment is also salient for this dataset. However, while experimenting on SegTrack-v2 and YoutubeObjects, we observed a significant drop in performance without using guidance from the previous frame mask as these two datasets have a weaker bias towards salient objects compared to DAVIS.

5.3. Single Frame Annotations

Table 2 presents results when the first frame is annotated with an object segmentation mask. This is the protocol

Method	Dataset, mIoU		
	DAVIS	YoutbObjs	SegTrack-v2
Box oracle	45.1	55.3	56.1
Grabcut oracle	67.3	67.6	74.2
ObjFlow [49]	71.4	70.1	67.5
BVS [29]	66.5	59.7	58.4
NLC [15]	64.1	-	-
FCP [35]	63.1	-	-
W16 [50]	-	59.2	-
Z15 [56]	-	52.6	-
TRS [53]	-	-	69.1
MaskTrack	74.8	71.7	67.4
MaskTrack _{Box}	73.7	69.3	62.4

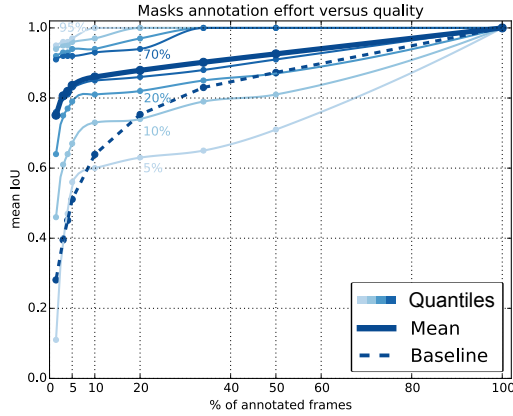
Table 2: Video object segmentation results on three datasets. Compared to related state of the art, our approach provides consistently good results. On DAVIS the extended version of our system MaskTrack+Flow+CRF reaches 80.3 mIoU. See §5.3 for details.

commonly used on DAVIS, SegTrack-v2, and YoutubeObjects. MaskTrack obtains competitive performance across all three datasets. This is achieved using our purely frame-by-frame feed-forward system, using the exact same model and parameters across all datasets. Our MaskTrack results are obtained in a single pass, do not use any global optimization, not even optical flow. We believe this shows the promise of formulating video object segmentation from the instance segmentation perspective.

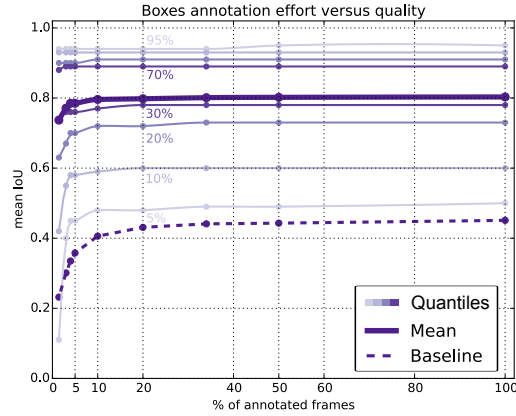
On SegTrack-v2, JOTS [52] reported higher numbers (71.3 mIoU), however, they report tuning their method parameters’ per video, and thus it is not comparable to our setup with fix-parameters. Table 2 also reports results for the MaskTrack_{Box} variant described in §3.3. Starting only from box annotations on the first frame, our system still generates comparably good results (see Figure 3), remaining on the top three best results in all the datasets covered.

By adding additional ingredients specifically tuned for different datasets, such as optical flow (see §3.3) and CRF post-processing, we can push the results even further, reaching 80.3 mIoU on DAVIS, 72.6 on YoutubeObjects and 70.3 on SegTrack-v2. The dataset specific tuning is described in the supplementary material. Figure 6 presents qualitative results of the proposed MaskTrack model across three different datasets.

Attribute-based analysis Figure 5 presents a more detailed evaluation on DAVIS [34] using video attributes. The attribute based analysis shows that our generic model, MaskTrack, is robust to various video challenges presented in DAVIS. It compares favourably on any subset of videos sharing the same attribute, except camera-shake, where ObjFlow [49] marginally outperforms our approach. We observe that MaskTrack handles fast-motion and motion-blur well, which are typical failure cases for methods rely-



a) Segment annotations.



b) Box Annotations.

