

Robust Visual Tracking by Segmentation

Matthieu Paul, Martin Danelljan, Christoph Mayer, and Luc Van Gool

Computer Vision Lab, ETH Zürich, Switzerland
{paulma, damartin, chmayer, vangool}@vision.ee.ethz.ch

Abstract. Estimating the target extent poses a fundamental challenge in visual object tracking. Typically, trackers are *box-centric* and fully rely on a bounding box to define the target in the scene. In practice, objects often have complex shapes and are not aligned with the image axis. In these cases, bounding boxes do not provide an accurate description of the target and often contain a majority of background pixels.

We propose a *segmentation-centric* tracking pipeline that not only produces a highly accurate segmentation mask, but also internally works with segmentation masks instead of bounding boxes. Thus, our tracker is able to better learn a target representation that clearly differentiates the target in the scene from background content. In order to achieve the necessary robustness for the challenging tracking scenario, we propose a separate instance localization component that is used to condition the segmentation decoder when producing the output mask. We infer a bounding box from the segmentation mask, validate our tracker on challenging tracking datasets and achieve the new state of the art on LaSOT with a success AUC score of 69.7%. Since most tracking datasets do not contain mask annotations, we cannot use them to evaluate predicted segmentation masks. Instead, we validate our segmentation quality on two popular video object segmentation datasets. The code and trained models are available at <https://github.com/visionml/pytracking>.

1 Introduction

Visual object tracking is the task of estimating the state of a target object for each frame in a video sequence. The target is solely characterized by its initial state in the video. Current approaches predominately characterize the state itself with a bounding box. However, this only gives a very coarse representation of the target in the image. In practice, objects often have complex shapes, undergo substantial deformations. Often, targets do not align well with the image axes, while most benchmarks use axis-aligned bounding boxes. In such cases, the majority of the image content inside the target’s bounding box often consists of background regions which provide limited information about the object itself. In contrast, a segmentation mask precisely indicates the object’s extent in the image (see Fig. 1 frames #1600 and #3200). Such information is vital in a variety of applications, including video analysis, video editing, and robotics. In this work, we therefore develop an approach for accurate and robust target object segmentation, even in the highly challenging tracking datasets [16,36].

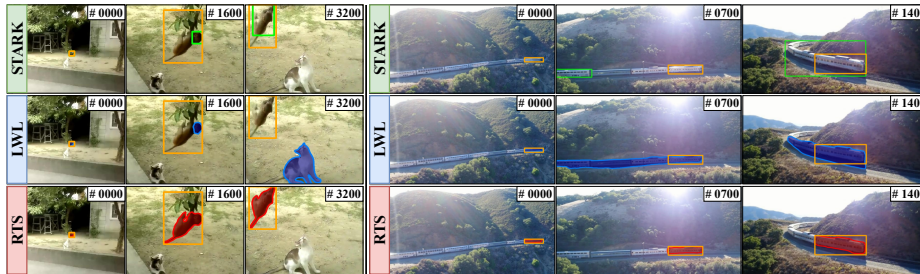


Fig. 1. Comparison between the VOT method Stark [55], the VOS method LWL [5] and our proposed method on two tracking sequences from the LaSOT [16] dataset. The ground-truth annotation (\square) is shown in each frame for reference. Our approach is more robust and predicts a more accurate target representation.

While severely limiting the information about the target’s state in the video, the aforementioned issues with the bounding box representation can itself lead to inaccurate bounding box predictions, or even tracking failure. To illustrate this, Fig. 1 shows two typical tracking sequences. The tracking method STARK [55] (first row) fails to regress bounding boxes that contain the entire object (#1600, #1400) or even starts tracking the wrong object (#0700). Conversely, segmentation masks are a better fit to differentiate pixels in the scene that belong to the background and the target. Therefore, a *segmentation-centric* tracking architecture designed to work internally with a segmentation mask of the target instead of a bounding box has the potential to learn better target representations, because it can clearly differentiate background from foreground regions in the scene.

A few recent tracking methods [47,54] have recognized the advantage of producing segmentation masks instead of bounding boxes as final output. However, these trackers are typically *bounding-box-centric* and the final segmentation mask is obtained by a separate *box-to-mask* post-processing network. These methods do not leverage the accurate target definition of segmentation masks to learn a more accurate and robust internal representation of the target.

In contrast, most Video Object Segmentation (VOS) methods [38,5] follow a *segmentation-centric* paradigm. However, these methods are not designed for the challenging tracking scenarios. Typical VOS sequences consist only of a few hundred frames [41] whereas multiple sequences of more than ten thousand frames exist in tracking datasets [16]. Due to this setup, VOS methods focus on producing highly accurate segmentation masks but are sensitive to distractors, substantial deformations and occlusions of the target object. Fig. 1 shows two typical tracking sequences where the VOS method LWL [5] (second row) produces a fine-grained segmentation mask of the wrong object (#3200) or is unable to detect only the target within a crowd (#0700, #1400).

We propose *Robust Visual Tracking by Segmentation* (RTS), a unified tracking architecture capable of predicting accurate segmentation masks. To design

a *segmentation-centric* approach, we take inspiration from the aforementioned LWL [5] method. However, to achieve robust and accurate segmentation on Visual Object Tracking (VOT) datasets, we introduce several new components. In particular, we propose an instance localization branch trained to predict a target appearance model, which allows occlusion detection and target identification even in cluttered scenes. The output of the instance localization branch is further used to condition the high-dimensional mask encoding. This allows the segmentation decoder to focus on the localized target, leading to a more robust mask prediction. Since our proposed method contains a segmentation and instance memory that need to be updated with previous tracking results, we design a memory management module. This module first assesses the prediction quality, decides whether the sample should enter the memory and, when necessary, triggers the model update.

Contributions Our contributions are the following: **(i)** We propose a unified tracking architecture capable of predicting robust classification scores and accurate segmentation masks. We design separate feature spaces and memories to ensure optimal receptive fields and update rates for segmentation and instance localization. **(ii)** To produce a segmentation mask which agrees with the instance prediction, we design a fusion mechanism that conditions the segmentation decoder on the instance localization output and leads to more robust tracking performance. **(iii)** We introduce an effective inference procedure capable of fusing the instance localization output and mask encoding to ensure both robust and accurate tracking. **(iv)** We perform comprehensive evaluation and ablation studies of the proposed tracking pipeline on multiple popular tracking benchmarks. Our approach achieves the new state of the art on LaSOT with an *area-under-the-curve* (AUC) score of 69.7%.

