

# A Discriminative Single-Shot Segmentation Network for Visual Object Tracking

Alan Lukežič<sup>1</sup>, Jiří Matas<sup>2</sup>, Matej Kristan<sup>1</sup>

<sup>1</sup>Faculty of Computer and Information Science, University of Ljubljana, Slovenia

<sup>2</sup>Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic

alan.lukezic@fri.uni-lj.si

**Abstract**—Template-based discriminative trackers are currently the dominant tracking paradigm due to their robustness, but are restricted to bounding box tracking and a limited range of transformation models, which reduces their localization accuracy. We propose a discriminative single-shot segmentation tracker – D3S<sub>2</sub>, which narrows the gap between visual object tracking and video object segmentation. A single-shot network applies two target models with complementary geometric properties, one invariant to a broad range of transformations, including non-rigid deformations, the other assuming a rigid object to simultaneously achieve robust online target segmentation. The overall tracking reliability is further increased by decoupling the object and feature scale estimation. Without per-dataset finetuning, and trained only for segmentation as the primary output, D3S<sub>2</sub> outperforms all published trackers on the recent short-term tracking benchmark VOT2020 and performs very close to the state-of-the-art trackers on the GOT-10k, TrackingNet, OTB100 and LaSoT. D3S<sub>2</sub> outperforms the leading segmentation tracker SiamMask on video object segmentation benchmarks and performs on par with top video object segmentation algorithms.

**Index Terms**—Visual object tracking, video object segmentation, discriminative tracking, single-shot segmentation.

## 1 INTRODUCTION

VISUAL object tracking is one of core computer vision problems. The most common formulation considers the task of reporting target location in each frame of the video given a single training image. Currently, the dominant tracking paradigm, performing best in evaluations [1], [2], is correlation bounding box tracking [3], [4], [5], [6], [7], [8] where the target represented by a multi-channel rectangular template is localized by cross-correlation between the template and a search region.

State-of-the-art template-based trackers apply an efficient brute-force search for target localization. Such strategy is appropriate for low-dimensional transformations like translation and scale change, but becomes inefficient for more general situations e.g. such that induce an aspect ratio change and rotation. As a compromise, modern trackers combine approximate exhaustive search with sampling and/or bounding box refinement/regression networks [9], [10] for aspect ratio estimation. However, these approaches are restricted to axis-aligned rectangles.

Estimation of high-dimensional template-based transformation is unreliable when a bounding box is a poor approximation of the target [11]. This is common – consider e.g. elongated, rotating, deformable objects, or a person with spread out hands. In these cases, the most accurate and well-defined target location model is a binary per-pixel segmentation mask. If such output is required, tracking becomes the video object segmentation task recently popularized by DAVIS [12], [13] and YoutubeVOS [14] challenges.

Unlike in tracking, video object segmentation challenges typically consider large targets observed for less than 100 frames with low background distractor presence. Top video object segmentation approaches thus fare poorly in short-term tracking scenarios [2] where the target covers a fraction of the image, substantially changes its appearance over a longer period and

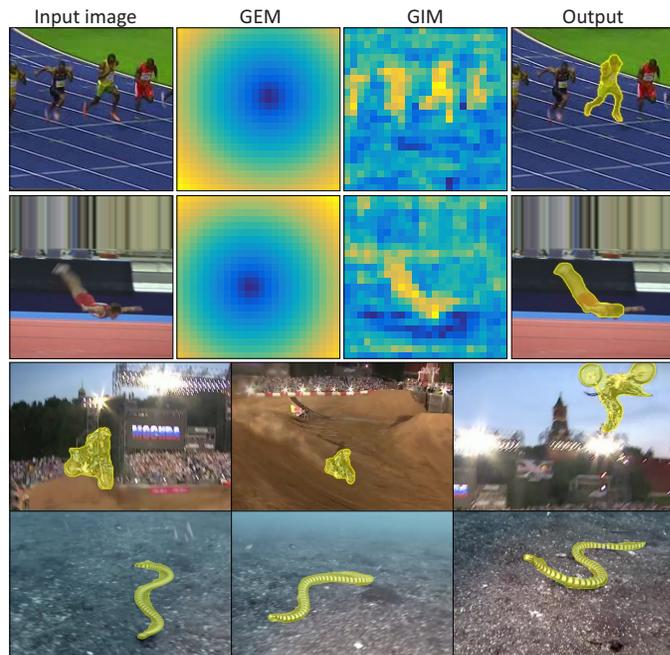


Fig. 1. The D3S<sub>2</sub> tracker represents the target by two models with complementary geometric properties, one invariant to a broad range of transformations, including non-rigid deformations (GIM - geometrically invariant model), the other assuming a rigid object with motion well approximated by an euclidean transformation (GEM - geometrically constrained Euclidean model). The D3S<sub>2</sub>, exploiting the complementary strengths of GIM and GEM, provides both state-of-the-art localisation and accurate segmentation, even in the presence of substantial deformation.

moves over a cluttered background. Best trackers apply visual model adaptation, but in the case of segmentation errors it leads to an irrecoverable tracking failure [15]. Because of this, in the past, segmentation has played only an auxiliary role in template-based trackers [16], constrained DCF learning [5] and tracking by 3D model construction [17].

Recently, the SiamRPN [8] tracker has been extended to address target localization and segmentation within a common framework called SiamMask [18]. The processing pipeline in SiamMask is split into two stages – the first stage localizes the bounding box by a SiamRPN branch and then another branch estimates the segmentation mask by only considering the features within the bounding box. Splitting the processing pipeline into two separate stages misses the opportunity to treat localization and segmentation in an end-to-end framework. In addition, a fixed template is used that cannot be discriminatively adapted to the changing scene.

We propose a new single-shot discriminative segmentation tracker, D3S<sub>2</sub>, that addresses the above-mentioned limitations. The target is encoded by two discriminative visual models – one is adaptive and highly discriminative, but geometrically constrained to an Euclidean motion (GEM), while the other is invariant to broad range of transformation (GIM, geometrically invariant model), see Figure 1.

GIM sacrifices spatial relations to allow target localization under significant deformation. On the other hand, GEM predicts only position, but discriminatively adapts to the target and acts as a selector between possibly multiple target segmentations inferred by GIM. In contrast to related trackers [9], [10], [18], [19], the primary output of D3S<sub>2</sub> is a segmentation map computed in a single pass through the network, which is trained end-to-end for segmentation only (Figure 2).

Tracking reliability depends significantly on the target scale estimation accuracy, since the latter directly affects the spatial scale of the extracted features. A standard approach in state-of-the-art is to approximate the scale by the predicted bounding box, which commonly corresponds to the visible parts of the target. Such approximation is invalid in many cases. For example, during a partial occlusion, the visible part of the target becomes smaller, but the inherent target size and the feature scale, may not change. Mismatching the scale leads to inaccurate model updates, unreliable target localization in the search region and eventual tracking failure. To address this issue, D3S<sub>2</sub> decouples the target inherent size (referred to the target scale in this paper) from the visible size (used for reporting target location in the image). The target scale is robustly predicted within the single-shot pipeline from the estimated segmentation mask, the target template and the appearance in the current frame, thus improving the feature extraction accuracy and tracking reliability. The target location in the current image is reported as the predicted segmentation mask or by fitting a bounding box to the mask if required by the evaluation protocol.

A preliminary version of the tracker was published in [20] under the name D3S. The tracker presented in this paper extends the conference version by (i) the scale estimation module SEM, (ii) trainable similarity computation and (iii) channel attention in GIM and by (iv) using a stronger backbone and (v) a higher number of features in GEM. To avoid confusion with the preliminary version, we denote the new tracker as D3S<sub>2</sub>. In addition to methodological improvements, we include the following new experiments: (i) a new experiment which exposes the ambiguity of manually

annotated bounding boxes in datasets, (ii) segmentation-based annotations for ten OTB sequences for additional segmentation performance insights, (iii) an extended ablation study and (iv) additional qualitative analysis.

