arXiv:2008.00992v2 [cs.CV] 13 Aug 2020

An Exploration of Target-Conditioned Segmentation Methods for Visual Object Trackers

Matteo Dunnhofer[0000−0002−1672−667X], Niki Martinel[0000−0002−6962−8643], and Christian Micheloni[0000−0003−4503−7483]

Machine Learning and Perception Lab, University of Udine, Italy

Abstract. Visual object tracking is the problem of predicting a target object's state in a video. Generally, bounding-boxes have been used to represent states, and a surge of effort has been spent by the community to produce efficient causal algorithms capable of locating targets with such representations. As the field is moving towards binary segmentation masks to define objects more precisely, in this paper we propose to extensively explore target-conditioned segmentation methods available in the computer vision community, in order to transform any bounding-box tracker into a segmentation tracker. Our analysis shows that such methods allow trackers to compete with recently proposed segmentation trackers, while performing quasi real-time.

Keywords: Visual Object Tracking, Video Object Segmentation, Target-Conditioned Segmentation, Deep Learning.

1 Introduction

In its simplest definition, visual object tracking corresponds to the prediction of the state of a target object in a stream of images. It is considered one of the most difficult problems in computer vision. Object occlusion and fast motion, light changes, and motion blur are some of the challenges that algorithms have to deal with. Additionally, constraints of real-time operation are often demanded by the many applications, such as video surveillance, behavior understanding, autonomous driving, and robotics.

Despite the many state representation one can use to model the target's state, the bounding-box has been the most used until now. This is a rectangle that encloses the object of interest, and it is defined by the coordinates of its top-left corner or of its center, and by its width and height. Based on this representation, many model-free tracking algorithms have been studied. Early solutions were based on mean shift algorithms [9], part-based methods [5,41], and SVM learning [21]. Later, the correlation filter approach gained interest thanks to its fast processing time [4,24,12,1,36]. More recently, the features of convolutional neural networks (CNNs) have been exploited to develop more efficient trackers. Trackers based on this image representation include deep regression networks

[23,20,17], online tracking-by-detection methods [40,26], solutions that use reinforcement learning [64,48,6,16,17], CNN-based discriminative correlation filters [13,10,11,3] and siamese CNNs [2,31,30,67,65,59]. The last two methods raised the state-of-the-art year-by-year, showing a remarkable performance across all the available tracking benchmarks [25,19,60,38,39,28,29] that almost reaches the 70% of bounding-box overlap accuracy. Such a high and generalized performance poses the question about if the bounding-box based measures have been now saturated. Moreover, it was proven in the object detection community that humans hardly distinguish a bounding-box prediction that has a 30% overlap from one with 50% [51]. Hence, it is fair to ask ourselves if a 100% overlap accuracy tracker is really necessary for applications. From such considerations, we can wonder if the time is done for bounding-box representations.

Furthermore, starting from 2020, tracking communities (VOT2020[1] and MOT[2]) raised the bar in their annual challenges by requesting trackers binary segmentation maps – precise location and shape definitions through pixel-wise background-target classifications – as target state representation. Segmentation representations are not new in the visual tracking panorama. In many applications, model-based algorithms used contours [62,32] or masks [37,27,15] for tracking particular objects. From a more general point of view, the recent video object segmentation (VOS) problem requires to produce the segmentation masks of generic target objects in a video, given the mask of each in the first frame. The currently available solutions propose highly accurate methods in terms of segmentation ability [55,42,56,34,63,57], but with the drawback of poor speed performance. This is due to characteristics of the available benchmarks [44,46,61], that do not include challenging situations from a tracking point of view. In fact, these datasets provide temporally short sequences where the target covers a large fraction of frames, its appearance does not suffer major changes, or low background distractors are present. The performance of such methods on standard VOT benchmarks was proven very poor [28] and to mitigate such behavior, the SiamMask [59] and D3S [35] algorithms have been proposed recently. These solutions adapted, respectively, the siamese correlation approach and discriminative correlation filters to segmentation outputs, and showed promising results while performing in real-time.

We believe that the huge effort spent by the tracking community in developing bounding-box based trackers can be still exploited in the segmentation tracking domain. With such an idea in mind, in this paper, we propose to explore what is currently available in the computer vision literature that can be adapted to make any bounding-box tracker output segmentation masks. In particular, we propose to extensively evaluate three methods: Box2Seg [57], SiamMask [59], and AMP [53]. Two were already proposed for this task [57,53], but their capabilities were little explored. The other is a recent segmentation tracker [59] that we reinterpret as a segmentation module. Our evaluations are based on a framework that requires a bounding-box tracker to provide a coarse localization of the object,

[1] `https://votchallenge.net/vot2020/`
[2] `https://motchallenge.net`

and then a segmentation module conditioned on the target object is employed to provide its precise localization. Along with practical considerations, we will show that this combination can produce trackers able to compete with the recently proposed methods [59,35] on the VOT2020 and DAVIS benchmarks.

1.1 Related Work

Combining segmentation methods and trackers has been increasingly tackled in the last two years. SiamMask [59] and D3S [35] employed a CNN decoder module [50,45] to refine a latent representation constructed by a cross-correlation operation and discriminative filter, respectively. Zhang *et al.* [66] proposed to use ECO tracker's [10] bounding-box predictions to improve the segmentation performance of the OSVOS [55] VOS method. Similarly, [57] adapted a deep CNN for semantic segmentation to generate a segmentation mask after the bounding-box proposal of a tracking-by-detection approach. The combination of these methods achieved promising results, but they were mainly focused on the VOS task. Additionally, they did not provide any extensive evaluation by considering different trackers and segmentation methods. In this paper, we aim to provide a deep analysis of such combination on both visual tracking and VOS benchmarks.

2 Methodology

In this paper, we study how state-of-the-art off-the-shelf bounding-box trackers can be augmented to track an object with the requirement of a segmentation representation. Our idea is based on the belief that the much effort spent in developing algorithms to predict the motion of a target is relevant even if a segmentation is required. To implement our analysis, we design a framework where a bounding-box tracker is first used to get a coarse localization of the target object, and then a target-conditioned segmentation method is executed to generate a pixel-wise map. Under this setup, any bounding-box based tracker can be transformed into a segmentation tracker. Considering separately tracking and mask generation carries practical advantages: (i) the performance of a segmentation tracker can be analyzed more consistently, by separating the error committed in the localization from the error in shape definition; (ii) flexibility of easily switch tracking and segmentation modules to adapt to application needs; (iii) availability of two different forms of output (bounding-box and mask) that are obtained with independent modules.

In the following of this section, we first introduce the framework employed for the analysis. Then, an abstract description of each of the selected segmentation methods will be given.

2.1 Segmentation Tracking Framework

We first define the key elements of the framework. A video

$$\mathcal{V} = \{F_t\}, t \in \{0, \cdots, T\}, T \in \mathbb{N} \tag{1}$$

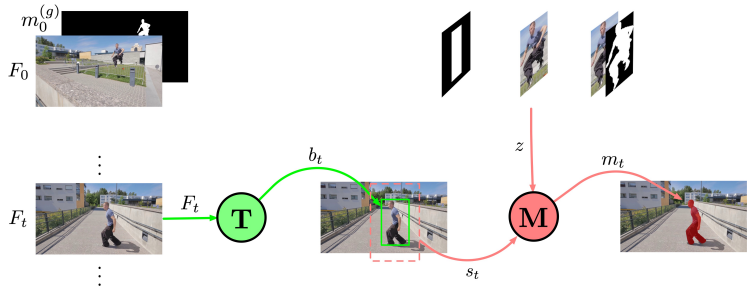

Fig. 1. Graphical representation of the framework used for the evaluation. $m_0^{(g)}$ outlines the target to be tracked in the first frame F_0 of a video. At every step t, the frame F_t is first given in input the the tracker $\mathbf{T}$ which outputs a bounding-box estimate b_t. This, together with a factor k, is employed to crop a searching area s_t in F_t. s_t is inputted to the mask generation algorithm $\mathbf{M}$ which is conditioned on the target template z computed in different form depending on $\mathbf{M}$'s input requirements. $\mathbf{M}$ returns the segmentation of the target inside s_t. The output mask m_t is finally built by placing $\mathbf{M}$'s output inside a zero-matrix at the location of s_t.