2 Related Work

Visual Object Tracking

Over the years, research in the field of visual tracking has been accelerated by the introduction of new and challenging benchmarks, such as LaSOT [16], GOT-10k [24], and TrackingNet [37]. This led to the introduction of new paradigms in visual object tracking, based on Discriminative Correlation Filters (DCF), Siamese networks and Transformers.

One of the most popular type of approaches, DCF-based visual trackers [6,22,15,32,12,48,61,3,14] essentially solve an optimization problem to estimate the weights of the DCF that allow to distinguish foreground from background regions. The DCF is often referred to as the target appearance model and allows to localize the target in the video frame. More recent DCF approaches [3,14] enable end-to-end training by unrolling a fixed number of the optimization iterations during *offline* training.

Siamese tracking methods have gained in popularity due to their simplicity, speed and end-to-end trainability [44,2,43,62,20,50,21,29,28]. These trackers learn a similarity metric using only the initial video frame and its annotation

that allows to clearly identify the target *offline*. Since no *online* learning component is involved, these trackers achieve high frame rates at the cost of limited *online* adaptability to changes of the target’s appearance. Nonetheless, several methods have been proposed to overcome these issues [44,2,29,28].

Very recently, Transformer-based trackers have achieved state-of-the-art performance on many datasets, often outperforming their rivals. This group of trackers typically uses a Transformer component in order to fuse information extracted from training and test frames. This produces discriminative features that allow to accurately localize and estimate the target in the scene [8,56,55,49,34].

Video Object Segmentation Semi-supervised VOS is the task of classifying all pixels belonging to the target in each video frame, given only the segmentation mask of the target in the initial frame. The cost of annotating accurate segmentation masks is limiting the sequence length and number of videos contained in available VOS datasets. Despite the relatively small size of VOS datasets compared to other computer vision problems, new benchmarks such as Youtube-VOS [53] and DAVIS [41] accelerated the research progress in the last years.

Some methods rely on a learnt target detector [7,46,33], others learn how to propagate the segmentation mask across frames [52,40,30,25]. Another group of methods uses feature matching techniques across one or multiple frames with or without using an explicit spatio-temporal memory [9,23,45,38]. Recently, Bhat *et al.* [5] employed meta-learning approach, introducing an end-to-end trainable VOS architecture. In this approach, a few-shot learner predicts a learnable labels encoding. It generates and updates *online* the parameters of a segmentation target model that produces the mask encoding used to generate the final segmentation mask.

Joint Visual Tracking and Segmentation A group of tracking methods have already identified the advantages of predicting a segmentation mask instead of a bounding box [54,60,47,31,51,42]. Siam-RCNN is a box-centric tracker that uses a pretrained *box2seg* network to predict the segmentation mask given a bounding box prediction. In contrast, AlphaRefine represents a novel *box2seg* method that has been evaluated with many recent trackers such as SuperDiMP [14] and SiamRPN++ [28]. Further, Zhao *et al.* [60] focus on generating segmentation masks from bounding box annotations in videos using a spatio-temporal aggregation module to mine consistencies of the scene across multiple frames. Conversely, SiamMask [51] and D3S [31] are segmentation-centric trackers that produce a segmentation mask directly, without employing a *box2seg* module. In particular, SiamMask [51] is a fully-convolutional Siamese network with a separate branch which predicts binary segmentation masks supervised by a segmentation loss.

From a high-level view, the single-shot segmentation tracker D3S [31] is most related to our proposed method. Both methods employ two dedicated modules or branches; one for localization and one for segmentation. D3S adopts the target classification component of ATOM [12], requiring online optimization of weights in a two-layer CNN. In contrast, we learn online the weights of a DCF similar

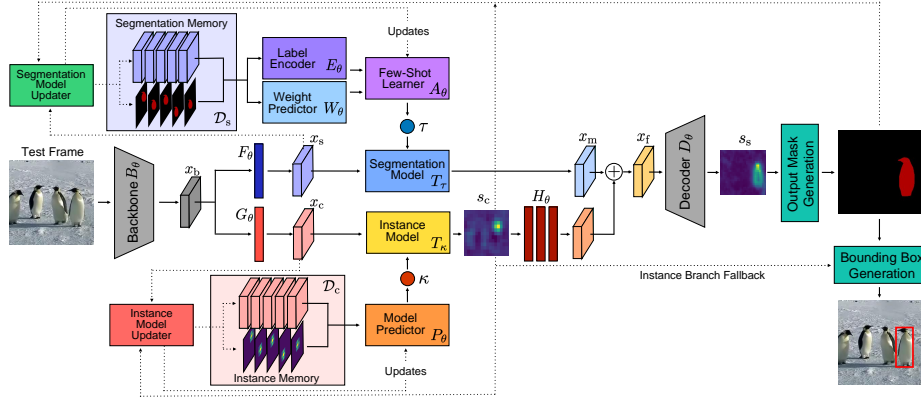


Fig. 2. Overview of our entire online tracking pipeline used for inference, see Sec 3.1.

to DiMP [3]. For segmentation, D3S [31] propose a feature matching technique that matches test frame features with background and foreground features corresponding to the initial frame. In contrast, we adopt the few-shot learning based model prediction proposed in LWL [5] to produce accurate segmentation masks. Furthermore, D3S proposes to simply concatenate the outputs of both modules whereas we learn a localization encoding to condition the segmentation mask decoding based on the localization information. Compared to D3S, we update not only the instance localization but also the segmentation models and memories. Hence, our method integrates specific memory management components.

3 Method

3.1 Overview

Video object segmentation methods can produce high quality segmentation masks but are typically not robust enough for video object tracking. Robustness becomes vital for medium and long sequences, which are most prevalent in tracking datasets [16,36]. In such scenarios, the target object frequently undergoes substantial appearance changes. Occlusions and similarly looking objects are common. Hence, we propose to adapt a typical VOS approach with tracking components to increase its robustness. In particular, we base our approach on the Learning What to Learn (LWL) [5] method and design a novel and segmentation-centric tracking pipeline that estimates accurate object masks instead of bounding boxes. During inference, a segmentation mask is typically not provided in visual object tracking. Hence, we use STA [60] to generate a segmentation mask from the provided initial bounding box. An overview of our RTS method is shown in Fig. 2. Our pipeline consists of a backbone network, a segmentation branch, an instance localization branch and a segmentation decoder. For each video frame, the backbone first extracts a feature map x_b . These features are

further processed into segmentation features x_s and classification features x_c to serve as input for their respective branch. The segmentation branch is designed to capture the details of the object with a high dimensional mask encoding, whereas the instance localization branch aims at providing a coarser but robust score map representing the target location. Both branches contain components learned online, trained on memories (\mathcal{D}_s and \mathcal{D}_c) that store features and predictions of past frames. The instance localization branch has two purposes. The first is to control models and memories updating. The second is used to condition the segmentation mask decoder. To do so, we add instance localization information with a learnt score encoding produced by H_θ . The obtained segmentation scores and the raw instance model score map are then used to generate the final segmentation mask output.