D3S<sub>2</sub> outperforms all published state-of-the-art trackers on the segmentation-based tracking dataset VOT2020 [21], while ranking second among all participated trackers on the challenge. On the bounding box-based short-term tracking datasets, D3S<sub>2</sub> outperforms most state-of-the-art trackers ([22], [23]) or performs on par [24] and performs close to the state-of-the-art on a long-term bounding-box tracking dataset [25] despite being a short-term tracker not trained for bounding box tracking. In video object segmentation benchmarks [12], [13], D3S<sub>2</sub> outperforms the leading segmentation tracker [18] and performs on par with top video object segmentation algorithms (often tuned to a specific domain), yet running orders of magnitude faster than most of them. Note that D3S<sub>2</sub> is not re-trained for different benchmarks – a single pre-trained version shows remarkable generalization ability and versatility<sup>1</sup>.

## 2 RELATED WORK

Robust localization depends on the discrimination capability between the target and the background distractors. This property has been studied in depth in discriminative template trackers called discriminative correlation filters (DCF) [26]. The template learning is formulated as a (possibly nonlinear) ridge regression problem and solved by circular correlation [26], [27], [28], [29]. Extensive effort has been invested by the research community to extend the DCF robustness by increasing the training and detection range via regularization [30] and masking [31] and by patch sampling [32].

While trackers based purely on color segmentation [15], [33] are inferior to DCFs, segmentation has been used for improved DCF tracking of non-rectangular targets [11], [16]. Lukežič et al. [5] used color segmentation to constrain DCF learning and proposed a real-time tracker with hand-crafted features which achieved performance comparable to trackers with deep features. The method was extended to long-term [34] and RGB-depth tracking [17] using color and depth segmentation. Further improvements in DCF tracking considered deep features: Danelljan et al. [3] used features pre-trained for detection and Valmadre et al. [35] proposed pre-training features for DCF localization. Recently Danelljan et al. [9] proposed a deep DCF training using backpropagation, while Bhat et al. [19] extended it to the end-to-end trainable tracking architecture.

Another class of trackers, called Siamese trackers [6], [36], [37], has evolved in direction of generative templates. Siamese trackers apply a backbone pre-trained offline with general targets such that object-background discrimination is maximized by correlation between the search region and target template extracted in the first frame [6]. The template and the backbone are fixed during tracking, leading to an excellent real-time performance [2]. Several multi-stage Siamese extensions have been proposed. These include addition of region proposal networks for improved target localization accuracy [8], [10] and addition of segmentation branches [18] for accurate target segmentation. Voigtlaender et al. [38], proposed a siamese two-stage tracking architecture, which combines a siamese tracker and a detector based on region

1. A PyTorch implementation will be made publicly available on GITHUB.

proposal network. A common drawback of the siamese trackers is that template update by standard temporal averaging leads to reduced performance. This was recently addressed by proposing a template adaptation technique based on backpropagation [39]. Recent works [40], [41] have shown benefits of replacing the region proposal networks by modules inspired by anchor-free object detection networks [42].

Segmentation of moving objects is a central problem in the emerging field of video object segmentation (VOS) [12], [14]. Most recent works [43], [44], [45], [46], [47] achieve impressive results, but involve large deep networks, which often require finetuning and are slow. Hu et al. [48] and Chen et al. [49] concurrently proposed segmentation by matching features extracted in the first frame, which considerably reduces the processing time. Further advancements of the VOS field focus mostly on target representation and online adaptation. Johnander et al. [50] proposed a network architecture that learns a generative model of target and background feature distributions and runs close to real-time. Recently, Robinson et al. [51] adopted the optimizer from [9] and used the estimated target module to compute the low-resolution segmentation. An upsampling module was proposed to predict a final high-resolution segmentation mask. Bhat et al. [52] proposed a differentiable VOS architecture for predicting a target parametric model by minimizing the segmentation error at model initialization, allows online adaptation and substantially improves the segmentation accuracy. The VOS task, however, considers segmentation of large objects with limited appearance changes in short videos. Thus, these methods fare poorly on the visual object tracking task with small, fast moving objects. The work proposed in this paper aims at narrowing the gap between visual object tracking and video object segmentation.

### 3 DISCRIMINATIVE SEGMENTATION NETWORK

The D3S<sub>2</sub> architecture is summarized in Figure 2. Two models are used to robustly cope with target appearance changes and background discrimination: a geometrically invariant model (GIM) presented in Section 3.1, which is initialized from the target template, and a geometrically constrained Euclidean model (GEM) presented in Section 3.2, which adapts online to target appearance changes. These models process the input in parallel pathways and produce several coarse target presence channels, which are fused into a detailed segmentation mask by a refinement pathway described in Section 3.3. Finally, the target scale is estimated from the mask, model template and the search region by the scale estimation module (SEM) presented in Section 3.4. The template is used at an early stage of the D3S architecture (in GIM) as well as at the late stage (SEM) – D3S thus considers early as well as late template feature fusion for accurate target localization. The following subsections detail the architecture.

#### 3.1 Geometrically invariant model pathway

Accurate segmentation of a deformable target requires loose spatial constraints in the discriminative model. Our geometrically invariant model (GIM) thus removes all positional information and is composed of two sets of deep feature vectors corresponding to the target and the background, i.e.,  $\mathbf{X}_{\text{GIM}} = \{\mathbf{X}^F, \mathbf{X}^B\}$ .

The pre-trained backbone features are first processed by a  $1 \times 1$  convolutional layer to reduce their dimensionality to 64 channels, which is followed by a  $3 \times 3$  convolutional layer (a ReLU is placed after each convolutional layer). Both these layers are

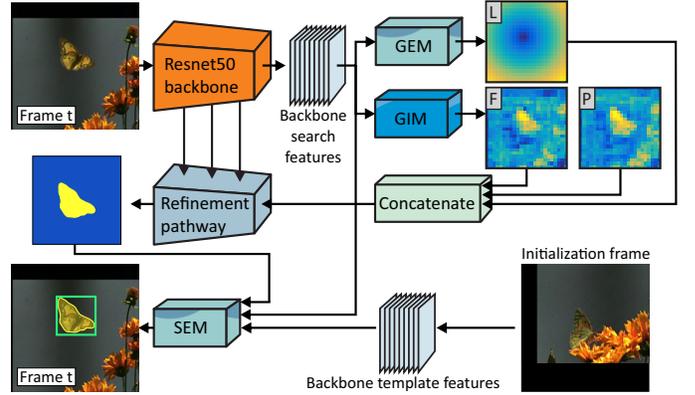


Fig. 2. The D3S<sub>2</sub> segmentation architecture. The backbone features are processed by the GEM and GIM pathways, producing the target location (L), foreground similarity (F) and target posterior (P) channels. The three channels are concatenated and refined into a detailed segmentation mask, which is used in SEM for robust target size estimation.

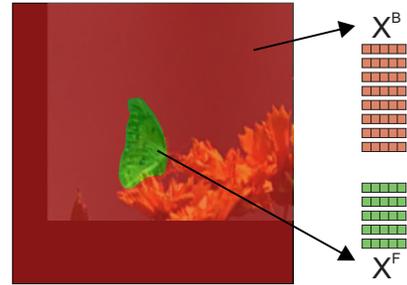


Fig. 3. Target and background feature vectors in GIM are extracted in the first frame at pixel locations corresponding to the target ( $\mathbf{X}^F$ ) and from the immediate neighborhood for the background ( $\mathbf{X}^B$ ).

adjusted in the network training stage to produce optimal features for segmentation. The target/background models are created in the first frame by extracting the segmentation feature vectors at pixel locations corresponding to the target ( $\mathbf{X}^F$ ) and from the immediate neighborhood for the background ( $\mathbf{X}^B$ ). The extraction process is shown in Figure 3.