is considered as a T long sequence of frames $F_t \in \mathcal{I}$, where $\mathcal{I} = \{0, \cdots, 255\}^{W \times H \times 3}$ is the space of RGB images. We treat a bounding-box tracking algorithm as a function

$$\mathbf{T} : \mathcal{I} \rightarrow \mathbb{R}^4 \tag{2}$$

that is inputted with frame F_t and produces a bounding-box estimate $b_t = [x_t, y_t, w_t, h_t]$ as a real-valued vector containing the center coordinates x_t, y_t, and the width and height w_t, h_t (in the image coordinate system).[3] In a similar fashion, we consider a target-based segmentation algorithm as the function

$$\mathbf{M} : \mathcal{P} \times \mathcal{Z} \rightarrow \{0, 1\}^{W' \times H'} \tag{3}$$

which is given an image patch $s_t \in \mathcal{P} = \{0, \cdots, 255\}^{W' \times H' \times 3} \subseteq \mathcal{I}$ extracted from F_t and a template image $z \in \mathcal{Z}$ of the target object, and outputs a binary segmentation mask with zero-elements belonging to the background and one-elements defining the pixels of the target.

Given these concepts, the segmentation tracking procedure works as follows. At every time step t of a video $\mathcal{V}$, F_t is first given to the tracker $\mathbf{T}$ to produce b_t. Then, F_t and b_t are used to extract a searching area

$$s_t = F_t[x_t, y_t, k \cdot w_t, k \cdot h_t] \tag{4}$$

which is the area of F_t localized by the coordinates of b_t and which width and height are scaled by the factor $k \in \mathbb{R}$. s_t and z are given to the segmentation

[3] At $t = 0$, $\mathbf{T}$ is initialized with F_0 and the ground-truth bounding-box $b_0^{(g)}$.

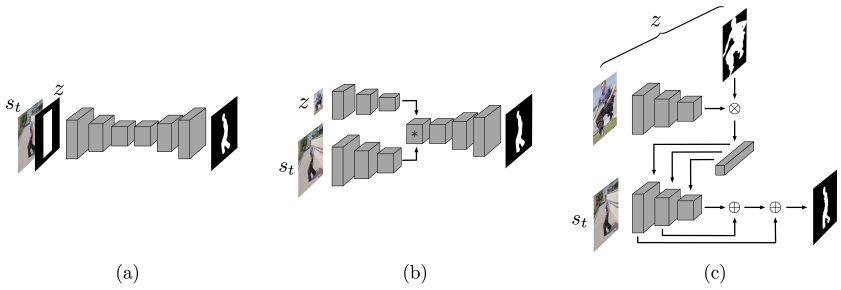

Fig. 2. Visual representation of the methodologies employed to generate target segmentations. (a) shows SemSeg, a deep segmentation model adapted to take in input a 4-channel tensor composed of s_t and z and produce the mask of the object. (b) presents SiamSeg, the siamese framework where a cross correlation operation between z and s_t features is employed to first locate the object, and then to produce its segmentation mask. (c) shows FewShotSeg, a few-shot segmentation algorithm that is adapted for visual segmentation tracking, by considering z as the support set and s_t as the query image.

algorithm $\mathbf{M}$ to produce the pixel-wise mask of the target inside s_t.[4] The output mask m_t is finally built by placing $\mathbf{M}$'s output at the s_t location of a zero-matrix with size $W \times H$.

A graphical representation of the described framework is shown in Figure 1.

2.2 Target-Conditioned Segmentation Methods

In this subsection we describe the target-conditioned segmentation methodologies we analyzed. Three conceptually different approaches were chosen:

- an adapted semantic segmentation network [57,7], that we name SemSeg;
- a module based on the siamese correlation framework [59], referred as SiamSeg;
- a few-shot segmentation algorithm [53], called FewShotSeg.

SemSeg. The first target-conditioned segmentation method we analyzed was proposed as Box2Seg in [34,57]. The idea is to adapt a state-of-the-art fully convolutional deep neural network for image segmentation to target segmentation. Given an RGB image and an additional input channel containing coarse information regarding the position of the target, this module produces a detailed segmentation of the latter. In the context of our framework, the RGB channels of the searching area s_t are concatenated with the template channel

$$z = \{0, 1\}^{k \cdot w_t, k \cdot h_t} \tag{5}$$

[4] The details to obtain z are described in each subsection describing the segmentation methods.

which is a binary mask of the same size of s_t, and which positive elements are located inside the area defined by b_t. z is computed at every time step t, and the 4-channel input resulting from the concatenation is given to the network which produces the segmentation of the target inside the searching area. A visualization of this approach is proposed in Figure 2 (a). The network is trained offline by exploiting object segmentation, instance segmentation, and/or VOS datasets. The training pairs are formed as batches of inputs-targets, where the first are composed by searching area and template (built using the bounding-box that encloses the ground-truth segmentation), and the second are the actual object masks. Optimization is done by solving a two class segmentation problem (foreground-background) defined as the minimization of a pixel-wise classification loss (cross-entropy, Dice loss, etc.). This approach has the advantage of requiring just the bounding-box as first-frame target definition.

SiamSeg. As second mask generation method, we reinterpreted the siamese correlation framework for segmentation tracking [59]. The general view of this scheme is to first locate the target template in the higher-level feature space of template and searching area, and then project the localization into the segmentation space. These steps are jointly implemented with an encoder-decoder CNN architecture, which capabilities are acquired through an end-to-end offline procedure in which the whole model is optimized by minimizing a foreground-background pixel-wise classification loss. The training examples are pairs of searching area-template inputs, and ground-truth target masks, where searching area and template are sampled without temporal correlation.

Following this intuition, we adapted such method in our framework as follows. The target template is the image crop

$$z = F_0[x_0, y_0, w_0, h_0] \tag{6}$$

extracted from the first frame F_0 of the video, using the ground-truth bounding-box $b_0^{(g)} = [x_0, y_0, w_0, h_0]$. $b_0^{(g)}$ is obtained as the box that encloses the ground-truth mask $m_0^{(g)}$. The features $\widehat{z}$ of z are computed with a forward pass through the encoder module just at $t = 0$. At every other t, $\widehat{z}$ is cross-correlated to the encoded representation $\widehat{s_t}$, and the resulting activation map is then refined by the decoder module, and ultimately placed into m_t. The procedure is depicted in Figure 2 (b) and, as for SemSeg, it just requires the target definition as a bounding-box.

FewShotSeg. The last analyzed methodology treats target-conditioned segmentation as a few-shot segmentation problem. In such a setting, the goal is to provide a pixel-wise segmentation of a target object inside a query image, given a so-called support-set, i.e. one or more (few-shot) image and mask examples of the target. Algorithms for this problem are generally designed as fully CNNs, where the segmentation ability is guided by other convolutional branches or by model parameters that are made dependent on the support-set.

This view of few-shot segmentation can be reframed for the purpose of segmentation tracking. In our setting, the support-set is considered as the target template

$$z = (F_0[x_0, y_0, w_0, h_0], m_0^{(g)}[x_0, y_0, w_0, h_0]) \tag{7}$$

that is the pair of the image crop that contains the visual appearance of the target in F_0, and the relative cropped ground-truth mask. The crops are constructed considering $b_0^{(g)} = [x_0, y_0, w_0, h_0]$. The searching area s_t is extracted after every b_t of $\mathbf{T}$ and it is considered as the query image. Together with the template (the support-set), they are given to the few-shot segmentation model to produce the target segmentation. A graphical example of this methodology is proposed in Figure 2 (c). With respect to the previous methods, employing FewShotSeg requires the definition of the target object through a mask.

3 Experimental Setup

In this section, we report the experimental procedures we performed to implement and analyze the previously presented methodologies. All experiments were run on a machine with an Intel Xeon E5-2690 v4 @ 2.60GHz CPU, 320 GB of RAM, and an NVIDIA TITAN V GPU. Code for tracker and segmentation methods was implemented in Python.