3.2 Segmentation Branch

The architecture of the segmentation branch is adopted from LWL [5], and we briefly review it here. It consists of a segmentation sample memory \mathcal{D}_s , a label generator E_θ , a weight predictor W_θ , a few-shot learner A_θ and a segmentation model T_τ . The goal of the few-shot learner A_θ is producing the parameters τ of the segmentation model T_τ such that the obtained mask encoding x_m contains the information needed to compute the final segmentation mask of the target object. The label mask encodings used by the few-shot learner are predicted by the label generator E_θ .

The few-shot learner is formulated through the following optimization problem, which is unrolled through steepest descent iterations in the network

$$L_s(\tau) = \frac{1}{2} \sum_{(x_s, y_s) \in \mathcal{D}_s} \|W_\theta(y_s) \cdot (T_\tau(x_s) - E_\theta(y_s))\|^2 + \frac{\lambda_s}{2} \|\tau\|^2, \quad (1)$$

where \mathcal{D}_s corresponds to the segmentation memory, x_s denotes the segmentation features, y_s the segmentation masks and λ_s is a learnable scalar regularization parameter. The weight predictor W_θ produces sample confidence weights for each spatial location in each memory sample. Applying the optimized model parameters τ^* within the segmentation model produces the mask encoding $x_m = T_{\tau^*}(x_s)$ for the segmentation features x_s .

LWL [5] feeds the mask encoding directly into the segmentation decoder to produce the segmentation mask. For long and challenging tracking sequences, only relying on the mask encoding may lead to an accurate segmentation mask, but often for the wrong object in the scene (see Fig 1). Since LWL [5] is only able to identify the target to a certain degree in challenging tracking sequences, we propose to condition the mask encoding based on an instance localization representation, described next.

3.3 Instance Localization Branch

The segmentation branch can produce accurate masks but typically lacks the necessary robustness for tracking in medium or long-term sequences. Especially

challenging are sequences where objects similar to the target appear, where the target object is occluded or vanishes from the scene for a short time. Therefore, we propose a dedicated branch for target instance localization, in order to robustly identify the target among distractors or to detect occlusions. A powerful tracking paradigm that learns a target-specific appearance model on both foreground and background information are discriminative correlation filters (DCF) [6,22,13,3]. These methods learn the weights of a filter that differentiates foreground from background pixels represented by a score map, where the maximal value corresponds to the target’s center.

Similar to the segmentation branch, we propose an instance localization branch that consists of a sample memory \mathcal{D}_c and a model predictor P_θ . The latter predicts the parameters κ of the instance model T_κ . The instance model is trained online to produce the target score map used to localize the target object. To obtain the instance model parameters κ we minimize the following loss function

$$L_c(\kappa) = \sum_{(x_c, y_c) \in \mathcal{D}_c} \|R(T_\kappa(x_c), y_c)\|^2 + \frac{\lambda_c}{2} \|\kappa\|^2, \quad (2)$$

where \mathcal{D}_c corresponds to the instance memory containing the classification features x_c and the Gaussian labels y_c . R denotes the robust hinge-like loss [3] and λ_c is a fixed regularization parameter. To solve the optimization problem we apply the method from [3], which unrolls steepest descent iterations of the Gauss-Newton approximation of (2) to obtain the final model parameters κ^* . The score map can then be obtained with $s_c = T_{\kappa^*}(x_c)$ by evaluating the target model on the classification features x_c .

3.4 Instance-Conditional Segmentation Decoder

In video object segmentation the produced mask encoding is directly fed into the segmentation decoder to generate the segmentation mask. However, solely relying on the mask encoding is not robust enough for the challenging tracking scenario, see Fig 1. Thus, we propose to integrate the instance localization information into the segmentation decoding procedure. In particular, we condition the mask encoding on a learned encoding of the instance localization score map.

First, we encode the raw score maps using a multi-layer Convolutional Neural Network (CNN) to learn a suitable representation. Secondly, we condition the mask encoding with the learned representation using element-wise addition. The entire conditioning procedure can be defined as $x_f = x_m + H_\theta(s_c)$, where H_θ denotes the CNN encoding the scores s_c , and x_m the mask encoding. The resulting features are then fed into the segmentation decoder that produces the segmentation scores of the target object.

3.5 Jointly Learning Instance Localization and Segmentation

In this section, we describe our general training strategy and parameters. In particular, we further detail the segmentation and classification losses that we use for offline training.

Segmentation Loss First, we randomly sample J frames from an annotated video sequence and sort them according to their frame IDs in increasing order to construct the training sequence $\mathcal{V} = \{(x_b^j, y_s^j, y_c^j)\}_{j=0}^{J-1}$, where $x_b^j = B_\theta(I^j)$ are the extracted features of the video frame I^j using the backbone B_θ , y_s^j is the corresponding segmentation mask and y_c^j denotes the Gaussian label at the target’s center location. We start with entry $v_0 \in \mathcal{V}$ and store it in the segmentation \mathcal{D}_s and instance memory \mathcal{D}_c and obtain parameters τ^0 and κ^0 of the segmentation and instance model. We use these parameters to compute the segmentation loss for $v_1 \in \mathcal{V}$. Using the predicted segmentation mask, we update the segmentation model parameters to τ^1 but keep the instance model parameters fixed. Segmentation parameters typically need to be updated frequently to enable accurate segmentation. Conversely, we train the model predictor on a single frame only. The resulting instance model generalizes to multiple unseen future frames, ensuring robust target localization. The resulting segmentation loss for the entire sequence \mathcal{V} can thus be described as follows

$$\mathcal{L}_s^{\text{seq}}(\theta; \mathcal{V}) = \sum_{j=1}^{J-1} \mathcal{L}_s \left(D_\theta \left(T_{\tau^{j-1}}(x_s^j) + H_\theta(T_{\kappa^0}(x_c^j)) \right), y_s^j \right), \quad (3)$$

where $x_s = F_\theta(x_b)$ and $x_c = G_\theta(x_b)$ and \mathcal{L}_s is the Lovasz segmentation loss [1].