During tracking, the pixel-level features extracted from the search region are compared to those of GIM ( $\mathbf{X}_{\text{GIM}}$ ) to compute foreground and background similarity channels  $\mathbf{F}$  and  $\mathbf{B}$  following [48]. Specifically, for the  $\mathbf{F}$  channel computation, each feature  $\mathbf{y}_i$  extracted at pixel  $i$  is compared to all features  $\mathbf{x}_j^F \in \mathbf{X}^F$  by a normalized dot product

$$s_{ij}^F(\mathbf{y}_i, \mathbf{x}_j^F) = \langle \tilde{\mathbf{y}}_i, \tilde{\mathbf{x}}_j^F \rangle, \quad (1)$$

where  $\tilde{(\cdot)}$  indicates an  $L_2$  normalization. The similarity scores computed at pixel indexed by  $i$  are sorted<sup>2</sup> in a descending order and passed to a multi-layer perceptron (MLP) for decoding into a single similarity score  $\mathbf{F}_i$ , i.e.,

$$\mathbf{F}_i = \text{MLP}(\text{sort}\{s_{ij}^F\}_{j=1:N_F}). \quad (2)$$

Note that various methods could be used to transform the set of per-pixel similarities  $s_{ij}^F$  into a single foreground similarity score

<sup>2</sup> Note that we use sorting which allows gradient backpropagation in the end-to-end network pre-training, even though the sorting operation is not differentiable on its own

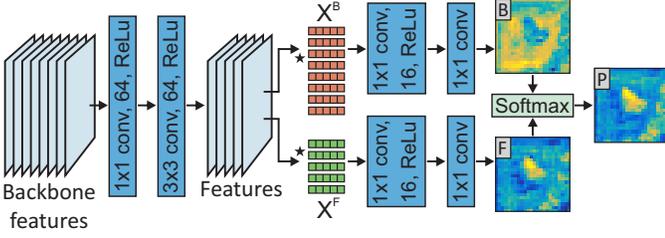


Fig. 4. GIM – the geometrically invariant model – features are matched to the features in the foreground-background model  $\{\mathbf{X}^F, \mathbf{X}^B\}$  to obtain the target (F) and background (B) similarity channels. The posterior channel (P) is the softmax of F and B.

$\mathbf{F}_i$ . We use the MLP as a good compromise between simplicity and flexibility to learn the optimal prototype set matching. In particular, we apply a two-layer MLP with  $N_F$  input nodes, 16 hidden neurons and a single output. Computation of the background similarity channel  $\mathbf{B}$  follows the same principle, but with similarities computed with the background model feature vectors, i.e.,  $\mathbf{x}_j^B \in \mathbf{X}^B$ . Finally, a softmax layer is applied to produce a target posterior channel  $\mathbf{P}$ . The GIM pathway architecture is shown in Figure 4.

### 3.2 Geometrically constrained model pathway

While GIM produces an excellent target-background separation, it cannot well distinguish the target from similar instances, leading to a reduced robustness (see Figure 1, first line). Robust localization, however, is a well-established quality of the discriminative correlation filters. Although they represent the target by a geometrically constrained model (i.e., a rectangular filter), efficient techniques developed to adapt to the target discriminative features [5], [9], [30] allow tracking reliably under considerable appearance changes.

We thus employ a recent deep DCF formulation [9] in the geometrically constrained Euclidean model (GEM) pathway. The backbone features are first reduced to 256 channels by  $1 \times 1$  convolutional layer. The reduced features are correlated by a 256 channel DCF followed by a PeLU nonlinearity [53]. The reduction layer and DCF are trained by an efficient backprop formulation (see [9] for details).

The maximum of the correlation response is considered as the most likely target position. The D3S<sub>2</sub> output (i.e., segmentation), however, requires specifying a belief of target presence at each pixel. Therefore, a target location channel is constructed by computing a (Euclidean) distance transform from the position of the maximum in the correlation map to the remaining pixels in the search region. The GEM pathway is shown in Figure 5.

### 3.3 Refinement pathway

The GIM and GEM pathways provide complementary information about the pixel-level target presence. GEM provides a robust, but rather inaccurate estimate of the target region, whereas the output channels from GIM show a greater detail, but are less discriminative (Figure 1). Furthermore, the individual outputs are low-resolution due to the backbone encoding. A refinement pathway is thus designed to combine the different information channels and upscale the solution into an accurate and detailed segmentation map. A network architecture of the refinement pathway is shown in Figure 6.

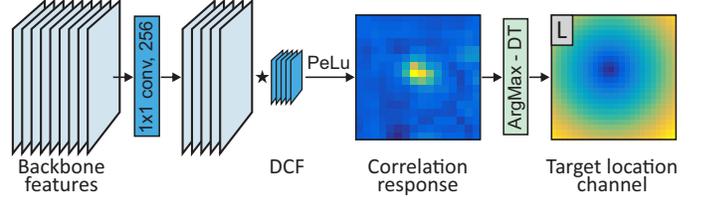


Fig. 5. GEM – the geometrically constrained Euclidean model – reduces the backbone features dimensionality and correlates them with a DCF. The target localisation channel (L) is the distance transform to the maximum correlation response, representing the per-pixel confidence of target presence.

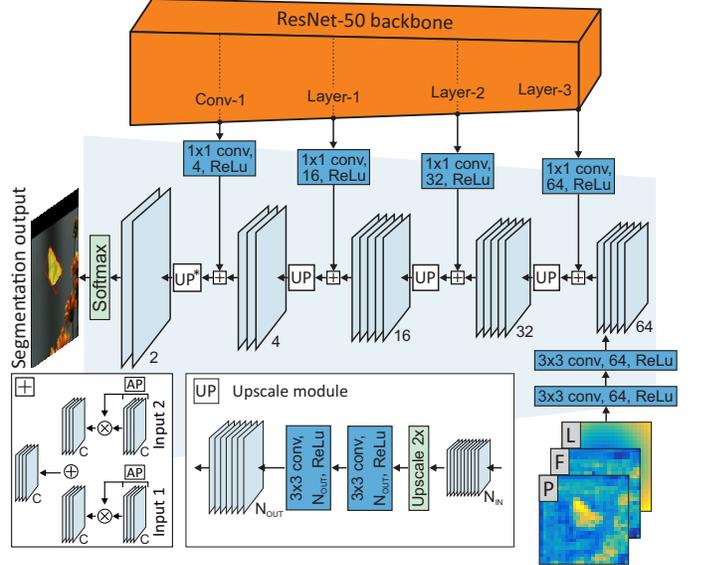


Fig. 6. The refinement pathway combines the GIM and GEM channels and gradually upscales them by using adjusted features from the backbone. The AP is an average pooling for channel attention in the feature combination process. The  $\oplus$  is summation of two tensors and  $\otimes$  represents multiplication of a tensor channels with the elements from a vector. The UP\* is a modified UP layer (see the text).

The refinement pathway takes the following inputs: the target location channel (L) from GEM and the foreground similarity and posterior channels (F and P) from the GIM. The channels are concatenated and processed by a  $3 \times 3$  convolutional layer followed by a ReLU, resulting in a tensor of 64 channels. Three stages of upscaling akin to [54], [55] are then applied to refine the details by considering the features in different layers computed in the backbone. An upscaling stage consists of doubling the resolution of the input channels, followed by two  $3 \times 3$  convolution layers (each followed by a ReLU). The resulting channels are summed with the adjusted features from the corresponding backbone layer. Specifically, the backbone features are adjusted for the upscaling task by a  $1 \times 1$  convolution layer, followed by a ReLU. A channel attention mechanism [56] is applied before summing the adjusted backbone feature tensor and the upscaled tensor from the previous refined layer. The channel attention operation involves re-weighting the tensor channels by their average pooled values. The last upscaling stage (which contains only resolution doubling, followed by a single  $3 \times 3$  convolution layer) is followed by a softmax to produce the final segmentation probability map.

### 3.4 Scale estimation module

As discussed in the introduction, D3S<sub>2</sub> decouples the *inherent target size* (i.e., the scale), responsible for determining the spatial scale for feature extraction in the model update target localization, from the *target visible size* responsible for reporting the target location in the current image. We define the target scale by an axis-aligned bounding box, which also includes all the invisible (potentially occluded) target parts. This bounding box is estimated by the scale estimation module (SEM), described in the following.

Since the inherent target location may not correspond to the visible target location, SEM predicts several possible target sizes related to different image positions along with corresponding target presence certainty scores (see Figure 7). Note that while similar modules have been proposed in [8], [10], [40], [42], their usage is different. Prior work use such modules for target localization (i.e., tracking output), while D3S<sub>2</sub> applies it for robust target scale estimation (not visible size).

The input to SEM are the segmentation mask predicted by the refinement pathway, the backbone template features and the search region features. The purpose of the mask and the template is to help focusing the network on the visible parts of the target. The predicted mask is first processed by a mask adjustment module (shown in Figure 8), to match the resolution of the backbone features and encode the mask into a form suitable for scale prediction by the downstream network. The number of channels in the template and search region features is reduced to 256 by  $1 \times 1$  convolutions. The reduced features and adjusted mask are then combined using the classification and target region heads. For simplicity, the two heads have the same architecture (Figure 9), but report a different number of output channels.