Figure 4: Percent of annotated frames versus video object segmentation quality. We report mean IoU, and quantiles at 5, 10, 20, 30, 70, 80, 90, and 95%. Results on DAVIS, using segment or box annotations. The baseline simply copies the annotations to adjacent frames. Discussion in §5.4.

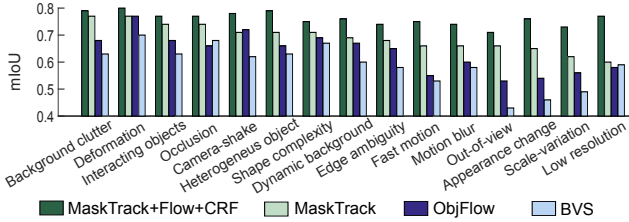


Figure 5: Attribute based evaluation on DAVIS.

ing on spatio-temporal connections [29, 49].

Due to the online fine-tuning on the first frame annotation of a new video, our system is able to capture the appearance of the specific object of interest. This allows it to better recover from occlusions, out-of-view scenarios and appearance changes, which usually affect methods that strongly rely on propagating segmentations on a per-frame basis.

Incorporating optical flow information into MaskTrack substantially increases robustness on all categories. As one could expect, MaskTrack+Flow+CRF better discriminates cases involving color ambiguity and salient motion. We also observe improvements in cases with scale-variation and low-resolution objects.

Conclusion. With our simple, generic system for video object segmentation we are able to achieve competitive results with existing techniques, on three different datasets. These results are obtained with fixed parameters, from a forward-pass only, using only static images during offline training. We also reach good results even when using only a box annotation as starting point.

5.4. Multiple Frames Annotations


In some applications, e.g. video editing for movie production, one may want to consider more than a single frame annotation on videos. Figure 4 shows the segmentation

quality result when considering different number of annotated frames on the DAVIS dataset. We evaluate both pixel accurate segmentation and bounding box annotations.

For these experiments, we run our method twice, forward and backwards; and for each frame we pick the result closest in time to the annotated frame (either from forward or backwards propagation). Here, the online fine-tuning uses all annotated frames instead of only the first one. For the experiments with box annotations in Figure 4, we use a similar procedure to MaskTrack_{Box}. Box annotations are first converted to segments, and then apply MaskTrack as-is, treating these as the original segment annotations. The evaluation reports the mean IoU results when annotating one frame only (same as Table 2), and every 40th, 30th, 20th, 10th, 5th, 3rd, and 2nd frame. Since DAVIS videos have length ~ 100 frames, 1 annotated frame corresponds to $\sim 1\%$, otherwise annotations every 20th is 5% of annotated frames, 10th 10%, 5th 20%, etc. We follow the same evaluation protocol as in §5.3, ignoring first and last frames, and including the annotated frames in the evaluation (this is particularly relevant for the box annotation results).

Other than mean IoU we also show the quantile curves indicating the cutting line for the 5%, 10%, 20%, etc. lowest quality video frame results. This gives a hint of how much targeted additional annotations might be needed. The higher mean IoU of these quantiles are, the better. The baseline for these experiments consists in directly copying the ground truth annotations from the nearest annotated neighbour. For visual clarity, in Figure 4, we only include the mean value for the baseline experiment.

Analysis. We can see that results in Figures 4 show slightly different trends. When using segment annotations the baseline quality increases steadily until reaching IoU 1 when all frames are annotated. Our MaskTrack approach provides large gains with 30% of annotated frames or less. For in-



1st frame, GT segment Results with MaskTrack, the frames are chosen equally distant based on the video sequence length

Figure 6: Qualitative results of three different datasets. Our algorithm is robust to challenging situations such as occlusions, fast motion, multiple instances of the same semantic class, object shape deformation, camera view change and motion blur.

stance when annotating 10% of the frames we reach mIoU 0.86, with the 20% quantile at 0.81 mIoU. This means with only 10% of annotated frames, 80% of all video frames will have a mean IoU above 0.8, which is good enough to be used for many applications, or can serve as initialization for a refinement process. With 10% of annotated frames the baseline only reaches 0.64 mIoU. When using box annotations the quality of the baseline and our method saturates. There is only so much information our instance segmenter can estimate from boxes. After 10% of annotated frames, not much additional gain is obtained. Interestingly, the mean IoU and 30% quantile here both reach ~ 0.8 mIoU. Additionally, 70% of the frames have IoU above 0.89.