Classification Loss Instead of training our tracker only with the segmentation loss, we add an auxiliary loss to ensure that the instance module produces score maps localizing the target via a Gaussian distribution. These score maps are essential to update the segmentation and instance memories and to generate the final output. As explained before, we use only the first training $v_0 \in \mathcal{V}$ to optimize the instance model parameters. To encourage fast convergence, we use not only the parameters corresponding to the final iteration N_{iter} of the optimization method $\kappa_{(N_{\text{iter}})}^0$ explained in Sec. 3.3, but also all the intermediate parameters $\kappa_{(i)}^0$ of loss computation. The final target classification loss for the whole sequence \mathcal{V} is defined as follows

$$\mathcal{L}_c^{\text{seq}}(\theta; \mathcal{V}) = \sum_{j=1}^{J-1} \left(\frac{1}{N_{\text{iter}}} \sum_{i=0}^{N_{\text{iter}}} \mathcal{L}_c \left(T_{\kappa_{(i)}^0}(x_c^j), y_c^j \right) \right), \quad (4)$$

where \mathcal{L}_c is the hinge loss defined in [3]. To train our tracker we combine the segmentation and classification losses using the scalar weight η and minimize both losses jointly

$$\mathcal{L}_{\text{tot}}^{\text{seq}}(\theta; \mathcal{V}) = \mathcal{L}_s^{\text{seq}}(\theta; \mathcal{V}) + \eta \cdot \mathcal{L}_c^{\text{seq}}(\theta; \mathcal{V}). \quad (5)$$

Training Details We use the train sets of LaSOT [16], GOT-10k [24], Youtube-VOS [53] and DAVIS [41]. For VOT datasets that only provide annotated bounding boxes, we use these boxes and STA [60] to generate segmentation masks and

treat them as ground truth annotations during training. STA [60] is trained separately on YouTube-VOS 2019 [53] and DAVIS 2017 [39]. For our model, we use ResNet-50 with pre-trained MaskRCNN weights as our backbone and initialize the segmentation model and decoder weights with the ones available from LWL [5]. We train for 200 epochs and sample 15'000 videos per epoch, which takes 96 hours to train on a single Nvidia A100 GPU. We use the ADAM [26] optimizer with a learning rate decay of 0.2 at epochs 25, 115 and 160. We weigh the losses such that the segmentation loss is predominant but in the same range as the classification loss. We empirically choose $\eta = 10$. Further details about training and the network architecture are given in the appendix.

3.6 Inference

Memory Management and Model Updating Our tracker consists of two different memory modules. A segmentation memory that stores segmentation features and predicted segmentation masks of previous frames. In contrast, an instance memory contains classification features and Gaussian labels marking the center location of the target in the predicted segmentation mask of the previous video frame. The quality of the predicted labels directly influences the localization and segmentation quality in future video frames. Hence, it is crucial to avoid contaminating the memory modules with predictions that do not correspond to the actual target. We propose the following strategy to keep the memory as clean as possible. (a) If the instance model is able to clearly localize the target (maximum value in the score map larger than $t_{sc} = 0.3$) and the segmentation model constructs a valid segmentation mask (at least one pixel above $t_{ss} = 0.5$) we update both memories with the current predictions and features. (b) If either the instance localization or segmentation fail to identify the target we omit updating the segmentation memory. (c) If only the segmentation mask fails to represent the target but the instance model can localize it, we update the instance memory only. (d) If instance localization fails we do not update either memory. Further, we trigger the few-shot learner and model predictor after 20 frames have passed, but only if the corresponding memory has been updated.

Final Mask Output Generation We obtain the final segmentation mask by thresholding the segmentation decoder output. To obtain the bounding box required for standard tracking benchmarks, we report the smallest axis-aligned box that contains the entire estimated object mask.

Inference Details We set the input image resolution such that the segmentation learner features have a resolution of 52×30 (stride 16), while the instance learner operates on features of size 26×15 (stride 32). The learning rate is set to 0.1 and 0.01 for the segmentation and instance learner respectively. We use a maximum buffer of 32 frames for the segmentation memory and 50 frames for the instance memory. We keep the samples corresponding to the initial frame in both memories and replace the oldest entries if the memory is full. We update both memories for the first 100 video frames and afterwards only after every 20th frame. We randomly augment the sample corresponding to the initial frame with vertical flip, random translation and blurring.

4 Evaluation

Our approach is developed within the PyTracking [11] framework. The implementation is done with PyTorch 1.9 with CUDA 11.1. Our model is evaluated on a single Nvidia GTX 2080Ti GPU. Our method achieves an average speed of 30 FPS on LaSOT [16]. Each number corresponds to the average of five runs with different random seeds.

4.1 Branch Ablation Study

For the ablation study, we analyze the impact of the instance branch on three datasets and present the results in Tab. 1. First, we report the performance of LWL [5] since we build upon it to design our final tracking pipeline. We use the network weights provided by Bhat *et al.* [5] and the corresponding inference settings. We input the same segmentation masks obtained from the initial bounding box for LWL as used for our method. We observe that LWL is not robust enough for challenging tracking scenarios. The second row in Tab. 1 corresponds to our method but we omit the proposed instance branch. Hence, we use the proposed inference components and settings and train the tracker as explained in Sec. 3.5, but with conditioning removed. We observe that even without the instance localization branch our tracker can achieve competitive performance on all three datasets (*e.g.* +5.6% on LaSOT). Fully integrating the instance localization branch increases the performance even more (*e.g.* +4.4 on LaSOT). Thus, we conclude that adapting the baseline method to the tracking domain improves the tracking performance. To boost the performance and achieve state-of-the-art results, an additional component able to increase the tracking robustness is required.

4.2 Inference Parameters

In this part, we ablate two key aspects of our inference strategy. First, we study the effect of relying on the instance branch if the segmentation decoder is unable to localize the target ($\max(s_s) < t_{s_s}$). Second, we study different values for t_{s_c} that determines whether the target is detected by the instance model, see Tab. 2.

If the segmentation branch cannot identify the target, using the instance branch improves tracking performance on all datasets (*e.g.* +1.3% on UAV123).

Table 1. Comparison between our segmentation network baseline LWL and our pipeline, with and without Instance conditioning on different VOT datasets.