3.1 Trackers

The trackers selected for the analysis were KCF [24], DCFNet [58], MDNet [40], MetaCrest [43], SiamFC [2], SiamRPN [31], ECO [10], ATOM [11], and DiMP [3]. Such algorithms were chosen because they tackle visual tracking by different approaches and so can provide performance of various quality. For each of them, we used the public code made available by the authors. We tried the best to respect default parameters and settings.

3.2 Segmentation Modules

SemSeg. To implement this methodology, we followed the details of the Box2Seg refinement module provided in [57,34]. The DeepLab-v3 architecture [7] for image segmentation was translated for the task of interest. ImageNet [14] pre-trained ResNet-50 [22] was employed as backbone network and adapted to receive the 4-channel tensor. Before being inputted, the concatenated RGB and template channels were resized to 385×385 pixels. During training, the searching area was enlarged by the factor k, chosen uniformly in $\{1, 1.25, 1.5, 1.75, 2\}$. Batches of 12 input-target mask pairs were sampled from a training set composed of the training sets of COCO [33], YouTube-VOS [61], and DAVIS 2016 and 2017 [44,46]. Learning rate was set to 10^{-5} for the backbone layers, and to 10^{-4} for all the others. Training was carried on until the mIoU [18], computed over the foreground and background classes, stopped improving on a custom validation set composed of the validation sets of the aforementioned datasets.

SiamSeg. The second approach introduced in subsection 2.2 was implemented through the segmentation tracker SiamMask [59]. We used the code provided by the authors along with the pre-trained models. For completion, we present to the reader some information about the training procedures performed by the authors. The SiamMask architecture model was trained in two-stages: first, the encoder module based on ResNet-50 [22] was trained for target localization by optimizing a multi-task loss for similarity maximization and RPN [49] detection. After that, the decoder module designed as [45] was attached to the intermediate cross-correlation map and trained by minimizing a foreground-background pixel-wise cross-entropy loss. The training set used was a combination of ImageNet-VID [14], COCO [33] and YouTube-VOS [61]. Before being inputted to the model, z and s_t were resized to 127×127 and 255×255 pixels respectively.

FewShotSeg. As a few-shot segmentation module, we employed the strategy proposed in [53], which is a recently introduced state-of-the-art method that has been shown to perform well also in VOS tasks. The authors proposed a sample efficient method to segment an unseen class object via a multi-resolution imprinting procedure of adaptive masked proxies (AMP). AMPs are constructed by a Normalized Masked Average Pooling (NMAP) operation between the CNN embeddings of the support set's RGB sample and its relative binary mask. The AMP representations are used to imprint [47], at multiple resolutions, the CNN embeddings computed on the query image. The VGG16 [54] architecture is employed as a backbone feature extractor, and skip connections are also exploited as done similarly in FCN8s [52]. Data extracted from the PASCAL-VOC dataset [18] was used to compose training samples as query image, support-set image, support-set mask, and target mask. Optimization was performed by minimizing the pixel-wise cross-entropy loss between predicted and ground-truth masks. Code and pre-trained model provided by the authors were adapted to our implementation needs.

3.3 Benchmarks and Performance Measures

We performed analysis on the VOT2020 benchmark, and the validation sets of the DAVIS 2016 [44] and DAVIS 2017 [46] VOS benchmarks. All provide segmentations as target representations.

For VOT2020 we employed the newly introduced protocol.[5] The novel baseline protocol requires running a tracker on shorter sequences determined by predefined points (anchors). From such starting points, the tracker is initialized with the ground-truth mask and run either forward or backward, depending on the longest sub-sequence yielded by the two directions. The new accuracy ($A_{\uparrow}$) measures the average pixel-wise intersection-over-union between predicted and ground-truth masks, for frames where the tracker did not fail (i.e. the accuracy did not decrease after a certain threshold). The new robustness ($R_{\uparrow}$) expresses the normalized average number of frames where the algorithm successfully tracked the target

[5] `https://data.votchallenge.net/vot2020/vot-2020-protocol.pdf`

Table 1. Results of the baseline experiment on VOT2020. Best segmentation method results, per tracker, are highlighted in red (Rectangular Mask results are excluded).

Tracker	SemSeg			SiamSeg			FewShotSeg			Rectangular Mask		
	$EAO_\uparrow$	$A_\uparrow$	$R_\uparrow$	$EAO_\uparrow$	$A_\uparrow$	$R_\uparrow$	$EAO_\uparrow$	$A_\uparrow$	$R_\uparrow$	$EAO_\uparrow$	$A_\uparrow$	$R_\uparrow$
DCFNet	0.203	0.616	0.426	**0.310**	**0.676**	0.558	0.230	0.491	**0.567**	0.184	0.441	0.523
KCF	0.199	0.648	0.371	**0.285**	**0.659**	**0.501**	0.200	0.459	0.499	0.155	0.402	0.432
SiamFC	0.218	0.602	0.446	**0.309**	**0.682**	**0.571**	0.228	0.491	0.563	0.183	0.418	0.537
MetaCrest	0.240	0.602	0.513	**0.336**	**0.657**	0.624	0.250	0.479	**0.647**	0.189	0.390	0.587
SiamRPN	0.356	0.692	0.639	**0.369**	**0.701**	0.651	0.311	0.551	**0.677**	0.247	0.452	0.663
MDNet	0.295	0.638	0.609	**0.371**	**0.662**	0.689	0.308	0.546	**0.723**	0.234	0.440	0.687
ATOM	0.402	0.678	**0.735**	**0.406**	**0.691**	0.723	0.337	0.560	0.731	0.277	0.467	0.738
DiMP	0.410	0.675	0.744	**0.410**	**0.691**	0.730	0.347	0.556	**0.749**	0.278	0.464	0.733
ECO	0.322	0.632	0.735	**0.414**	**0.694**	0.729	0.349	0.561	**0.759**	0.275	0.459	0.746
b-oracle	**0.806**	**0.809**	**0.996**	0.697	0.744	0.970	0.541	0.623	0.941	0.516	0.519	1.0

before drifting. The two measures are joined in a refreshed single performance score known as expected average overlap ($EAO_\uparrow$). Version 0.4.2 of the Python toolkit was used to obtain the results.

The protocol used for DAVIS datasets is similar to the One-Pass evaluation (OPE) employed in OTB [60] benchmarks: the tracker is initialized with the mask of the target object in the first frame, and then it is run until the end of the sequence. Performance is measured in terms of the Jaccard index $\mathcal{J}$ which measures the pixel-wise intersection-over-union between the predicted and ground-truth masks. Along with this index, the F-measure $\mathcal{F}$ is employed to evaluate contour accuracy. For both measures, mean ($\mathcal{J}_{\mathcal{M}\uparrow}, \mathcal{F}_{\mathcal{M}\uparrow}$), recall ($\mathcal{J}_{\mathcal{R}\uparrow}, \mathcal{F}_{\mathcal{R}\uparrow}$), and decay ($\mathcal{J}_{\mathcal{D}\downarrow}, \mathcal{F}_{\mathcal{D}\downarrow}$) values are reported. For DAVIS 2017, where multiple objects must be tracked and segmented, we run the trackers independently for each object and then fuse the prediction masks by assigning each pixel to the object that received higher confidence in that location.

4 Results

General Performance. Results on VOT2020 benchmark are presented in Table 1. Trackers combined with SiamSeg achieve the best overall performance in $EAO_\uparrow$ and $A_\uparrow$. This is explained by the fact the VOT benchmarks include difficult tracking scenarios for trackers, resulting in lower quality bounding-boxes that affect SemSeg and FewShotSeg. Thanks to its more robust segmentation method, SiamSeg allows to recover (to some extent) from inaccurate b_t estimates and so produce more accurate target segmentations. Interestingly, FewShotSeg is the approach that achieves the highest $R_\uparrow$, showing to be the method less susceptible to failure. For all the methods, employing a better tracker is fundamental to improve the overall performance.