The target classification head predicts two channels, representing per-pixel target and background presence scores. Position of the maximum in the per-pixel softmax between these two channels is the most likely target position, denoted as  $\tilde{\mathbf{p}}$ . The target region head predicts four channels  $\mathbf{d}_T$ ,  $\mathbf{d}_B$ ,  $\mathbf{d}_R$  and  $\mathbf{d}_L$ , corresponding to the per-pixel distances to the top, bottom, right and left edges of the target bounding box, respectively. The target width and height ( $\tilde{w}$ ,  $\tilde{h}$ ), determining the target inherent scale, are thus defined as

$$\tilde{w} = \mathbf{d}_R(\tilde{\mathbf{p}}) - \mathbf{d}_L(\tilde{\mathbf{p}}), \quad (3)$$

$$\tilde{h} = \mathbf{d}_T(\tilde{\mathbf{p}}) - \mathbf{d}_B(\tilde{\mathbf{p}}), \quad (4)$$

where  $\mathbf{d}_{(\cdot)}(\tilde{\mathbf{p}})$  is the corresponding channel value at most likely target inherent position  $\tilde{\mathbf{p}}$ .

## 4 DISCRIMINATIVE SEGMENTATION TRACKER

This section outlines application of the discriminative segmentation network from Section 3 to online general object tracking. Given a single supervised training example from the first frame, the network produces target segmentation masks in all the remaining frames. However, some applications and most tracking benchmarks require target location represented by a bounding box. For most benchmarks, the bounding box is trivially obtained by fitting an axis-aligned bounding box that tightly fits a segmentation mask. The tracking steps are outlined in Section 4.1.

### 4.1 Tracking with D3S<sub>2</sub>

**Initialization.** D3S<sub>2</sub> is initialized on the first frame using the ground truth target location. The GEM and GIM initialization details depend on whether the target ground truth is presented by a

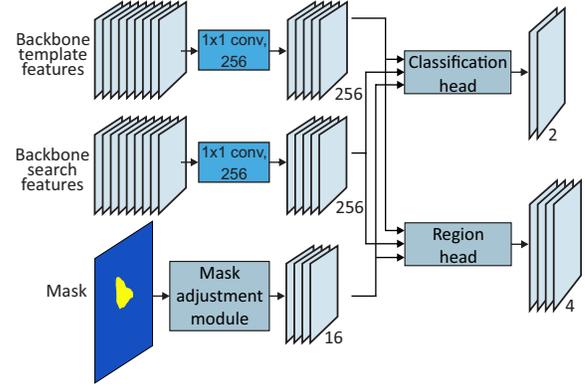


Fig. 7. The target scale estimation network (SEM) adjusts the mask by an adjustment module and combines it with the template and search region features to predict per-pixel likelihood of target inherent position along with the size using the target classification head and target region head.

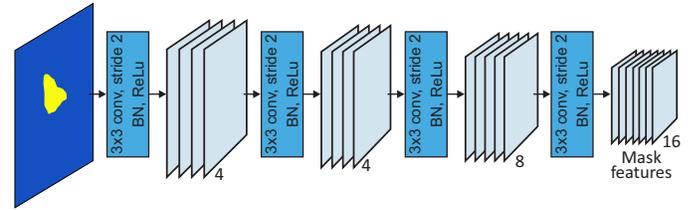


Fig. 8. The mask adjustment module extracts mask features adjusted for scale prediction.

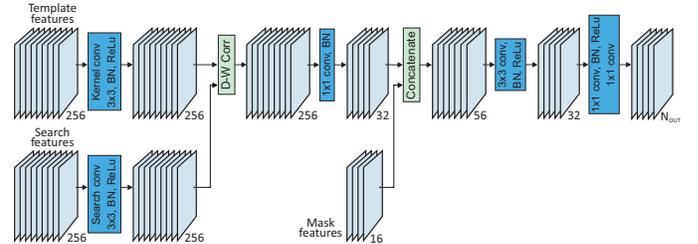


Fig. 9. The same architecture is used in the target classification head and the target region head with  $N_{OUT} = 2$  and  $N_{OUT} = 4$  output channels, respectively.

bounding box or a segmentation mask. If a ground truth bounding box is available, the GEM follows the initialization procedure proposed in [9], which involves training both the dimensionality reduction network and the DCF by backprop on the first frame by considering the region four times the target size. On the other hand, if a segmentation mask is available, the ground truth target bounding box is first approximated by an axis-aligned rectangle encompassing the segmented target.

In case a segmentation mask is available, the GIM is initialized by extracting foreground samples from the target mask and background samples from the neighborhood four times the target size. However, if only a bounding box is available, an approximate ground truth segmentation mask is constructed first. Foreground samples are extracted from within the bounding box, while the background samples are extracted from a four times larger neighborhood. A tracking iteration of D3S<sub>2</sub> is then run on the initialization region to infer a *proxi* ground truth segmentation

mask. The final foreground and background samples are extracted from this mask. This process might be iterated a few times (akin to GrabCut [57]), however, we did not observe improvements and chose only a single iteration for initialization speed and simplicity. Note that GIM initialization does not involve additional network fine tuning.

**Tracking.** During tracking, when a new frame arrives, a region four times the target size is extracted at previous target location. The region is processed by the discriminative segmentation network from Section 3 to produce the output segmentation mask. The target size is estimated next by the scale estimation module from Section 3.4. The estimated size is used to update the DCF in GEM by backprop procedure [9] and determines the target search region size in the next frame. In principle the search region size can be defined using (i) the size of the visible part of the target only, or (ii) the whole target size, including (potentially occluded) target parts. We use the latter option, which means that all training samples are consistently aligned. An axis-aligned bounding box is fitted to the mask if required by the evaluation protocol.

## 5 EXPERIMENTS

### 5.1 Implementation details

The backbone network in D3S<sub>2</sub> is composed of the first four layers of ResNet50 [58] with octave convolution architecture, pre-trained on ImageNet for object classification [59]. The backbone features are extracted from the target search region resized to  $384 \times 384$  pixels.

**Network pre-training.** The GIM pathway and the refinement pathway are pre-trained on 3471 training segmentation sequences from Youtube-VOS [14]. A training sample is constructed by uniformly sampling a pair of images and the corresponding segmentation masks from the same sequence within a range of 50 frames. To increase the diversity of training examples, we include image-based COCO [60] and randomly perturb an image to construct a training pair. Training samples in a batch are sampled from the VOS dataset with probability 80% and from COCO with probability 20% to focus on the more challenging training samples from video. To increase the robustness to possibly inaccurate GEM localization, the target location channel was constructed by perturbing ground truth locations uniformly from  $[-\frac{1}{8}\sigma, \frac{1}{8}\sigma]$ , where  $\sigma$  is target size. The network was trained by 64 image pairs batches for 40 epochs with 1000 iterations per epoch using the ADAM optimizer [61] with learning rate set to  $10^{-3}$  and with 0.2 decay every 15 epochs. The training loss was a crossentropy between the predicted and ground truth segmentation mask. The training takes 20 hours on a single GPU.

The scale estimation module (Section 3.4) is trained next on the training subsets of the following datasets: GoT-10k [22] (8335 sequences), Trackingnet [23] (30131 sequences) and LaSoT [25] (1120 sequences), where training pairs are sampled with the probability 40%, 50% and 10%, respectively. Each batch is composed of 92 pairs of image patches sampled 250 frames apart.

**Tracking speed.** D3S<sub>2</sub> is implemented in Pytorch and a single NVidia GTX 1080 GPU was used for network training and evaluation. Loading the network to a GPU and initialization of the tracker takes approximately 1.3s, while the average speed is 11.3fps. The most time-consuming parts of D3S<sub>2</sub> are feature extraction and DCF filter update in GEM. The backbone selection and the DCF complexity can thus be considered as hyperparameters that can be adjusted depending on the performance-to-

speed trade-offs application requires. To demonstrate the impact of these two components, we design a D3S<sub>2</sub> with a so-called intermediate configuration, which contains the vanilla ResNet-50 backbone and 64 channels in DCF, denoted as D3S<sub>2f</sub>. This configuration achieves 0.458 EAO on VOT2020, which represents an approximate 11% performance drop compared to D3S<sub>2</sub>, but runs at approximately 24 FPS, which is more than two times faster than D3S<sub>2</sub>. The two configurations allow user to choose one of two operating points, one providing a significantly more accurate, but slower, tracker than the original D3S, the other with slightly improved accuracy, but equally fast.