Method	Seg. Branch	Inst. Branch Conditioning	LaSOT [16]			NFS [19]			UAV123 [36]		
			AUC	P	NP	AUC	P	NP	AUC	P	NP
LWL [5]	✓	-	59.7	60.6	63.3	61.5	75.1	76.9	59.7	78.8	71.4
RTS	✓	✗	65.3	68.5	71.5	65.8	84.0	85.0	65.2	85.6	78.8
RTS	✓	✓	69.7	73.7	76.2	65.4	82.8	84.0	67.6	89.4	81.6

Table 2. Ablation on inference strategies. The first column analyzes the effect of using the instance branch as fallback for target localization if the segmentation branch is unable to detect the target ($\max(s_s) < t_{s_c}$). The second column shows the impact of different confidence thresholds t_{s_c} .

Inst. Branch Fallback	t_{s_c}	LaSOT [16]			NFS [19]			UAV123 [36]		
		AUC	P	NP	AUC	P	NP	AUC	P	NP
✗	0.30	69.3	73.1	75.9	65.3	82.7	84.0	66.3	87.2	80.4
✓	0.30	69.7	73.7	76.2	65.4	82.8	84.0	67.6	89.4	81.6
✓	0.20	68.6	72.3	75.0	65.3	82.7	83.9	67.0	88.7	80.7
✓	0.30	69.7	73.7	76.2	65.4	82.8	84.0	67.6	89.4	81.6
✓	0.40	69.1	72.7	75.6	63.3	79.7	81.7	67.1	89.1	80.7

Table 3. Comparison to the state of the art on the LaSOT [16] test set in terms of AUC score. The methods are ordered by AUC score.

	RTS	ToMP 101 [34]	ToMP 50 [34]	Keep Track [35]	STARK ST-101 [55]	Alpha Refine [54]	TransT [8]	Siam R-CNN [47]	Tr DiMP [49]	Super DiMP [11]	STM Track [18]	Pr DiMP [14]	DM LWL [5]	DM Track [59]	LTMU [10]	DiMP [3]	Ocean [58]	D3S [31]
Precision	73.7	73.5	72.2	70.2	72.2	68.8	69.0	68.4	66.3	65.3	63.3	60.8	60.6	59.7	57.2	56.7	56.6	49.4
Norm. Prec	76.2	79.2	78.0	77.2	76.9	73.8	73.8	72.2	73.0	72.2	69.3	68.8	63.3	66.9	66.2	65.0	65.1	53.9
Success (AUC)	69.7	68.5	67.6	67.1	67.1	65.9	64.9	64.8	63.9	63.1	60.6	59.8	59.7	58.4	57.2	56.9	56.0	49.2
Δ AUC to Ours	-	↑1.2	↑2.1	↑2.6	↑2.6	↑3.8	↑4.8	↑4.9	↑5.8	↑6.6	↑9.1	↑9.9	↑10.0	↑11.3	↑12.5	↑12.8	↑13.7	↑20.5

Furthermore, Tab. 2 shows that our tracking pipeline achieves the best performance when setting $t_{s_c} = 0.3$ whereas smaller or larger values for t_{s_c} decrease the tracking accuracy. Hence, it is important to find a suitable trade-off between frequently updating the model and memory to quickly adapt to appearance changes and updating only rarely to avoid contaminating the memory and model based on wrong predictions.

4.3 Comparison to the state of the art

Assessing segmentation accuracy on tracking datasets is not possible since only bounding box annotations are provided. Therefore, we compare our approach on six VOT benchmarks and validate the segmentation masks quality on two VOS datasets.

LaSOT [16] We evaluate our method on the test set of the LaSOT dataset, consisting of 280 sequences with 2500 frames on average. Thus, the benchmark challenges the long term adaptability and robustness of trackers. Fig. 3 shows the success plot reporting the overlap precision OP with respect to the overlap threshold T . Trackers are ranked by AUC score. In addition, Tab. 3 reports the precision and normalized precision for all compared methods. Our method outperforms the state-of-the-art ToMP-50 [34] and ToMP-101 [34] by large margins (+1.2% and +2.1% AUC respectively). Our method is not only as robust as KeepTrack (see the success plot for $T < 0.2$) but also estimates far more accurate bounding boxes than any tracker ($0.8 < T < 1.0$).




Fig. 3. Success (left) and Precision (right) plots on LaSOT [16] with other state-of-the-art methods. The AUCs for all methods are ordered and reported in the legend. Our method outperforms all existing approaches, both in Overlap Precision (left) and Distance Precision (right).

Table 4. Results on the GOT-10k validation set [24] in terms of Average Overlap (AO) and Success Rates (SR) for overlap thresholds of 0.5 and 0.75.

	RTS	STA	LWL	PrDiMP-50	DiMP-50	SiamRPN++
	[60]	[5]	[14]	[3]	[28]	
SR _{0.50} (%)	94.5	95.1	92.4	89.6	88.7	82.8
SR _{0.75} (%)	82.6	85.2	82.2	72.8	68.8	-
AO(%)	85.2	86.7	84.6	77.8	75.3	73.0

Table 5. Comparison to the state of the art on the TrackingNet [37] test set in terms of AUC scores, Precision and Normalized Precision.

	RTS	ToMP 101	ToMP 50	Keep Track	STARK ST101	STARK ST50	STA	LWL	TransT	Siam R-CNN	Alpha Refine	STM Track	DTT	DiMP	Super DiMP	Pr DiMP	D3S
	[34]	[34]	[35]	[55]	[55]	[60]	[5]	[8]	[47]	[54]	[18]	[56]	[49]	[11]	[14]	[31]	
Precision	79.4	78.9	78.6	73.8	-	-	79.1	78.4	80.3	80.0	78.3	76.7	78.9	73.1	73.3	70.4	66.4
Norm. Prec	86.0	86.4	86.2	83.5	86.9	86.1	84.7	84.4	86.7	85.4	85.6	85.1	85.0	83.3	83.5	81.6	76.8
Success (AUC)	81.6	81.5	81.2	78.1	82.0	81.3	81.2	80.7	81.4	81.2	80.5	80.3	79.6	78.4	78.1	75.8	72.8
Δ AUC to Ours	-	$\uparrow 0.1$	$\uparrow 0.4$	$\uparrow 3.5$	$\downarrow 0.4$	$\uparrow 0.3$	$\uparrow 0.4$	$\uparrow 0.9$	$\uparrow 0.2$	$\uparrow 0.4$	$\uparrow 1.1$	$\uparrow 1.3$	$\uparrow 2.0$	$\uparrow 3.2$	$\uparrow 3.5$	$\uparrow 5.8$	$\uparrow 8.8$

GOT-10k [24] The large-scale GOT-10k dataset contains over 10,000 shorter sequences. Since we train our method on several datasets instead of only GOT-10k *train*, we evaluate it on the *val* set only, which consists of 180 short videos. We compile the results in Tab. 4. Our method ranks second for all metrics, falling between two VOS-oriented methods, +0.6% over LWL [5] and -1.5% behind STA [60]. Our tracker outperforms other trackers by a large margin.