Results on the DAVIS 2016 benchmark are reported in Tables 2 and 3. More weak trackers like DCFNet, KCF, MDNet, MetaCrest, and SiamFC, benefit of FewShotSeg for pixel-wise accuracy. When more precise bounding-box estimates are provided, through ECO, SiamRPN, ATOM, DiMP, SemSeg allows the best $\mathcal{J}_{\mathcal{M}\uparrow}$ performance. For $\mathcal{J}_{\mathcal{R}\uparrow}$ and $\mathcal{J}_{\mathcal{D}\downarrow}$, FewShotSeg is almost always the best approach. For contour accuracy, SemSeg is generally the best method at $\mathcal{F}_{\mathcal{M}\uparrow}$. Better trackers also benefit the same for $\mathcal{F}_{\mathcal{R}\uparrow}$ and $\mathcal{F}_{\mathcal{D}\downarrow}$. For the others, FewShotSeg

Table 2. $\mathcal{J}$ results on DAVIS 2016 validation set. Best segmentation method results, per tracker, are highlighted in red (Rectangular Mask results are excluded).

Tracker	SemSeg			SiamSeg			FewShotSeg			Rectangular Mask		
	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$
KCF	0.527	0.570	0.174	0.557	0.616	0.199	**0.580**	**0.688**	**0.162**	0.302	0.200	0.153
DCFNet	0.531	0.574	0.209	0.551	0.627	0.229	**0.564**	**0.674**	**0.178**	0.313	0.183	0.130
MetaCrest	0.574	0.624	0.169	0.595	0.672	0.145	**0.598**	**0.712**	**0.136**	0.323	0.151	0.108
MDNet	0.582	0.635	0.177	0.593	0.656	0.196	**0.610**	**0.717**	**0.143**	0.342	0.198	0.149
SiamFC	0.607	0.661	**0.159**	0.611	0.694	0.177	**0.621**	**0.738**	0.163	0.356	0.234	0.140
ECO	0.615	0.679	**0.099**	0.623	0.744	0.108	**0.626**	**0.748**	0.113	0.375	0.243	0.070
SiamRPN	**0.689**	0.772	**0.089**	0.663	0.782	0.111	0.681	**0.859**	**0.089**	0.417	0.340	0.066
ATOM	**0.723**	**0.846**	**0.074**	0.658	0.785	0.105	0.669	0.845	0.081	0.415	0.345	0.053
DiMP	**0.723**	0.827	**0.086**	0.704	0.844	0.100	0.699	**0.886**	0.095	0.443	0.379	0.027
b-oracle	**0.812**	0.920	**0.020**	0.732	0.896	0.044	0.739	**0.946**	0.052	0.455	0.418	0.008

Table 3. $\mathcal{F}$ results on DAVIS 2016 validation set. Best segmentation method results, per tracker, are highlighted in red (Rectangular Mask results are excluded).

Tracker	SemSeg			SiamSeg			FewShotSeg			Rectangular Mask		
	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$
DCFNet	**0.553**	0.587	0.210	0.536	0.596	0.187	0.530	**0.599**	**0.158**	0.155	0.017	0.068
KCF	**0.559**	0.577	0.180	0.525	0.572	0.190	0.542	**0.598**	**0.149**	0.136	0.018	0.119
MetaCrest	**0.599**	0.632	0.193	0.561	**0.637**	**0.136**	0.572	0.634	0.137	0.139	0.019	0.063
MDNet	**0.603**	0.623	**0.170**	0.570	0.616	0.197	0.582	**0.635**	**0.170**	0.163	0.050	0.112
SiamFC	**0.633**	0.665	**0.152**	0.592	0.663	0.159	0.597	**0.675**	0.157	0.156	0.037	0.126
ECO	**0.637**	**0.696**	**0.097**	0.590	0.692	0.102	0.592	0.673	0.117	0.170	0.020	0.066
SiamRPN	**0.713**	**0.783**	**0.105**	0.629	0.707	0.127	0.642	0.752	**0.105**	0.186	0.059	0.081
ATOM	**0.739**	**0.856**	0.098	0.628	0.697	0.111	0.626	0.751	**0.090**	0.178	0.025	0.060
DiMP	**0.744**	**0.821**	**0.108**	0.658	0.754	0.130	0.657	0.767	0.118	0.191	0.071	0.015
b-oracle	**0.843**	**0.918**	**0.033**	0.693	0.805	0.064	0.717	0.873	0.056	0.219	0.073	0.015

gets the best results. SiamSeg is the weakest method on this benchmark, justified by the presence of easy tracking situations that put the focus on providing more accurate target segmentations.

On DAVIS 2017, which results are presented in Tables 4 and 5, SemSeg is still the best approach to use with stronger bounding-box trackers for $\mathcal{J}_{\mathcal{M}\uparrow}$ and $\mathcal{J}_{\mathcal{R}\uparrow}$. FewShotSeg is the method that achieves the most consistent masks across time. For low-performance tracking algorithms, SiamSeg results to be better than the others in $\mathcal{J}_{\mathcal{M}\uparrow}$ and $\mathcal{J}_{\mathcal{R}\uparrow}$, mitigating the lower tracking performance with its target search strategy and showing the increased difficulty of this benchmark than its previous version. In terms of contour performance, SemSeg is the most appropriate method for $\mathcal{F}_{\mathcal{M}\uparrow}$ and $\mathcal{F}_{\mathcal{R}\uparrow}$ performance. For $\mathcal{F}_{\mathcal{D}\downarrow}$, FewShotSeg results in the best solution. Overall, as for VOT2020, in both DAVIS 2016 and 2017 employing better trackers lets achieve the best performances.

Comparison with a Rectangular Segmentation Tracker. In the last block of columns of Tables 1, 2, 3, 4, 5, we report the performance of the trackers considering their b_t predictions as segmentation masks, i.e. binary mask where the rectangular area defined by b_t is filled with 1. Overall, all the considered segmentation methods improve those baseline results on all the benchmarks and across all measures. This proves that employing the approaches presented in this paper lets bounding-box trackers improve their accuracy in terms of precise target definition. SemSeg is the method that achieves generally the best improvement, followed by SiamSeg and FewShotSeg.

Table 4. $\mathcal{J}$ results on DAVIS 2017 validation set. Best segmentation method results, per tracker, are highlighted in red (Rectangular Mask results are excluded).

Tracker	SemSeg			SiamSeg			FewShotSeg			Rectangular Mask		
	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$	$\mathcal{J}_{\mathcal{M}\uparrow}$	$\mathcal{J}_{\mathcal{R}\uparrow}$	$\mathcal{J}_{\mathcal{D}\downarrow}$
DCFNet	0.443	0.474	0.299	**0.455**	**0.497**	0.281	0.424	0.434	**0.214**	0.283	0.166	0.176
KCF	0.433	0.464	0.277	**0.461**	**0.517**	0.272	0.425	0.451	**0.209**	0.268	0.167	0.198
MDNet	0.444	0.478	0.284	**0.465**	**0.515**	0.260	0.444	0.493	**0.216**	0.284	0.156	0.168
MetaCrest	0.447	0.468	0.276	**0.468**	**0.518**	0.262	0.426	0.443	**0.178**	0.273	0.145	0.155
SiamFC	0.466	0.499	0.260	**0.468**	**0.523**	0.277	0.431	0.454	**0.225**	0.280	0.176	0.196
ECO	**0.498**	0.556	0.244	0.503	**0.567**	0.222	0.458	0.501	**0.178**	0.310	0.220	0.132
SiamRPN	**0.536**	**0.600**	0.233	0.506	0.578	0.237	0.470	0.518	**0.180**	0.321	0.248	0.141
ATOM	**0.566**	**0.659**	**0.148**	0.544	0.626	0.188	0.488	0.547	0.168	0.321	0.251	0.103
DiMP	**0.583**	**0.671**	**0.148**	0.553	0.639	0.170	0.498	0.555	0.162	0.323	0.251	0.093
b-oracle	**0.762**	**0.891**	**0.0**	0.618	0.738	0.073	0.578	0.694	0.059	0.408	0.340	0.0

Table 5. $\mathcal{F}$ results on DAVIS 2017 validation set. Best segmentation method results, per tracker, are highlighted in red (Rectangular Mask results are excluded).