### 5.2 Evaluation on Tracking Datasets

D3S<sub>2</sub> is evaluated on the most recent tracking dataset VOT2020 [21], where targets are annotated by segmentation masks. In addition to the VOT2020, the following major bounding-box based short-term tracking datasets are used in the experimental validation: GOT-10k [22], TrackingNet [23], LaSoT [25] and OTB100 [24]. In the following we discuss the results obtained on each of the datasets.

**The VOT2020** dataset consists of 60 sequences. Targets are annotated by segmentation masks to enable a more thorough localization accuracy evaluation than on classical datasets with subjective bounding box annotations. A new VOT evaluation protocol [21] is used. The tracker is initialized on the fixed pre-defined initialization frames and let to track forward or backward through the sequence. Performance is measured by accuracy (average overlap over successfully tracked frames), robustness (percentage of the successfully tracked frames until tracking failure) and the EAO (expected average overlap), which is a principled combination of the former two measures [62].

D3S<sub>2</sub> is compared to the top-ten trackers on the most recent VOT2020 challenge: RPT, OceanPlus, AlphaRef, AFOD, LWTL, fastOcean, DET50, Ocean and D3S, which is a preliminary version of the D3S<sub>2</sub> [20]. Results reported in Table 1 show that D3S<sub>2</sub> ranks second in terms of the VOT primary measure – EAO, achieving only a 3% lower score than top-performing RPT. This is a remarkable result since most of the methods are complicated combinations of existing methods from visual tracking, video object segmentation and even object detection. For example, the top-performing RPT tracker, which substantially outperformed all other entries on VOT2020, is a combination of several existing methods like RepPoints [63], DIMP [19] and a segmentation method from the GIM model from the preliminary D3S [20] and runs at 4FPS. In contrast, D3S<sub>2</sub> performs comparably with a much cleaner and lean architecture, running over twice as fast.

In addition to the top-performing trackers on the VOT2020 challenge, we compare the proposed D3S<sub>2</sub> to the state-of-the-art published trackers on the VOT2020 dataset. The comparison includes the recent trackers Siam R-CNN [38] (denoted as RCNN), Siam++ [64], Ocean [65], SiamMask [18], DIMP [19] and ATOM [9], older tracking baselines CSRDCF [66] and SiamFC [6] and a state-of-the-art video object segmentation method STM [67]. D3S<sub>2</sub> is the top-performing method among the published state-of-the-art methods outperforming the most recent second-best performing Ocean by 19% in EAO. Results show that performance increases are due to both, accuracy and robustness, in particular D3S<sub>2</sub> outperforms the second-most robust tracker by 12% and achieves second-most accurate performance, which is 5% lower than the state-of-the-art video-object segmentation method STM [67], whose robustness is 32% lower.

	RPT	D3S <sub>2</sub>	Ocean+	AlphaRef	AFOD	LWTL	fOcean	DET50	D3S	Ocean
EAO ↑	① <b>0.530</b>	② 0.513	③ 0.491	0.482	0.472	0.463	0.461	0.441	0.439	0.430
Acc. ↑	0.700	③ 0.714	0.685	① <b>0.754</b>	0.713	② 0.719	0.693	0.679	0.699	0.693
Rob. ↑	① <b>0.869</b>	② 0.842	② 0.842	0.777	0.795	0.798	0.803	0.787	0.769	0.754

TABLE 1

Comparison with the top-ten trackers on the VOT2020 Challenge. Note that the D3S denotes the preliminary version [20] of the proposed tracker.

	D3S <sub>2</sub>	Ocean	SiamMask	STM	DIMP	ATOM	Siam++	RCNN	CSRDCF	SiamFC
EAO ↑	① <b>0.513</b>	② 0.430	③ 0.321	0.308	0.274	0.271	0.255	0.230	0.193	0.179
Acc. ↑	② 0.714	③ 0.639	0.624	① <b>0.751</b>	0.457	0.462	0.424	0.622	0.406	0.418
Rob. ↑	① <b>0.842</b>	② 0.754	0.648	0.574	③ 0.740	0.734	0.730	0.484	0.582	0.502

TABLE 2

VOT2020 – comparison with the published state-of-the-art trackers.

	RCNN	D3S <sub>2</sub>	DIMP	D3S	Siam++	ATOM	SiamMask	SiamFCv2	SiamFC	GOTURN	CCOT	MDNet
AO ↑	<b>64.9</b>	63.9	61.1	59.7	59.5	55.6	51.4	37.4	34.8	34.2	32.5	29.9
SR <sub>0.75</sub> ↑	<b>59.7</b>	50.9	49.2	46.2	47.9	40.2	36.6	14.4	9.8	12.4	10.7	9.9
SR <sub>0.5</sub> ↑	72.8	<b>73.1</b>	71.7	67.6	69.5	63.5	58.7	40.4	35.3	37.5	32.8	30.3

TABLE 3

GOT-10k test set – comparison with state-of-the-art trackers.

	RCNN	MAML	Siam++	D3S <sub>2</sub>	DIMP	SiamRPN++	D3S	SiamMask	ATOM	MDNet	CFNet	SiamFC	ECO
AUC ↑	<b>81.2</b>	75.7	75.4	75.1	74.0	73.3	72.8	72.5	70.3	60.6	57.8	57.1	55.4
Prec. ↑	<b>80.0</b>	-	70.5	69.9	68.7	69.4	66.4	66.4	64.8	56.5	53.3	53.3	49.2
Prec. <sub>N</sub> ↑	<b>85.4</b>	82.2	80.0	80.4	80.1	80.0	76.8	77.8	77.1	70.5	65.4	66.3	61.8

TABLE 4

TrackingNet test set – comparison with state-of-the-art trackers.

	RCNN	DIMP	Siam++	D3S <sub>2</sub>	MAML	ATOM	SiamRPN++	D3S	MDNet	VITAL	SiamFC	SSiam	DSiam	ECO
LaSoT	<b>65.0</b>	56.9	54.4	54.0	52.0	51.5	49.6	49.2	39.7	39.0	33.6	33.5	33.3	32.4
OTB100	<b>70.1</b>	68.4	68.3	67.2	70.4	66.9	69.6	60.8	67.8	68.2	58.2	62.1	-	69.1
OTB100s	54.9	56.6	-	<b>58.0</b>	-	56.2	52.5	52.2	49.3	49.5	41.2	-	-	45.6

TABLE 5

Comparison with state-of-the-art trackers on the LaSoT and OTB100 datasets. The OTB100s denotes a set of ten randomly selected OTB100 sequences, which we manually annotated with segmentation masks.

**GOT-10k** [22] is the recent large-scale high-diversity dataset consisting of 10k video sequences with targets annotated by axis-aligned bounding boxes. The trackers are evaluated on 180 test sequences with 84 different object classes and 32 motion patterns, while the rest of the sequences form a training set. A tracker is initialized on the first frame and let to track to the end of the sequence. Trackers are ranked according to the average overlap, but success rates (SR<sub>0.5</sub> and SR<sub>0.75</sub>) are reported at two overlap thresholds 0.5 and 0.75, respectively, for detailed analysis<sup>3</sup>. The following top-performing sota trackers are used in comparison [22]: SiamFCv2 [35], SiamFC [6], GOTURN [68], CCOT [30], MDNet [69] and the most-recent ATOM [9], SiamMask [18], DIMP [19], Siam R-CNN [38] (denoted as RCNN) and Siam++ [64]. Some of the top-performing sota trackers we use in a comparison utilize the GOT-10k training set for target localization, while the D3S<sub>2</sub> is not fine-tuned on the training set, but trains only scale estimation module on the tracking datasets.

Results on GOT-10k are reported in Table 3. D3S<sub>2</sub> outperforms most of the top-performing state-of-the-art trackers by a large

margin in all performance measures, and achieves more than 70% boost in average overlap compared to the SiamFCv2, which is a top-performer on the [22] benchmark. The recent Siam R-CNN outperforms D3S<sub>2</sub> by 1.6% in average overlap, which is mostly due to the strong target re-detection capability. It is also worth pointing out that Siam R-CNN is trained for improved localization on the GoT-10k dataset. D3S<sub>2</sub> outperforms the most recent DIMP and SiamMask trackers by 4.6% and over 24% in average overlap, respectively. This demonstrates considerable generalization ability over a diverse set of target types without fine-tuning.