TrackingNet [37] We compare our approach on the test set of the TrackingNet dataset, consisting of 511 sequences. Tab. 5 shows the results obtained from the online evaluation server. Our method outperforms most of the existing approaches and ranks second in terms of AUC, close behind STARK-ST101 [55] which is based on a ResNet-101 backbone. Note that we outperform STARK-ST50 [55] that uses a ResNet-50 as backbone. Also, we achieve a higher precision score than other methods that produce a segmentation mask output such as LWL [5], STA [60], Alpha-Refine [54] and D3S [31].

Table 6. Comparison with state-of-the-art on the UAV123 [36] and NFS [19] datasets in terms of AUC score.

		ToMP	ToMP	Keep				STARK			STARK	Super	Pr	STM	Siam	Siam			
	RTS	101	50	Track	CRACK	ST101	TrDiMP	TransT		ST50	DiMP	DiMP	Track	AttN	R-CNN	KYS	DiMP	LWL	
		[34]	[34]	[35]	[17]	[55]	[49]	[8]		[55]	[11]	[14]	[18]	[57]	[47]	[4]	[3]	[5]	
UAV123	67.6	66.9	69.0	69.7	66.4	68.2	67.5	69.1	69.1	67.7	68.0	64.7	65.0	64.9	–	65.3	59.7		
NFS	65.4	66.7	66.9	66.4	62.5	66.2	66.2	65.7	65.2	64.8	63.5	–	–	63.9	63.5	62.0	61.5		

Table 7. Results on the VOT2020-ST [27] challenge in terms of Expected Average Overlap (EAO), Accuracy and Robustness.

		STARK	STARK-														
	RTS	ST-50	ST-101-	LWL	STA	Ocean	Fast	Alpha									
		+AR	+AR	[27]	[60]	Plus	Ocean	Refine	RPT	AFOD	D3S	STM					
		[55]	[55]	[27]	[60]	[27]	[27]	[27]	[27]	[27]	[27]	[27]					
Robustness	0.845	0.817	0.789	0.798	0.824	0.842	0.803	0.777	0.869	0.795	0.769	0.574					
Accuracy	0.710	0.759	0.763	0.719	0.732	0.685	0.693	0.754	0.700	0.713	0.699	0.751					
EAO	0.506	0.505	0.497	0.463	0.510	0.491	0.461	0.482	0.530	0.472	0.439	0.308					
Δ EAO to Ours	-	\uparrow 0.001	\uparrow 0.009	\uparrow 0.043	\downarrow 0.004	\uparrow 0.015	\uparrow 0.045	\uparrow 0.024	\downarrow 0.024	\uparrow 0.034	\uparrow 0.067	\uparrow 0.198					

UAV123 [36] The UAV dataset consists of 123 test videos that contain small objects, target occlusion, and distractors. Small objects are particularly challenging in a segmentation setup. Tab. 6 shows the achieved results in terms of success AUC. Our method achieves competitive results on UAV123, close to TrDiMP [49] or SuperDiMP [11]. It outperforms LWL [5] by a large margin.

NFS [19] The NFS dataset (30FPS version) contains 100 test videos with fast motions and challenging sequences with distractors. Our method achieves an AUC score that is only 1% below the current best method KeepTrack [35] while outperforming numerous other trackers, including STARK-ST50 [55] (+0.2) SuperDiMP [3] (+0.6) and PrDiMP [14] (+1.9).

VOT 2020 [27] Finally, we evaluate our method on the VOT2020 short-term challenge. It consists of 60 videos and provides segmentation mask annotations. For the challenge, the multi-start protocol is used and the tracking performance is assessed based on accuracy and robustness. We compare with the top methods on the leader board and include more recent methods in Tab. 7. In this setup, our method ranks 2nd in Robustness, thus outperforming most of the other methods. In particular, we achieve a higher EAO score than STARK [55], LWL [5], AlphaRefine [54] and D3S [31].

YouTube-VOS 2019 [53] We use the validation set which consist of 507 sequences. They contain 91 object categories out of which 26 are *unseen* in the training set. The results presented in Tab. 8 were generated by an online server after uploading the raw results. On this benchmark, we want to validate the quality of the produced segmentation masks rather than to achieve the best accuracy possible. Hence, we use the same model weight as for VOT without further fine tuning.

When using the provided segmentation masks for initialization, we observe that our method performs slightly worse than LWL [5] and STA [60] (-1.3 \mathcal{G} , -0.9 \mathcal{G}) but still outperforms the VOS method STM [38] (+0.5 \mathcal{G}). We conclude that

Table 8. Results on the Youtube-VOS 2019 [53] and DAVIS 2017 [41] datasets. The table is split in two parts to separate methods using bounding box initialization or segmentation masks initialization, in order to enable a fair comparison.

Method	YouTube-VOS 2019 [53]					DAVIS 2017 [41]		
	\mathcal{G}	$\mathcal{J}_{\text{seen}}$	$\mathcal{J}_{\text{unseen}}$	$\mathcal{F}_{\text{seen}}$	$\mathcal{F}_{\text{unseen}}$	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
RTS	79.7	77.9	75.4	82.0	83.3	80.2	77.9	82.6
LWL [5]	81.0	79.6	76.4	83.8	84.2	81.6	79.1	84.1
STA [60]	80.6	-	-	-	-	-	-	-
STM [38]	79.2	79.6	73.0	83.6	80.6	81.8	79.2	84.3
RTS (Box)	70.8	71.1	65.2	74.0	72.8	72.6	69.4	75.8
LWL (Box) [5]	-	-	-	-	-	70.6	67.9	73.3
Siam-RCNN [47]	67.3	68.1	61.5	70.8	68.8	70.6	66.1	75.0
D3S [51]	-	-	-	-	-	60.8	57.8	63.8
SiamMask [31]	52.8	60.2	45.1	58.2	47.7	56.4	54.3	58.5

our method can generate accurate segmentation masks. When using bounding boxes to predict both the initialization and segmentation masks, we outperform all other methods by a large margin. This confirms that even with our bounding-box initialization strategy, RTS produces accurate segmentation masks.