Tracker	SemSeg			SiamSeg			FewShotSeg			Rectangular Mask		
	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$	$\mathcal{F}_{\mathcal{M}\uparrow}$	$\mathcal{F}_{\mathcal{R}\uparrow}$	$\mathcal{F}_{\mathcal{D}\downarrow}$
KCF	**0.517**	0.542	0.307	0.500	0.544	0.287	0.506	**0.565**	**0.266**	0.172	0.035	0.178
DCFNet	**0.532**	0.567	0.322	0.512	0.561	0.284	0.511	**0.572**	**0.237**	0.194	0.049	0.140
MDNet	0.525	0.563	0.288	0.513	0.565	0.270	**0.526**	**0.598**	**0.255**	0.184	0.059	0.170
MetaCrest	**0.545**	**0.593**	0.305	0.520	0.567	0.278	0.521	0.583	**0.241**	0.176	0.042	0.146
SiamFC	**0.556**	**0.611**	0.299	0.523	0.583	0.294	0.524	0.601	**0.267**	0.184	0.064	0.184
ECO	**0.592**	**0.663**	0.255	0.553	0.620	0.236	0.553	0.637	**0.226**	0.214	0.055	0.128
SiamRPN	**0.626**	**0.713**	0.259	0.552	0.628	0.257	0.567	0.670	**0.219**	0.210	0.070	0.150
DiMP	**0.663**	**0.765**	**0.181**	0.591	0.675	0.215	0.584	0.685	0.195	0.206	0.059	0.093
ATOM	**0.640**	**0.751**	**0.195**	0.584	0.664	0.218	0.574	0.666	0.201	0.203	0.053	0.104
b-oracle	**0.829**	**0.945**	**0.017**	0.654	0.779	0.097	0.685	0.847	0.078	0.280	0.116	0.028

Comparison with a Bounding-box Oracle Tracker. In the last row of Tables 1, 2, 3, 4, 5, the performance of a bounding-box oracle based tracker, *b*-oracle (i.e. the tracker that returns the ground-truth bounding-box $b_t^{(g)}$ at every t), is presented. Given this ground-truth information, SemSeg is the approach that best segments the target object, on every considered benchmark and performance measure. On VOT2020, accuracy and robustness performances reach almost 80% and 100%, meaning that its segmentation capabilities are effective for the objects contained in this dataset. SiamSeg follows with a decrease of 9% and 2.6%, while FewShotSeg shows a much bigger performance loss in $A_\uparrow$ (-25% than SemSeg) than in $R_\uparrow$ (-5.5%). FewShotSeg comes after SemSeg in terms of $\mathcal{J}$ and $\mathcal{F}$ on DAVIS 2016, and in terms of $\mathcal{F}$ on DAVIS 2017. SiamSeg gets the weakest performance on DAVIS 2016 but surpasses FewShotSeg in $\mathcal{J}$ on DAVIS 2017.

SemSeg is also the method that suffers the major gap between the *b*-oracle performance and the best tracker DiMP ($EAO_\uparrow$ loss -48%, average $\mathcal{J}_{\mathcal{M}\uparrow}$ loss -17.2%, average $\mathcal{F}_{\mathcal{M}\uparrow}$ loss -15.9%). This shows the susceptibility to misaligned bounding-box predictions (we hypothesize this can be mitigated introducing some noise to the input bounding-boxes in the training procedure). The performance decrease happens also for the other methods, although with less magnitude.

Separating Localization and Segmentation Error. The results obtained with *b*-oracle and the rectangular mask output allow us to determine the tracking and segmentation error committed by **T** and **M** respectively. The error $e_{\mathbf{T}}$ committed by the tracker is just the performance difference between *b*-oracle and **T**, both considered with rectangular mask output. The error $e_{\mathbf{M}}$ of **M** can be computed as the performance difference between *b*-oracle with **M** and **T** with

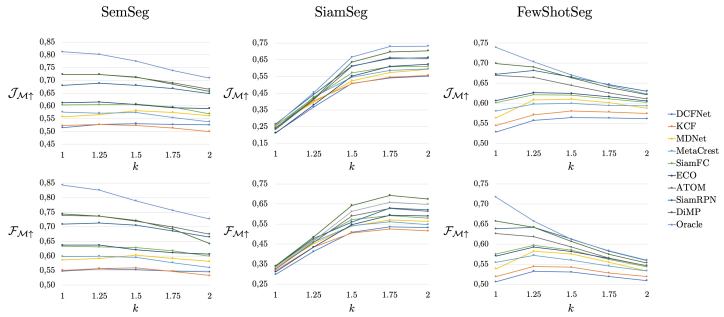

Fig. 3. Results on the DAVIS 2016 validation set of the sensibility of the target segmentation methods to the size of the searching area. Performance is evaluated in terms of $\mathcal{J}_{\mathcal{M}\uparrow}$ and $\mathcal{F}_{\mathcal{M}\uparrow}$.

M which tracking performance is corrected by summing $e_{\mathbf{T}}$. In this setting, $e_{\mathbf{M}}$ is considered as the distance from **M**'s maximum achievable performance, that happens when b-oracle is employed as tracker. For example, when MDNet and SemSeg are executed together, the $\text{A}_\uparrow$ error $e_{\mathbf{T}}$ is computed as $e_{\mathbf{T}} = 0.519 - 0.440 = 0.079$, while the $e_{\mathbf{M}}$ is obtained as $e_{\mathbf{M}} = 0.809 - (0.638 + 0.079) = 0.092$. So, it results that the highest loss in accuracy is due to the segmentation than to tracking. If DiMP and SiamSeg are considered, we have an $\text{A}_\uparrow$ error $e_{\mathbf{T}} = 0.055$ and $e_{\mathbf{T}} = -0.002$ meaning that SiamSeg compensates the tracking error and even improves the performance of the combination.

Impact of the Searching Area Size. We analyzed how sensible the three segmentation methods are to different sizes of the searching area. In particular, the factor k was studied across the values $\{1, 1.25, 1.5, 1.75, 2\}$ on the DAVIS 2016 benchmark, and the results are shown in Figure 3. With SemSeg, all the trackers show a slow decrease in $\mathcal{J}_{\mathcal{M}\uparrow}$ and $\mathcal{F}_{\mathcal{M}\uparrow}$ performance by enlarging the searching area. Best performance are obtained with $k = 1$ or $k = 1.25$ (proven also by the b-oracle based tracker). Similar conclusions can be made for FewShotSeg. The highest $\mathcal{J}_{\mathcal{M}\uparrow}$ is achieved with $k = 1.25$. For larger k, the performance of more weak trackers remains constant, while the performance of stronger trackers slightly decreases. $\mathcal{F}_{\mathcal{M}\uparrow}$ tends to decrease for all the trackers. SiamSeg shows the opposite trend. Better results are obtained with larger searching areas. Specifically, best $\mathcal{J}_{\mathcal{M}\uparrow}$ and $\mathcal{F}_{\mathcal{M}\uparrow}$ performance are obtained with $k \geq 1.75$. Weaker trackers have a smaller performance decrease between 1.75 and 1.5 than stronger ones, while for $k < 1.5$ the performance of all the trackers quickly drops. This can be explained by SiamSeg's training methodology, where the objective is set as target localization and segmentation in large image patches.

Speed Analysis. In Table 6 an analysis of the speed of the algorithms is presented. The fastest method to produce a segmentation is SiamSeg which runs

Table 6. Results on the speed analysis (in seconds and FPS) of the combined tracker-segmentation methods. The original tracker speeds are reported in the last two columns (times of the employed implementations). The last row shows the average speed of running just the segmentation methods.

Tracker	SemSeg		SiamSeg		FewShotSeg		Tracker speed	
	s	FPS	s	FPS	s	FPS	s	FPS
MDNet	0.628	1.6	0.580	1.7	0.704	1.4	0.550	1.8
MetaCrest	0.181	5.5	0.134	7.5	0.258	3.9	0.109	9.1
ECO	0.126	8	0.081	12.4	0.220	4.6	0.059	17.0
ATOM	0.123	8.1	0.079	12.7	0.198	5.1	0.050	20.0
DiMP	0.109	9.1	0.068	14.7	0.189	5.3	0.038	26.3
SiamRPN	0.026	12.1	0.047	21.4	0.172	5.8	0.026	38.4
KCF	0.070	14.3	0.034	29.5	0.167	6.0	0.013	78.8
SiamFC	0.064	15.6	0.030	33.5	0.157	6.4	0.008	125.3
DCFNet	0.062	16.2	0.026	39.0	0.143	7.0	0.004	227.8
No tracker	0.062	16.3	0.024	42.8	0.151	6.7	-	-

Table 7. State-of-the-art comparison for the best combinations. Best results are highlighted in red, second-best in blue.