**TrackingNet** is another large-scale dataset for training and testing trackers. The training set consists of over 30k video sequences, while the testing set contains 511 sequences. A tracker is initialized on the first frame and let to track to the end of the sequence. Trackers are ranked according to the area under the success rate curve (AUC), precision (Prec.) and normalized precision (Prec.<sub>N</sub>). The reader is referred to [23] for further details about the performance measures.

D3S<sub>2</sub> is compared with the top-performing sota trackers according to [23]: ECO [3], SiamFC [6], CFNet [35], MDNet [69] and the most recent sota trackers ATOM [9], SiamMask [18], SiamRPN++ [10], DIMP [19], Siam R-CNN [38] (denoted as RCNN), FCOS-MAML [70] (denoted as MAML)

3. Success rate denotes percentage of frames where predicted region overlaps with the ground-truth region more than the threshold.

and Siam++ [64]. D3S<sub>2</sub> significantly outperforms the sota reported in [23] as well as most of the recent published state-of-the-art trackers. Note that D3S<sub>2</sub> is trained for localization by segmentation only on 3471 sequences from YouTube-VOS [14], while ATOM, SiamRPN++, DIMP, MAML, Siam R-CNN and Siam++ are fine-tuned on much larger datasets (31k, and over 380k sequences, respectively), some of them include the TrackingNet training set as well. This further supports a considerable generalization capability of D3S<sub>2</sub>, which is primarily pre-trained for segmentation, not tracking.

**OTB100** is one of the first tracking datasets from 2015, containing 100 tracking sequences with targets annotated by axis-aligned bounding boxes. The dataset has largely saturated over the years and more challenging datasets have been developed since. Nevertheless, we include evaluation on this dataset since it is still widely used. A tracker is initialized on the first frame and let to track to the end of the sequence. Trackers are ranked according to the area under the success rate curve (AUC).

D3S<sub>2</sub> is compared with the trackers which are top performers on OTB: ECO [3], SSiam [71], SiamFC [6], MDNet [69] and the recent state-of-the-art trackers VITAL [72], SiamRPN++ [10], ATOM [9], DIMP [19], Siam R-CNN [38] (denoted as RCNN), Siam++ [64] and FCOS-MAML [70] (denoted as MAML). The results are presented in Table 5. D3S<sub>2</sub> performs comparably to the top-performing Siam R-CNN with approximately 4% lower AUC. Success plots in Figure 10 show that all top trackers perform very similarly for overlaps lower than 0.4. The main difference comes from the overlaps greater than 0.5. This would imply that some trackers predict the target position more accurately. However, we observed that the bounding box annotations on this dataset are ambiguous in certain cases. Figure 10 shows several examples in which the target location is in fact more accurately estimated by the tracker than the ground truth. For this reason we randomly selected 10 sequences from OTB100 and manually per-frame re-annotated them with target segmentation masks. We re-evaluated the bounding box predictions of the trackers on the segmentation ground truth. Tracking performance, measured as the average overlap (AUC), is shown in Table 5 (OTB100s). D3S<sub>2</sub> achieves the largest AUC, which shows high accuracy of its predicted bounding boxes.

In addition to comparison of the trackers, we measure the overlap between the OTB bounding box ground-truths and the segmentation annotations. The average overlap on the ten sequences is 0.572, which is comparable to the top-performing trackers. This result implies that trackers have advanced so much over the last years that the annotation accuracy on this benchmark may no longer differentiate faithfully between the performance of modern trackers.

**LaSoT** is a large scale tracking dataset with total 1400 sequences from 70 diverse categories, where 280 sequences are used for evaluation and each category is equally represented. A tracker is initialized at the beginning of the sequence and let to track to the end of the sequence. Note that since the targets frequently disappear in the sequences, the dataset is appropriate for evaluation of long-term trackers, i.e., trackers that implement target absence and re-detection capabilities. D3S<sub>2</sub> is a short-term tracker, but we nevertheless report the performance on this dataset for additional insights. Tracking performance is measured by the area under the success rate curve (AUC). D3S<sub>2</sub> is compared to the top-performing trackers according to [25] (MDNet [69], VITAL [72], SiamFC [6], StructSiam [71], DSiam [73], ECO [3])

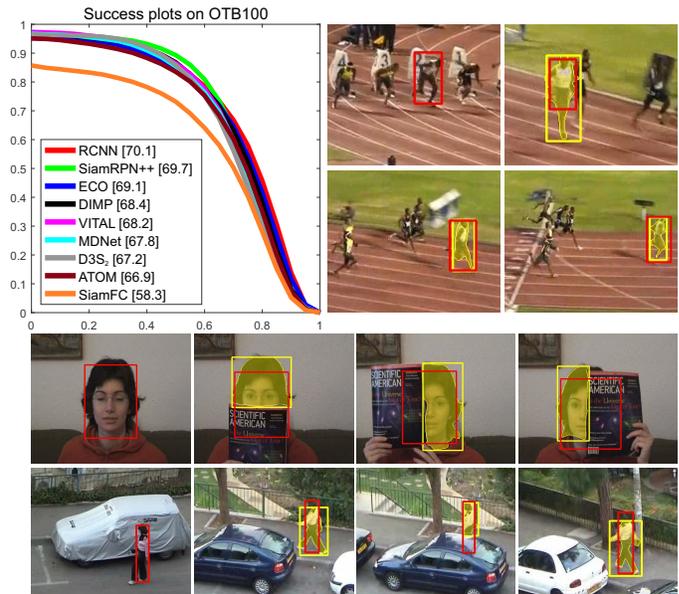


Fig. 10. Success plots of the state-of-the-art trackers on OTB100 [24] are shown on the top-left side. Qualitative examples on three sequences from OTB100 are shown top-right and at the bottom. Yellow bounding boxes and segmentation masks denote predictions made by D3S<sub>2</sub>, while red bounding boxes represent ground-truth. Note the inconsistencies and large ambiguity in ground-truth annotations.

and the recent state-of-the-art trackers (DIMP [19], ATOM [9], SiamRPN++ [10], Siam R-CNN [38] (denoted as RCNN), FCOS-MAML [70] (denoted as MAML) and Siam++ [64]). The results are shown in Table 5.

Even though LaSoT is a long-term dataset, the state-of-the-art short-term trackers perform relatively well. The reasons are in large search regions that often cover a substantial portion of the image and in that targets mostly re-appear close to their point of disappearance. D3S<sub>2</sub> ranks fourth, which is a remarkable result for a short-term tracker without target re-detection capability. The top-performing Siam R-CNN is a long-term tracker with a strong re-detection capability. Some trackers are fine-tuned on LaSoT categories for target localization, while D3S<sub>2</sub> uses LaSoT training set for training target size estimation module only. This shows great generalization capability of the proposed segmentation-based D3S<sub>2</sub>.

### 5.2.1 Qualitative analysis

The D3S<sub>2</sub> performance under challenging conditions is further visualized in Figure 11. The *gymnastics* sequence contains a rotating target that undergoes a scale change in a relatively short time. The *butterfly* significantly changes the size and shape, while keeping the inherent scale of features unchanged. Owing to decoupling of scale and visible target size, D3S<sub>2</sub> faithfully tracks and segments the target in both cases. Distractor robustness is demonstrated in the *birds* sequence, which contains several near-by deformable targets. D3S<sub>2</sub> robustly adapts to the appearance changes of the tracked bird via the GEM module and is not distracted by the other birds, while still delivering an accurate segmentation. The *diver* sequence demonstrates robustness to low-contrast scenarios, while the *deer* sequence demonstrates a challenging scenario in which only a part of the target (deer’s head) is tracked. Even though the head undergoes significant pose changes, and is partially occluded

by the deer torso, which is of a similar color, D3S<sub>2</sub> adapts to the appearance changes robustly and faithfully segments the non-occluded part of the head.