DAVIS 2017 [41] Similarly, we compare our method on the validation set of DAVIS 2017 [41], which contains 30 sequences. We do not fine tune the model for this benchmark. The results are shown in Tab. 8 and confirm the observation made above that RTS is able to generate accurate segmentation masks. Our method is competitive in the mask-initialization setup. In the box-initialization setup however, our approach outperforms all other methods in $\mathcal{J}\&\mathcal{F}$, in particular the segmentation trackers like SiamMask [51] (+16.2) and D3S [31] (+11.8).

5 Conclusion

We introduced RTS, a robust, end-to-end trainable, segmentation-driven tracking method that is able to generate accurate segmentation masks. Compared to the traditional bounding box outputs of classical visual object trackers, segmentation masks enable a more accurate representation of the target’s shape and extent. The proposed instance localization branch helps increasing the robustness of our tracker to enable reliable tracking even for long sequences of thousands of frames. Our method outperforms previous segmentation-driven tracking methods by a large margin, and it is competitive on several VOT benchmarks. In particular, we set a new state of the art on the challenging LaSOT [16] dataset with a success AUC of 69.7%. Competitive results on two VOS datasets confirm the high quality of the generated segmentation masks.

Acknowledgements This work was partly supported by uniqFEED AG and the ETH Future Computing Laboratory (EFCL) financed by a gift from Huawei Technologies.

References

1. Berman, M., Triki, A.R., Blaschko, M.B.: The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) 8
2. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: Proceedings of the European Conference on Computer Vision Workshops (ECCVW) (October 2016) 3, 4
3. Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019) 3, 5, 7, 8, 11, 12, 13
4. Bhat, G., Danelljan, M., Van Gool, L., Timofte, R.: Know your surroundings: Exploiting scene information for object tracking. In: Proceedings of the European Conference on Computer Vision (ECCV) (August 2020) 13
5. Bhat, G., Lawin, F.J., Danelljan, M., Robinson, A., Felsberg, M., Gool, L.V., Timofte, R.: Learning what to learn for video object segmentation. In: European Conference on Computer Vision ECCV (2020) 2, 3, 4, 5, 6, 9, 10, 11, 12, 13, 14
6. Bolme, D.S., Beveridge, J.R., Draper, B.A., Lui, Y.M.: Visual object tracking using adaptive correlation filters. In: CVPR (2010) 3, 7
7. Caelles, S., Maninis, K.K., Pont-Tuset, J., Leal-Taixé, L., Cremers, D., Van Gool, L.: One-shot video object segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 221–230 (2017) 4
8. Chen, X., Yan, B., Zhu, J., Wang, D., Yang, X., Lu, H.: Transformer tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2021) 4, 11, 12, 13
9. Chen, Y., Pont-Tuset, J., Montes, A., Van Gool, L.: Blazingly fast video object segmentation with pixel-wise metric learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1189–1198 (2018) 4
10. Dai, K., Zhang, Y., Wang, D., Li, J., Lu, H., Yang, X.: High-performance long-term tracking with meta-updater. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) 11
11. Danelljan, M., Bhat, G.: PyTracking: Visual tracking library based on PyTorch. <https://github.com/visionml/pytracking> (2019) 10, 11, 12, 13
12. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: ATOM: Accurate tracking by overlap maximization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019) 3, 4
13. Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M.: ECO: efficient convolution operators for tracking. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2017) 7
14. Danelljan, M., Gool, L.V., Timofte, R.: Probabilistic regression for visual tracking. In: CVPR (2020) 3, 4, 11, 12, 13
15. Danelljan, M., Robinson, A., Shahbaz Khan, F., Felsberg, M.: Beyond correlation filters: Learning continuous convolution operators for visual tracking. In: Proceedings of the European Conference on Computer Vision (ECCV) (October 2016) 3
16. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019) 1, 2, 3, 5, 8, 10, 11, 12, 14

17. Fan, H., Ling, H.: Cract: Cascaded regression-align-classification for robust visual tracking. arXiv preprint arXiv:2011.12483 (2020) [13](#)
18. Fu, Z., Liu, Q., Fu, Z., Wang, Y.: Stmtrack: Template-free visual tracking with space-time memory networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2021) [11](#), [12](#), [13](#)
19. Galoogahi, H.K., Fagg, A., Huang, C., Ramanan, D., Lucey, S.: Need for speed: A benchmark for higher frame rate object tracking. In: ICCV (2017) [10](#), [11](#), [13](#)
20. Guo, Q., Feng, W., Zhou, C., Huang, R., Wan, L., Wang, S.: Learning dynamic siamese network for visual object tracking. In: ICCV (2017) [3](#)
21. He, A., Luo, C., Tian, X., Zeng, W.: Towards a better match in siamese network based visual object tracker. In: ECCV workshop (2018) [3](#)
22. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **37**(3), 583–596 (2015) [3](#), [7](#)
23. Hu, Y.T., Huang, J.B., Schwing, A.G.: Videomatch: Matching based video object segmentation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 54–70 (2018) [4](#)
24. 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 (TPAMI) **43**(5), 1562–1577 (2021) [3](#), [8](#), [12](#)
25. Khoreva, A., Benenson, R., Ilg, E., Brox, T., Schiele, B.: Lucid data dreaming for object tracking. In: The DAVIS Challenge on Video Object Segmentation (2017) [4](#)
26. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of the International Conference on Learning Representations (ICLR) (2014) [9](#)
27. Kristan, M., Leonardis, A., Matas, J., Felsberg, M., Pflugfelder, R., Kämäräinen, J.K., Danelljan, M., Zajc, L.Č., Lukežič, A., Drbohlav, O., He, L., Zhang, Y., Yan, S., Yang, J., Fernández, G., et al: The eighth visual object tracking vot2020 challenge results. In: Proceedings of the European Conference on Computer Vision Workshops (ECCVW) (August 2020) [13](#)
28. Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., Yan, J.: Siamrpn++: Evolution of siamese visual tracking with very deep networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019) [3](#), [4](#), [12](#)
29. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) [3](#), [4](#)
30. Li, X., Change Loy, C.: Video object segmentation with joint re-identification and attention-aware mask propagation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 90–105 (2018) [4](#)
31. Lukezic, A., Matas, J., Kristan, M.: D3s - a discriminative single shot segmentation tracker. In: CVPR (2020) [4](#), [5](#), [11](#), [12](#), [13](#), [14](#)
32. Lukezic, A., Vojir, T., Zajc, L.C., Matas, J., Kristan, M.: Discriminative correlation filter tracker with channel and spatial reliability. International Journal of Computer Vision (IJCV) **126**(7), 671–688 (2018) [3](#)
33. Maninis, K.K., Caelles, S., Chen, Y., Pont-Tuset, J., Leal-Taixé, L., Cremers, D., Van Gool, L.: Video object segmentation without temporal information. IEEE Transactions on Pattern Analysis and Machine Intelligence **41**(6), 1515–1530 (2018) [4](#)