Method	DAVIS 2016				DAVIS 2017				VOT2020			FPS
	$\mathcal{J}_{\mathcal{M}}\uparrow$	$\mathcal{J}_{\mathcal{R}}\uparrow$	$\mathcal{F}_{\mathcal{M}}\uparrow$	$\mathcal{F}_{\mathcal{R}}\uparrow$	$\mathcal{J}_{\mathcal{M}}\uparrow$	$\mathcal{J}_{\mathcal{R}}\uparrow$	$\mathcal{F}_{\mathcal{M}}\uparrow$	$\mathcal{F}_{\mathcal{R}}\uparrow$	EAO$\uparrow$	A$\uparrow$	R$\uparrow$	
OSMN [63]	0.740	0.876	0.729	0.840	0.525	0.609	0.571	0.661	-	-	-	7
BoLTVOS [57]	0.781	-	0.812	-	**0.684**	-	**0.754**	-	-	-	-	1
OSVOS [55]	0.798	0.936	0.806	**0.926**	0.566	**0.638**	0.639	**0.738**	-	-	-	0.1
FAVOS [8]	0.824	**0.965**	0.795	0.894	0.546	0.611	0.618	0.723	-	-	-	0.8
RGMP [42]	0.815	0.917	0.820	0.908	0.648	-	0.686	-	-	-	-	8
OnAVOS [56]	**0.857**	-	**0.842**	-	0.610	-	0.661	-	-	-	-	0.1
PReMVOS [34]	**0.849**	**0.961**	**0.886**	**0.947**	**0.739**	**0.831**	**0.817**	**0.889**	-	-	-	0.03
SiamMask[6]	0.692	0.848	0.639	0.743	0.522	0.597	0.559	0.645	0.321	0.686	0.569	43
D3S [35]	0.754	-	0.726	-	0.578	-	0.638	-	**0.439**	**0.699**	**0.769**	25
SiamRPN-SiamSeg	0.663	0.782	0.629	0.707	0.506	0.578	0.552	0.628	0.369	**0.701**	0.651	21
ECO-SiamSeg	0.623	0.744	0.590	0.692	0.503	0.567	0.553	0.620	**0.414**	0.694	0.729	12
ATOM-SiamSeg	0.658	0.785	0.628	0.697	0.544	0.626	0.584	0.664	0.406	0.691	0.723	13
ATOM-SemSeg	0.723	0.846	0.739	0.856	0.566	0.659	0.640	0.751	0.402	0.678	0.735	8
DiMP-SemSeg	0.723	0.827	0.744	0.821	0.583	0.671	0.663	0.765	0.410	0.675	**0.744**	9
DiMP-SiamSeg	0.704	0.844	0.658	0.754	0.553	0.639	0.591	0.675	0.410	0.691	0.730	15

at 43 FPS. With this method, DCFNet and SiamFC run in real-time (39 and 34 FPS respectively). Stronger trackers like ECO, ATOM, and DiMP, achieve a speed of 12, 13, and 15 FPS respectively. SemSeg runs independently at 16 FPS, and combined with SiamRPN and DiMP allows a speed of 12 and 9 FPS respectively. FewShotSeg is the slowest method and takes around 7 FPS. In this setup, the speed performance is almost completely taken by the segmentation method and best trackers reach a speed of 5-6 FPS.

State-of-the-art Comparison. Comparison with the state-of-the-art is presented in Table 7. The VOS methods outperform every studied $\mathbf{T}-\mathbf{M}$ combination on DAVIS 2016 and 2017, but they show poor speed results. DiMP and ATOM with SemSeg perform better than SiamMask in $\mathcal{J}$ on both DAVIS 2016 and 2017. In terms of $\mathcal{F}$ they outperform also D3S. On DAVIS 2017, D3S is improved by DiMP-SemSeg in every measure. On VOT2020, SiamMask is largely beaten by all the best trackers, combined both with SemSeg and SiamSeg. All the trackers using the second method improves SiamMask, showing its limitations in target localization. ECO and SiamSeg reaches an EAO$_\uparrow$ of 0.414, slightly improving DiMP and ATOM. With the same segmentation method, SiamRPN outperforms D3S in A$_\uparrow$, achieving the best 0.701, while maintaining a quasi real-time speed of 21 FPS.

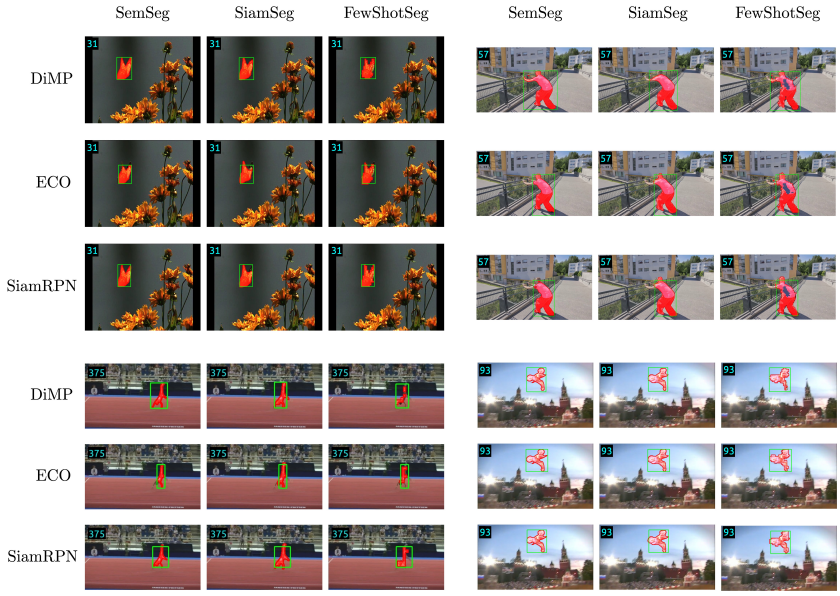

Fig. 4. Qualitative examples of the segmentation (red superimposed mask) proposed by the three target-conditioned segmentation methods, based on the bounding-box proposals (green rectangles) given by three different trackers.

In Figure 4 some qualitative examples of the segmentation methods are proposed.[7]

5 Conclusions

In this paper, three target-conditioned segmentation methods, SemSeg, SiamSeg, and FewShotSeg, were extensively analyzed to transform any bounding-box tracker into a segmentation tracker. SemSeg and SiamSeg resulted in the stronger methods, and their combination with trackers like SiamRPN, ECO, ATOM, and DiMP, allows to compete with the most recent segmentation trackers SiamMask and D3S on the DAVIS 2016 and 2017, and VOT2020 benchmarks.

Acknowledgements. This work was supported by the ACHIEVE-ITN H2020 project.

[6] Since we used SiamMask to implement SiamSeg, for fair comparison we report the results of the same implementation used for segmentation tracking, which has slightly worse performance than presented in the original paper.

[7] For more, please see `https://youtu.be/SODiKBD84_g`.