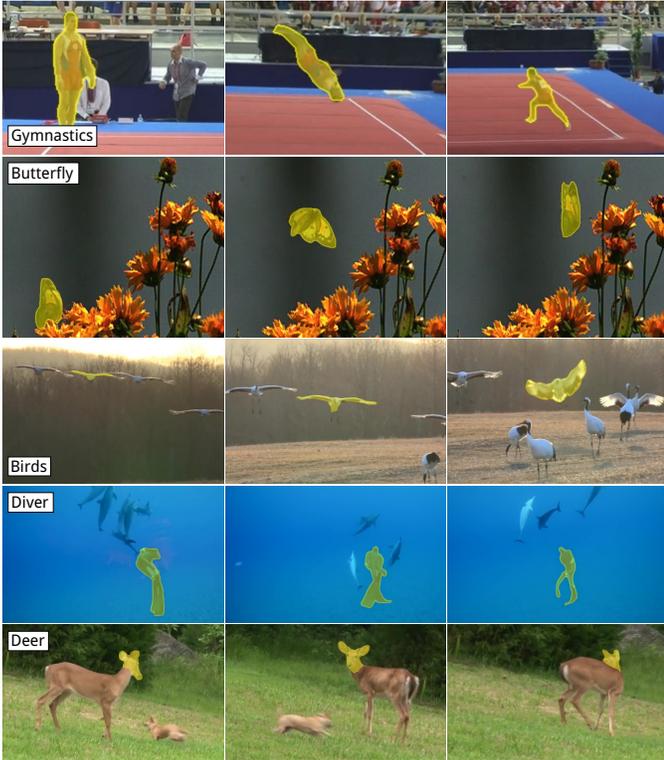


Fig. 11. Qualitative examples of D3S<sub>2</sub> segmentation masks on video sequences with various challenges: *gymnastics* contains a rotating target with significant scale changes, in *butterfly* the target shape changes dramatically, *bird* contains several nearby similar deformable targets, *diver* contains a low-contrast underwater scene, and in *deer* only part of the target under partial occlusions is tracked.

### 5.3 Ablation Study

An ablation study was performed on VOT2020 using the fixed-reset protocol [21] to analyze the contributions of different components of the D3S<sub>2</sub> architecture. The following variants were considered: (i) D3S<sub>2</sub> without the geometrically invariant module (D3S<sub>2</sub><sup>GIM</sup>); (ii) D3S<sub>2</sub> without the geometrically constrained Euclidean module (D3S<sub>2</sub><sup>GEM</sup>); (iii) D3S<sub>2</sub> without the channel attention mechanism in refinement pathway (D3S<sub>2</sub><sup>CA</sup>); (iv) D3S<sub>2</sub> without the multi-layer perceptron in GIM (D3S<sub>2</sub><sup>MLP</sup>); (v) D3S<sub>2</sub> without mask adjustment module (D3S<sub>2</sub><sup>MAM</sup>); (vi) D3S<sub>2</sub> without segmentation mask in SEM (D3S<sub>2</sub><sup>M</sup>); (vii) D3S<sub>2</sub> without target size estimation module (D3S<sub>2</sub><sup>SEM</sup>); (viii) D3S<sub>2</sub> with 64 channels in DCF, instead of 256 (D3S<sub>2</sub><sup>64C</sup>); (ix) D3S<sub>2</sub> with a vanilla ResNet50 backbone, instead of ResNet50 with octave convolution and (x) preliminary D3S version with the SEM module (D3S<sub>SEM</sub>). We also include the results of the preliminary D3S version [20] for reference.

Results of the ablation study are presented in Table 6. Removing the channel attention mechanism (D3S<sub>2</sub><sup>CA</sup>) from the refinement pathway reduces the tracking performance of the D3S<sub>2</sub> by 3%. Inspection of the accuracy/robustness measures, shows that channel attention does not crucially impact the number of

failures (robustness), but is responsible for improving the segmentation masks, thus increasing the accuracy. The improvement comes mostly from the additional channel re-weighting at feature upsampling in the refinement pathway.

To evaluate the quality of mask adjustment in the scale estimation module (SEM), the mask adjustment module is replaced by a bilinear resampling (D3S<sub>2</sub><sup>MAM</sup>). The replacement results in a 3% performance drop. A detailed inspection of the results showed that the main reason for the drop is a reduced robustness due to less accurate target scale estimation in GEM visual model update.

Segmentation mask is removed from the SEM to demonstrate its impact in the scale estimation (D3S<sub>2</sub><sup>M</sup>). The removal results in an approximate 1% performance drop, mostly due to the decreased accuracy. This results confirms the positive effect of the segmentation mask for robust scale estimation.

To measure the impact of the target scale estimation module, the target scale was estimated instead by the tracker output (i.e., fitting the bounding box to the segmentation mask). The resulting tracker, D3S<sub>2</sub><sup>SEM</sup> experiences a 4% performance drop, which is primarily due to increased number of tracking failures (an over 3% robustness drop). This strongly supports our hypothesis that robust target scale estimation, which effects mostly the visual model update process, is crucial for robust tracking and speaks in favor of the proposed scale estimation module. For a more detailed demonstration of the importance of the SEM module, we perform a per-attribute comparison of D3S<sub>2</sub><sup>SEM</sup> on the LaSoT dataset. The results presented in Table 7 show that removing SEM from D3S<sub>2</sub> reduces tracking performance under all 14 visual attributes, while the most significant performance drop is observed under fast motion and rotation. This conclusively emphasises the importance of SEM. To further demonstrate the impact of the SEM, we add it to the preliminary D3S version, and observe increase both in accuracy and robustness and in overall tracking performance improvement by over 8% in EAO.

The preliminary D3S version [20] applied a handcrafted average-top-K rule in computing the target foreground/background similarity in GIM. To evaluate the contribution of the new multi-layer perceptron method from Section 3.1, we replace it by the average-top-K rule with  $K = 3$  as in [20] (D3S<sub>2</sub><sup>MLP</sup>). The replacement results in a 9% performance drop, affecting both accuracy and robustness and confirms strength of the new trainable similarity computation module over the manually-selected top-K rule.

To demonstrate the impact of the larger number of channels in DCF compared to the preliminary version [20] the version with 64 channels (64C) is included in Table 6. Tracking performance is reduced for approximately 2%, mostly due to the lower robustness. This experiment demonstrates that adding more feature channels in GEM reduces tracking speed, but increases the overall tracking performance.

To demonstrate the impact of the backbone, the proposed D3S<sub>2</sub> using the ResNet50 with octave convolution is compared to the vanilla ResNet50 backbone. Using different backbone results in 6.6% drop in tracking performance and effects both, accuracy and robustness, which is an expected result since the backbone is shared between GIM and GEM. This result confirms that tracking performance of D3S<sub>2</sub> may be further improved by using more powerful backbones.

Finally we evaluate the individual contributions of the two visual models GIM and GEM. Removing GEM (D3S<sub>2</sub><sup>GEM</sup>) affects both accuracy and robustness and results in an approximately

	D3S <sub>2</sub>	$\overline{M}$	64C	$\overline{CA}$	$\overline{MAM}$	$\overline{SEM}$	ResNet50	$\overline{MLP}$	$\overline{GEM}$	$\overline{GIM}$	D3S <sub>SEM</sub>	D3S
EAO	0.513	0.508	0.504	0.497	0.496	0.492	0.479	0.468	0.463	0.439	0.475	0.439
Acc.	0.714	0.709	0.712	0.718	0.721	0.709	0.696	0.698	0.686	0.691	0.709	0.699
Rob.	0.842	0.840	0.832	0.826	0.821	0.813	0.820	0.803	0.805	0.767	0.805	0.769

TABLE 6

VOT2020 – ablation study. Removing: the segmentation mask from the SEM module ( $\overline{M}$ ), the channel attention mechanism in refinement pathway ( $\overline{CA}$ ), the mask adjustment module ( $\overline{MAM}$ ), the multi-layer perceptron in GIM and replacing it by the averaging top-K values ( $\overline{MLP}$ ) and the target size estimation module ( $\overline{SEM}$ ). The  $\overline{GEM}$  and  $\overline{GIM}$  denote removing the whole GEM and GIM modules, respectively. The 64C represents a version with 64 channels in DCF (instead of 256) and ResNet50 uses vanilla ResNet50 backbone, instead of ResNet50 with octave convolution. The D3S denotes the preliminary version [20] of D3S<sub>2</sub>, while D3S<sub>SEM</sub> denotes the preliminary D3S with SEM module.



Fig. 12. Tracking results of D3S<sub>2</sub> (yellow mask) compared to D3S<sub>2</sub> variations (red) with the individual modules, GIM, GEM, SEM and MAM, removed. The arrow indicates the tracked target.