34. Mayer, C., Danelljan, M., Bhat, G., Paul, M., Paudel, D.P., Yu, F., Van Gool, L.: Transforming model prediction for tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8731–8740 (June 2022) [4](#), [11](#), [12](#), [13](#)
35. Mayer, C., Danelljan, M., Paudel, D.P., Van Gool, L.: Learning target candidate association to keep track of what not to track. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 13444–13454 (October 2021) [11](#), [12](#), [13](#)
36. Mueller, M., Smith, N., Ghanem, B.: A benchmark and simulator for uav tracking. In: Proceedings of the European Conference on Computer Vision (ECCV) (October 2016) [1](#), [5](#), [10](#), [11](#), [13](#)
37. Müller, M., Bibi, A., Giancola, S., Al-Subaihi, S., Ghanem, B.: Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In: ECCV (2018) [3](#), [12](#)
38. Oh, S.W., Lee, J.Y., Xu, N., Kim, S.J.: Video object segmentation using space-time memory networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019) [2](#), [4](#), [13](#), [14](#)
39. Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: Computer Vision and Pattern Recognition (2016) [9](#)
40. Perazzi, F., Khoreva, A., Benenson, R., Schiele, B., Sorkine-Hornung, A.: Learning video object segmentation from static images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2663–2672 (2017) [4](#)
41. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv:1704.00675 (2017) [2](#), [4](#), [8](#), [14](#)
42. Son, J., Jung, I., Park, K., Han, B.: Tracking-by-segmentation with online gradient boosting decision tree. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (December 2015) [4](#)
43. Tao, R., Gavves, E., Smeulders, A.W.M.: Siamese instance search for tracking. In: CVPR (2016) [3](#)
44. Valmadre, J., Bertinetto, L., Henriques, J., Vedaldi, A., Torr, P.H.S.: End-to-end representation learning for correlation filter based tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017) [3](#), [4](#)
45. Voigtlaender, P., Chai, Y., Schroff, F., Adam, H., Leibe, B., Chen, L.C.: Feelvos: Fast end-to-end embedding learning for video object segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9481–9490 (2019) [4](#)
46. Voigtlaender, P., Leibe, B.: Online adaptation of convolutional neural networks for video object segmentation. In: BMVC (2017) [4](#)
47. Voigtlaender, P., Luiten, J., Torr, P.H., Leibe, B.: Siam R-CNN: Visual tracking by re-detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) [2](#), [4](#), [11](#), [12](#), [13](#), [14](#)
48. Wang, G., Luo, C., Sun, X., Xiong, Z., Zeng, W.: Tracking by instance detection: A meta-learning approach. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) [3](#)
49. Wang, N., Zhou, W., Wang, J., Li, H.: Transformer meets tracker: Exploiting temporal context for robust visual tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2021) [4](#), [11](#), [12](#), [13](#)

50. Wang, Q., Teng, Z., Xing, J., Gao, J., Hu, W., Maybank, S.J.: Learning attentions: Residual attentional siamese network for high performance online visual tracking. In: CVPR (2018) [3](#)
51. Wang, Q., Zhang, L., Bertinetto, L., Hu, W., Torr, P.H.: Fast online object tracking and segmentation: A unifying approach. In: Proceedings of the IEEE conference on computer vision and pattern recognition (2019) [4](#), [14](#)
52. Wug Oh, S., Lee, J.Y., Sunkavalli, K., Joo Kim, S.: Fast video object segmentation by reference-guided mask propagation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7376–7385 (2018) [4](#)
53. Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., Huang, T.: Youtube-vos: A large-scale video object segmentation benchmark (2018) [4](#), [8](#), [9](#), [13](#), [14](#)
54. Yan, B., Wang, D., Lu, H., Yang, X.: Alpha-refine: Boosting tracking performance by precise bounding box estimation. In: CVPR (2021) [2](#), [4](#), [11](#), [12](#), [13](#)
55. Yan, B., Peng, H., Fu, J., Wang, D., Lu, H.: Learning spatio-temporal transformer for visual tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10448–10457 (October 2021) [2](#), [4](#), [11](#), [12](#), [13](#)
56. Yu, B., Tang, M., Zheng, L., Zhu, G., Wang, J., Feng, H., Feng, X., Lu, H.: High-performance discriminative tracking with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9856–9865 (October 2021) [4](#), [12](#)
57. Yu, Y., Xiong, Y., Huang, W., Scott, M.R.: Deformable siamese attention networks for visual object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) [13](#)
58. Zhang, Z., Peng, H., Fu, J., Li, B., Hu, W.: Ocean: Object-aware anchor-free tracking. In: Proceedings of the European Conference on Computer Vision (ECCV) (August 2020) [11](#)
59. Zhang, Z., Zhong, B., Zhang, S., Tang, Z., Liu, X., Zhang, Z.: Distractor-aware fast tracking via dynamic convolutions and mot philosophy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2021) [11](#)
60. Zhao, B., Bhat, G., Danelljan, M., Van Gool, L., Timofte, R.: Generating masks from boxes by mining spatio-temporal consistencies in videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 13556–13566 (October 2021) [4](#), [5](#), [8](#), [9](#), [12](#), [13](#), [14](#)
61. Zheng, L., Tang, M., Chen, Y., Wang, J., Lu, H.: Learning feature embeddings for discriminant model based tracking. In: Proceedings of the European Conference on Computer Vision (ECCV) (August 2020) [3](#)
62. Zhu, Z., Wang, Q., Bo, L., Wu, W., Yan, J., Hu, W.: Distractor-aware siamese networks for visual object tracking. In: ECCV (2018) [3](#)