Conclusion. Results indicate that with only 10% of annotated frames we can reach satisfactory quality, even when using only bounding box annotations. We see that with moving from one annotation per video to two or three frames ($1\% \rightarrow 3\% \rightarrow 4\%$) quality increases sharply, showing that our system can adequately leverage a few extra annotations per video.

6. Conclusion

We have presented a novel approach to video object segmentation. By treating video object segmentation as a

guided instance segmentation problem, we have proposed to use a pixel labelling convnet for frame-by-frame segmentation. By exploiting both offline and online training with image annotations only our approach is able to produce highly accurate video object segmentation. The proposed system is generic and reaches competitive performance on three extremely heterogeneous video segmentation benchmarks, using the same model and parameters across all videos. The method can handle different types of input annotations and our results are competitive even when using only bounding box annotations (instead of segmentation masks).

We provided a detailed ablation study, and explored the effect of varying the amount of annotations per video. Our results show that with only one annotation every 10th frame we can reach 85% mIoU quality. Considering we only do per-frame instance segmentation without any form of global optimization, we deem these results encouraging to achieve high quality via additional post-processing.

We believe the use of labeling convnets for video object segmentation is a promising strategy. Future work should consider exploring more sophisticated network architectures, incorporating temporal dimension and global optimization strategies.

References

- [1] X. Bai, J. Wang, D. Simons, and G. Sapiro. Video snap-cut: robust video object cutout using localized classifiers. In *TOG*, 2009. 2
- [2] C. Bailer, B. Taetz, and D. Stricker. Flow fields: Dense correspondence fields for highly accurate large displacement optical flow estimation. In *ICCV*, 2015. 4
- [3] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr. Fully-convolutional siamese networks for object tracking. *arXiv preprint arXiv:1606.09549*, 2016. 3
- [4] F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *PAMI*, 1989. 3, 4
- [5] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixà, D. Cremers, and L. V. Gool. One-shot video object segmentation. *arXiv:1611.05198*, 2016. 2, 3
- [6] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik. Human pose estimation with iterative error feedback. In *CVPR*, 2016. 2
- [7] J. Chen, S. Paris, and F. Durand. Real-time edge-aware image processing with the bilateral grid. *ACM Trans. Graph.*, 26(3), July 2007. 2
- [8] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv:1606.00915*, 2016. 1, 3, 4
- [9] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. S. Torr, and S.-M. Hu. Global contrast based salient region detection. *TPAMI*, 2015. 4
- [10] P. Chockalingam, S. N. Pradeep, and S. Birchfield. Adaptive fragments-based tracking of non-rigid objects using level sets. In *ICCV*, 2009. 2
- [11] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg. Convolutional features for correlation filter based visual tracking. In *ICCV workshop*, pages 58–66, 2015. 2
- [12] M. Danelljan, A. Robinson, F. S. Khan, and M. Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *ECCV*. Springer, 2016. 2, 3
- [13] S. Duffner and C. Garcia. Pixeltrack: A fast adaptive algorithm for tracking non-rigid objects. In *ICCV*, 2013. 2
- [14] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 2015. 4
- [15] A. Faktor and M. Irani. Video segmentation by non-local consensus voting. In *BMVC*, 2014. 2, 6
- [16] Q. Fan, F. Zhong, D. Lischinski, D. Cohen-Or, and B. Chen. Jumpcut: Non-successive mask transfer and interpolation for video cutout. *SIGGRAPH Asia*, 2015. 2
- [17] M. Godec, P. M. Roth, and H. Bischof. Hough-based tracking of non-rigid objects. In *ICCV*, 2011. 2
- [18] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph-based video segmentation. In *CVPR*, 2010. 2
- [19] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *CVPR*, 2015. 2, 3
- [20] D. Held, S. Thrun, and S. Savarese. Learning to track at 100 fps with deep regression networks. In *ECCV*, 2016. 2, 3
- [21] S. D. Jain and K. Grauman. Supervoxel-consistent foreground propagation in video. In *ECCV*, 2014. 5
- [22] S. D. Jain and K. Grauman. Click carving: Segmenting objects in video with point clicks. In *HCOMP*, 2016. 2
- [23] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*. 2011. 5
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [25] Y. J. Lee, J. Kim, and K. Grauman. Key-segments for video object segmentation. In *Proc. ICCV*, 2011. 2
- [26] F. Li, T. Kim, A. Humayun, D. Tsai, and J. M. Rehg. Video segmentation by tracking many figure-ground segments. In *ICCV*, 2013. 2, 5
- [27] Y. Li, X. Hou, C. Koch, J. M. Rehg, and A. L. Yuille. The secrets of salient object segmentation. In *CVPR*, 2014. 4
- [28] T. 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
- [29] N. Maerki, F. Perazzi, O. Wang, and A. Sorkine-Hornung. Bilateral space video segmentation. In *CVPR*, 2016. 1, 2, 5, 6, 7
- [30] K. Maninis, J. Pont-Tuset, P. Arbeláez, and L. V. Gool. Convolutional oriented boundaries. 4
- [31] V. Movahedi and J. H. Elder. Design and perceptual validation of performance measures for salient object segmentation. In *CVPR Workshops*, 2010. 4
- [32] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 2, 3
- [33] A. Papazoglou and V. Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, 2013. 2
- [34] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. V. Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 5, 6
- [35] F. Perazzi, O. Wang, M. Gross, and A. Sorkine-Hornung. Fully connected object proposals for video segmentation. In *ICCV*, 2015. 2, 6
- [36] P. O. Pinheiro, R. Collobert, and P. Dollár. Learning to segment object candidates. In *NIPS*, 2015. 2, 3
- [37] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár. Learning to refine object segments. In *ECCV*, 2016. 2, 3
- [38] J. Pont-Tuset, P. Arbeláez, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping for image segmentation and object proposal generation. *PAMI*, 2016. 2
- [39] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *CVPR*, 2012. 5
- [40] X. Ren and J. Malik. Tracking as repeated figure/ground segmentation. In *CVPR*, 2007. 2
- [41] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *CVPR*, 2015. 4
- [42] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM Trans. Graphics*, 2004. 2