References

1. Bertinetto, L., Valmadre, J., Golodetz, S., Miksik, O., Torr, P.H.: Staple: Complementary learners for real-time tracking. In: IEEE Conference on Computer Vision and Pattern Recognition (2016)
2. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. European Conference on Computer Vision (2016)
3. Bhat, G., Danelljan, M., Van Gool, L., Timofte, R.: Learning Discriminative Model Prediction for Tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2019)
4. Bolme, D.S., Beveridge, J.R., Draper, B.A., Lui, Y.M.: Visual object tracking using adaptive correlation filters. In: IEEE Conference on Computer Vision and Pattern Recognition. IEEE (2010)
5. Čehovin, L., Kristan, M., Leonardis, A.: Robust visual tracking using an adaptive coupled-layer visual model. IEEE Transactions on Pattern Analysis and Machine Intelligence (2013)
6. Chen, B., Wang, D., Li, P., Wang, S., Lu, H.: Real-time 'Actor-Critic' Tracking. In: European Conference on Computer Vision (2018)
7. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking Atrous Convolution for Semantic Image Segmentation (jun 2017)
8. Cheng, J., Tsai, Y.H., Hung, W.C., Wang, S., Yang, M.H.: Fast and Accurate Online Video Object Segmentation via Tracking Parts. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE Computer Society (dec 2018)
9. Comaniciu, D., Ramesh, V., Meer, P.: Real-time tracking of non-rigid objects using mean shift. IEEE Conference on Computer Vision and Pattern Recognition (2000)
10. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: ECO: Efficient Convolution Operators for Tracking. In: IEEE Conference on Computer Vision and Pattern Recognition (nov 2017)
11. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: ATOM: Accurate Tracking by Overlap Maximization. In: IEEE Conference on Computer Vision and Pattern Recognition (2019)
12. Danelljan, M., Hager, G., Khan, F.S., Felsberg, M.: Discriminative Scale Space Tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence (2017)
13. Danelljan, M., Robinson, A., Khan, F.S., Felsberg, M.: Beyond Correlation Filters: Learning Continuous Convolution Operators for Visual Tracking. In: European Conference on Computer Vision (aug 2016)
14. Deng, J., Dong, W., Socher, R., Li, L.J., Kai Li, Li Fei-Fei: ImageNet: A large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition. IEEE (jun 2009)
15. Dunnhofer, M., Antico, M., Sasazawa, F., Takeda, Y., Camps, S., Martinel, N., Micheloni, C., Carneiro, G., Fontanarosa, D.: Siam-U-Net: encoder-decoder siamese network for knee cartilage tracking in ultrasound images. Medical Image Analysis (feb 2020)
16. Dunnhofer, M., Martinel, N., Foresti, G.L., Micheloni, C.: Visual Tracking by means of Deep Reinforcement Learning and an Expert Demonstrator. In: Proceedings of The IEEE/CVF International Conference on Computer Vision Workshops (2019)
17. Dunnhofer, M., Martinel, N., Micheloni, C.: A distilled model for tracking and tracker fusion (2020)

18. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (VOC) challenge. International Journal of Computer Vision (2010)
19. Fan, H., Lin, L., Yang, F., Chu, P., Deng, G., Yu, S., Bai, H., Xu, Y., Liao, C., Ling, H.: LaSOT: A High-quality Benchmark for Large-scale Single Object Tracking. In: IEEE Conference on Computer Vision and Pattern Recognition (sep 2019)
20. Gordon, D., Farhadi, A., Fox, D.: Re 3 : Real-time recurrent regression networks for visual tracking of generic objects. IEEE Robotics and Automation Letters (2018)
21. Hare, S., Golodetz, S., Saffari, A., Vineet, V., Cheng, M.M., Hicks, S.L., Torr, P.H.: Struck: Structured Output Tracking with Kernels. IEEE Transactions on Pattern Analysis and Machine Intelligence (2016)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (2016)
23. Held, D., Thrun, S., Savarese, S.: Learning to Track at 100 FPS with Deep Regression Networks. European Conference on Computer Vision (2016)
24. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. IEEE Transactions on Pattern Analysis and Machine Intelligence (2015)
25. Huang, L., Zhao, X., Huang, K.: GOT-10k: A Large High-Diversity Benchmark for Generic Object Tracking in the Wild. IEEE Transactions on Pattern Analysis and Machine Intelligence (oct 2019)
26. Jung, I., Son, J., Baek, M., Han, B.: Real-Time MDNet. In: European Conference on Computer Vision (2018)
27. Kim, C., Hwang, J.N.: Fast and automatic video object segmentation and tracking for content-based applications. IEEE Transactions on Circuits and Systems for Video Technology (feb 2002)
28. Kristan, M., Leonardis, A., Matas, J., Felsberg, M., Pflugfelder, R., Zajc, L.Č., Vojíř, T., Bhat, G., Lukežič, A., Eldesokey, A., Fernández, G., García-Martín, Á., Iglesias-Arias, Á., Alatan, A.A., González-García, A., Petrosino, A., Memarmoghadam, A., Vedaldi, A., Muhič, A., He, A., Smeulders, A., Perera, A.G., Li, B., Chen, B., Kim, C., Xu, C., Xiong, C., Tian, C., Luo, C., Sun, C., Hao, C., Kim, D., Mishra, D., Chen, D., Wang, D., Wee, D., Gavves, E., Gundogdu, E., Velasco-Salido, E., Khan, F.S., Yang, F., Zhao, F., Li, F., Battistone, F., De Ath, G., Subrahmanyam, G.R.K.S., Bastos, G., Ling, H., Galoogahi, H.K., Lee, H., Li, H., Zhao, H., Fan, H., Zhang, H., Possegger, H., Li, H., Lu, H., Zhi, H., Li, H., Lee, H., Chang, H.J., Drummond, I., Valmadre, J., Martin, J.S., Chahl, J., Choi, J.Y., Li, J., Wang, J., Qi, J., Sung, J., Johnander, J., Henriques, J., Choi, J., van de Weijer, J., Herranz, J.R., Martínez, J.M., Kittler, J., Zhuang, J., Gao, J., Grm, K., Zhang, L., Wang, L., Yang, L., Rout, L., Si, L., Bertinetto, L., Chu, L., Che, M., Maresca, M.E., Danelljan, M., Yang, M.H., Abdelpakey, M., Shehata, M., Kang, M., Lee, N., Wang, N., Miksik, O., Moallem, P., Vicente-Moñivar, P., Senna, P., Li, P., Torr, P., Raju, P.M., Ruihe, Q., Wang, Q., Zhou, Q., Guo, Q., Martín-Nieto, R., Gorthi, R.K., Tao, R., Bowden, R., Everson, R., Wang, R., Yun, S., Choi, S., Vivas, S., Bai, S., Huang, S., Wu, S., Hadfield, S., Wang, S., Golodetz, S., Ming, T., Xu, T., Zhang, T., Fischer, T., Santopietro, V., Štruc, V., Wei, W., Zuo, W., Feng, W., Wu, W., Zou, W., Hu, W., Zhou, W., Zeng, W., Zhang, X., Wu, X., Wu, X.J., Tian, X., Li, Y., Lu, Y., Law, Y.W., Wu, Y., Demiris, Y., Yang, Y., Jiao, Y., Li, Y., Zhang, Y., Sun, Y., Zhang, Z., Zhu, Z., Feng, Z.H., Wang, Z., He, Z.: The Sixth Visual Object Tracking VOT2018 Challenge Results. In: European Conference on Computer Vision. Springer, Cham (sep 2018)