- [43] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 1
- [44] X. Shen, A. Hertzmann, J. Jia, S. Paris, B. Price, E. Shechtman, and I. Sachs. Automatic portrait segmentation for image stylization. In *Computer Graphics Forum*, 2016. 2
- [45] J. Shi, Q. Yan, L. Xu, and J. Jia. Hierarchical image saliency detection on extended CSSD. *TPAMI*, 2016. 4
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 4
- [47] T. V. Spina and A. X. Falcão. Fomtrace: Interactive video segmentation by image graphs and fuzzy object models. *arXiv preprint arXiv:1606.03369*, 2016. 2
- [48] D. Tsai, M. Flagg, and J. M. Rehg. Motion coherent tracking with multi-label MRF optimization. In *BMVC*, 2010. 2
- [49] Y.-H. Tsai, M.-H. Yang, and M. J. Black. Video segmentation via object flow. In *CVPR*, 2016. 1, 2, 5, 6, 7
- [50] H. Wang, T. Raiko, L. Lensu, T. Wang, and J. Karhunen. Semi-supervised domain adaptation for weakly labeled semantic video object segmentation. In *ACCV*, 2016. 2, 6
- [51] T. Wang, B. Han, and J. Collomosse. Touchcut: Fast image and video segmentation using single-touch interaction. *CVIU*, 2014. 2
- [52] L. Wen, D. Du, Z. Lei, S. Z. Li, and M.-H. Yang. Jots: Joint online tracking and segmentation. In *CVPR*, 2015. 2, 6
- [53] F. Xiao and Y. J. Lee. Track and segment: An iterative unsupervised approach for video object proposals. In *CVPR*, 2016. 1, 2, 6
- [54] N. Xu, B. Price, S. Cohen, J. Yang, and T. S. Huang. Deep interactive object selection. In *CVPR*, 2016. 2
- [55] D. Zhang, O. Javed, and M. Shah. Video object segmentation through spatially accurate and temporally dense extraction of primary object regions. In *CVPR*, 2013. 2
- [56] Y. Zhang, X. Chen, J. Li, C. Wang, and C. Xia. Semantic object segmentation via detection in weakly labeled video. In *CVPR*, 2015. 2, 6
- [57] F. Zhong, X. Qin, Q. Peng, and X. Meng. Discontinuity-aware video object cutout. *TOG*, 2012. 2