29. Kristan, M., Matas, J., Leonardis, A., Felsberg, M., Pflugfelder, R., Kämäräinen, J.K., Zajc, L., Drbohlav, O., Lukežič, A., Berg, A., Eldesokey, A., Käpylä, J., Fernández, G., Gonzalez-Garcia, A., Memarmoghadam, A., Lu, A., He, A., Varfolomieiev, A., Chan, A., Shekhar Tripathi, A., Smeulders, A., Suraj Pedasingu, B., Xin Chen, B., Zhang, B., Wu, B., Li, B., He, B., Yan, B., Bai, B., Li, B., Li, B., Hak Kim, B., Ma, C., Fang, C., Qian, C., Chen, C., Li, C., Zhang, C., Tsai, C.Y., Luo, C., Micheloni, C., Zhang, C., Tao, D., Gupta, D., Song, D., Wang, D., Gavves, E., Yi, E., Shahbaz Khan, F., Zhang, F., Wang, F., Zhao, F., De Ath, G., Bhat, G., Chen, G., Wang, G., Li, G., Cevikalp, H., Du, H., Zhao, H., Saribas, H., Min Jung, H., Bai, H., Yu, H., Peng, H., Lu, H., Li, H., Li, J., Li, J., Fu, J., Chen, J., Gao, J., Zhao, J., Tang, J., Li, J., Wu, J., Liu, J., Wang, J., Qi, J., Zhang, J., Tsotsos, J.K., Hyuk Lee, J., van de Weijer, J., Kittler, J., Ha Lee, J., Zhuang, J., Zhang, K., Wang, K., Dai, K., Chen, L., Liu, L., Guo, L., Zhang, L., Wang, L., Wang, L., Zhang, L., Wang, L., Zhou, L., Zheng, L., Rout, L., Van Gool, L., Bertinetto, L., Danelljan, M., Dunnhofer, M., Ni, M., Young Kim, M., Tang, M., Yang, M.H., Paluru, N., Martinel, N., Xu, P., Zhang, P., Zheng, P., Zhang, P., Torr, P.H., Zhang Qiang Wang, Q., Guo, Q., Timofte, R., Krishna Gorthi, R., Everson, R., Han, R., Zhang, R., You, S., Zhao, S.C., Zhao, S., Li, S., Li, S., Ge, S., Bai, S., Guan, S., Xing, T., Xu, T., Yang, T., Zhang, T., Ojí, T., Feng, W., Hu, W., Wang, W., Tang, W., Zeng, W., Liu, W., Chen, X., Qiu, X., Bai, X., Wu, X.J., Yang, X., Chen, X., Li, X., Sun, X., Chen, X., Tian, X., Tang, X., Zhu, X.F., Huang, Y., Chen, Y., Lian, Y., Gu, Y., Liu, Y., Chen, Y., Zhang, Y., Xu, Y., Wang, Y., Li, Y., Zhou, Y., Dong, Y., Xu, Y., Zhang, Y., Li, Y., Wang Zhao Luo, Z., Zhang, Z., Feng, Z.H., He, Z., Song, Z., Chen, Z., Zhang, Z., Wu, Z., Xiong, Z., Huang, Z., Teng, Z., Ni, Z.: The Seventh Visual Object Tracking VOT2019 Challenge Results. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (2019)
30. Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., Yan, J.: SIAMRPN++: Evolution of siamese visual tracking with very deep networks. IEEE Conference on Computer Vision and Pattern Recognition (2019)
31. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High Performance Visual Tracking with Siamese Region Proposal Network. In: IEEE Conference on Computer Vision and Pattern Recognition. IEEE (jun 2018)
32. Li, M., Kambhamettu, C., Stone, M.: Automatic contour tracking in ultrasound images. Clinical Linguistics and Phonetics (sep 2005)
33. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: European Conference on Computer Vision (2014)
34. Luiten, J., Voigtlaender, P., Leibe, B.: PReMVOS: Proposal-Generation, Refinement and Merging for Video Object Segmentation. In: Asian Conference on Computer Vision. Springer Verlag (jul 2018)
35. Lukežič, A., Matas, J., Kristan, M.: D3S – A Discriminative Single Shot Segmentation Tracker. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (nov 2020)
36. Lukežič, A., Vojíř, T., Čehovin Zajc, L., Matas, J., Kristan, M.: Discriminative Correlation Filter Tracker with Channel and Spatial Reliability. International Journal of Computer Vision (2018)
37. McFarlane, N.J., Schofield, C.P.: Segmentation and tracking of piglets in images. Machine Vision and Applications (may 1995)
38. Mueller, M., Smith, N., Ghanem, B.: A Benchmark and Simulator for UAV Tracking. In: European Conference on Computer Vision. Springer, Cham (2016)

39. Müller, M., Bibi, A., Giancola, S., Alsubaihi, S., Ghanem, B.: TrackingNet: A Large-Scale Dataset and Benchmark for Object Tracking in the Wild. In: European Conference on Computer Vision. Springer Verlag (mar 2018)
40. Nam, H., Han, B.: Learning Multi-domain Convolutional Neural Networks for Visual Tracking. IEEE Conference on Computer Vision and Pattern Recognition (2016)
41. Nam, H., Hong, S., Han, B.: Online graph-based tracking. In: European Conference on Computer Vision. Springer Verlag (2014)
42. Oh, S.W., Lee, J.Y., Sunkavalli, K., Kim, S.J.: Fast Video Object Segmentation by Reference-Guided Mask Propagation. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. IEEE (jun 2018)
43. Park, E., Berg, A.C.: Meta-tracker: Fast and Robust Online Adaptation for Visual Object Trackers. In: European Conference on Computer Vision. Springer Verlag (jan 2018)
44. Perazzi, F., Pont-Tuset, J., McWilliams, B., Gool, L.V., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2016)
45. Pinheiro, P.O., Lin, T.Y., Collobert, R., Dollár, P.: Learning to refine object segments. In: European Conference on Computer Vision. Springer Verlag (mar 2016)
46. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 DAVIS Challenge on Video Object Segmentation (apr 2017)
47. Qi, H., Brown, M., Lowe, D.G.: Low-Shot Learning with Imprinted Weights. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2018)
48. Ren, L., Yuan, X., Lu, J., Yang, M., Zhou, J.: Deep Reinforcement Learning with Iterative Shift for Visual Tracking. In: European Conference on Computer Vision (2018)
49. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (2017)
50. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. Springer (2015)
51. Russakovsky, O., Li, L.J., Fei-Fei, L.: Best of both worlds: Human-machine collaboration for object annotation. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2015)
52. Shelhamer, E., Long, J., Darrell, T.: Fully Convolutional Networks for Semantic Segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence (apr 2017)
53. Siam, M., Oreshkin, B., Jagersand, M.: AMP: Adaptive masked proxies for few-shot segmentation. In: Proceedings of the IEEE International Conference on Computer Vision (feb 2019)
54. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: 3rd International Conference on Learning Representations (2015)
55. Voigtlaender, P., Leibe, B.: Online Adaptation of Convolutional Neural Networks for Video Object Segmentation (2017)

56. Voigtlaender, P., Leibe, B.: Online adaptation of convolutional neural networks for video object segmentation. In: British Machine Vision Conference 2017. BMVA Press (jun 2017)
57. Voigtlaender, P., Luiten, J., Leibe, B.: BoLTVOS: Box-Level Tracking for Video Object Segmentation (apr 2019)
58. Wang, Q., Gao, J., Xing, J., Zhang, M., Hu, W.: DCFNet: Discriminant Correlation Filters Network for Visual Tracking (apr 2017)
59. Wang, Q., Zhang, L., Bertinetto, L., Hu, W., Torr, P.H.S.: Fast Online Object Tracking and Segmentation: A Unifying Approach. In: IEEE Conference on Computer Vision and Pattern Recognition (2019)
60. Wu, Y., Lim, J., Yang, M.H.: Online object tracking: A benchmark. In: IEEE Conference on Computer Vision and Pattern Recognition. IEEE Computer Society (2013)
61. Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., Huang, T.: YouTube-VOS: A Large-Scale Video Object Segmentation Benchmark (sep 2018)
62. Yang, F., Mackey, M.A., Ianzini, F., Gallardo, G., Sonka, M.: Cell segmentation, tracking, and mitosis detection using temporal context. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). Springer Verlag (oct 2005)
63. Yang, L., Wang, Y., Xiong, X., Yang, J., Katsaggelos, A.K.: Efficient Video Object Segmentation via Network Modulation. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE Computer Society (feb 2018)
64. Yun, S., Choi, J., Yoo, Y., Yun, K., Choi, J.Y.: Action-decision networks for visual tracking with deep reinforcement learning. In: IEEE Conference on Computer Vision and Pattern Recognition. IEEE (jul 2017)
65. Zhang, Z., Peng, H.: Deeper and Wider Siamese Networks for Real-Time Visual Tracking. IEEE Conference on Computer Vision and Pattern Recognition (jan 2019)
66. Zhang, Z., Hua, Y., Song, T., Xue, Z., Ma, R., Robertson, N., Guan, H.: Tracking-assisted weakly supervised online visual object segmentation in unconstrained videos. In: MM 2018 - Proceedings of the 2018 ACM Multimedia Conference. Association for Computing Machinery, Inc, New York, New York, USA (oct 2018)
67. Zhu, Z., Wang, Q., Li, B., Wu, W., Yan, J., Hu, W.: Distractor-aware siamese networks for visual object tracking. In: European Conference on Computer Vision (aug 2018)