10% overall performance reduction. Note, however, that even with the discriminative model removed, a solid overall tracking performance is still achieved, which indicates that GIM already affords a moderately robust target tracking, while robustness is further increased by GEM. Removing GIM (D3S<sub>2</sub><sup>GIM</sup>), however, reduces the accuracy and robustness by 3% and 9%, respectively, causing a 14% overall performance drop. This means that GIM is crucial for segmentation accuracy and localization robustness.

To gain further insights into importance of each module, we report qualitative comparison of four D3S<sub>2</sub> variants, each variant with one of the modules excluded, in Figure 12. Removing the geometrically invariant model ( $\overline{GIM}$ ) leads to over-segmentation in situations with nearby similar objects due to missing local target-specific and background-specific representation. Removing the geometrically constrained model ( $\overline{GEM}$ ) reduces capability of learning a target-specific discriminative model, which leads to drifting to a wrong target. Removing the scale estimation module ( $\overline{SEM}$ ) leads to incorrectly estimating the target size; in Figure 12, for example, the tracker region collapses from a person’s face to their eye. Removing the mask adjustment module ( $\overline{MAM}$ ) reduced the accuracy of feature scale estimation, which leads to an over-segmentation in the case of fish tracking in Figure 12.

#### 5.4 Evaluation on Segmentation Datasets

Segmentation capabilities of D3S<sub>2</sub> were analyzed on two popular video object segmentation benchmarks DAVIS16 [12] and DAVIS17 [13]. Under the DAVIS protocol, the segmentation algorithm is initialized on the first frame by a segmentation mask. The algorithm is then required to output the segmentation mask for all the remaining frames in the video. Performance is evaluated by

two measures averaged over the sequences: mean Jaccard index ( $\mathcal{J}_M$ ) and mean F-measure ( $\mathcal{F}_M$ ). Jaccard index represents a per-pixel intersection over union between the ground-truth and the predicted segmentation mask. The F-measure is a harmonic mean of precision and recall calculated between the contours extracted from the ground-truth and the predicted segmentation masks. For further details on these performance measures, the reader is referred to [12], [74].

D3S<sub>2</sub> is compared to several sota video object segmentation methods specialized to the DAVIS challenge setup: OSVOS [44], OnAVOS [45], OSMN [47], FAVOS [46], VM [48], PML [49] and A-GAME [50]. In addition, we include the most recent segmentation-based tracker SiamMask [18], which is the only published method that performs well on both, short-term tracking as well as on video object segmentation benchmarks. We include also the preliminary D3S version for a complete comparison.

Results are shown in Table 8. D3S<sub>2</sub> performs on par with most of the video object segmentation top performers on DAVIS. Compared to top performer on DAVIS2016, the performance of D3S<sub>2</sub> is 14% and 17% lower in the average Jaccard index and the F-measure, respectively. This is quite remarkable, considering that D3S<sub>2</sub> is more than 100 times faster. On DAVIS2017 the performance drop is 18% in Jaccard index and 14% in F-measure compared to the top-performer A-GAME. In addition, A-GAME was trained on YT-VOS and DAVIS training set, which demonstrates a considerable generalization capability of D3S<sub>2</sub>, which was not trained on the DAVIS training set. Furthermore, D3S<sub>2</sub> delivers a comparable segmentation accuracy as pure segmentation methods ASMN and PML, while being orders of magnitude faster, which is particularly important for many video editing applications.

D3S<sub>2</sub> also outperforms the only tracking-by-segmentation method SiamMask with respect to all measures. On average the segmentation is improved by over 3% in the Jaccard index and the contour accuracy-based F-measure. See Figure 13 for further qualitative comparison of D3S<sub>2</sub> and SiamMask on challenging targets.

The performance of D3S<sub>2</sub> is slightly lower on video object segmentation compared to its preliminary version D3S [20]. Videos in video object segmentation benchmarks are fairly short (on average less than 100 frames) with targets quite large, which means that already fairly robust trackers will not fail and performance measures will be dominated by the accuracy of object segmentation. The segmentation from D3S appears slightly better for the large objects. However, in more challenging segmentation scenarios with small objects undergoing significant appearance changes on long sequences (see Table 1), D3S<sub>2</sub> significantly outperforms D3S both in robustness and segmentation accuracy. Thus, compared to its preliminary version, D3S<sub>2</sub> only slightly sacrifices the segmentation accuracy of large objects on relatively

	AR	LR	OV	FM	FO	SV	VC	BC	RO	CM	MB	DE	PO	IL	Overall
D3S <sub>2</sub>	51.9	44.9	47.8	41.4	41.3	53.8	50.7	47.4	52.0	57.6	50.9	56.3	49.2	56.5	54.0
SEM	47.1	41.6	43.4	36.2	38.0	48.9	45.0	43.1	45.4	53.2	47.8	50.5	44.0	50.7	49.2

TABLE 7

Comparison of the D3S<sub>2</sub> and its variant without the SEM module (SEM) on the LaSoT dataset wrt. the 14 visual attributes: aspect ratio change (AR), low resolution (LR), Out-of-view (OV), fast motion (FM), full occlusion (FO), scale variation (SV), viewpoint change (VC), background clutter (BC), rotation (RO), camera motion (CM), motion blur (MB), deformation (DE), partial occlusion (PO) and illumination variation (IL).



Fig. 13. D3S<sub>2</sub> vs. SiamMask segmentation quality. *Hand*: the rigid template in SiamMask fails on a deforming target, the GIM model in D3S<sub>2</sub> successfully tracks despite a significant deformation. *Basketball*: SiamMask drifts to a similar object, while D3S<sub>2</sub> leverages the discriminative learning in GEM to robustly track the selected target. The yellow arrow denotes the target which should be tracked.

	$\mathcal{J}_M^{16}$	$\mathcal{F}_M^{16}$	$\mathcal{J}_M^{17}$	$\mathcal{F}_M^{17}$	FPS
D3S <sub>2</sub>	74.2	70.6	55.9	62.2	11.0
D3S [20]	75.4	72.6	57.8	63.8	25.0
SiamMask [18]	71.7	67.8	54.3	58.5	55.0
OnAVOS [45]	86.1	84.9	61.6	69.1	0.1
FAVOS [46]	82.4	79.5	54.6	61.8	0.8
A-GAME [50]	82.0	82.2	67.2	72.2	14.0
VM [48]	81.0	-	56.6	-	3.1
OSVOS [44]	79.8	80.6	56.6	63.9	0.1
PML [49]	75.5	79.3	-	-	3.6
OSMN [47]	74.0	72.9	52.5	57.1	8.0

TABLE 8

State-of-the-art comparison on the DAVIS16 and DAVIS17 segmentation datasets. Average Jaccard index and F-measure are denoted as  $\mathcal{J}_M^{16}$  and  $\mathcal{F}_M^{16}$  on DAVIS16 dataset and  $\mathcal{J}_M^{17}$  and  $\mathcal{F}_M^{17}$  on DAVIS17 dataset, respectively.

short video object segmentation sequences for significant robustness and segmentation performance boosts in challenging tracking scenarios.

## 6 CONCLUSION

A deep single-shot discriminative segmentation tracker – D3S<sub>2</sub> – was introduced. The tracker leverages two models from the extremes of the spectrum: a geometrically invariant model and a

geometrically restricted Euclidean model. The two models localize the target in parallel pathways and complement each other to achieve high segmentation accuracy of deformable targets and robust discrimination of the target from distractors. D3S<sub>2</sub> decouples the target scale estimation, which is related to the inherent target size, from the target region, which is related to the reported target position. The scale is robustly estimated by a new module and is used in model update and in setting the search region scale. The target position is reported as a segmentation mask or by a fitted axis-aligned bounding box if required by the evaluation method. The end-to-end trainable network architecture is the first single-shot pipeline with online adaptation that tightly connects discriminative tracking with accurate segmentation.

D3S<sub>2</sub> outperforms most state-of-the-art trackers on the VOT2020, GOT-10k and TrackingNet benchmarks and performs on par with top trackers on OTB100 and LaSoT, despite the fact that some of the tested trackers were re-trained for specific datasets. In contrast, D3S<sub>2</sub> was trained once on Youtube-VOS for segmentation only, the scale estimation module is trained on the public tracking training sequences, and the same version was used in *all benchmarks*. To cover a large range of potential applications, we present two versions of the tracker with different performance-to-speed trade-offs. The one provides a significantly more accurate, but slower performance, denoted as D3S<sub>2</sub> and the other named D3S<sub>2f</sub> with 11% lower tracking performance, while being more than two times faster. We envision future extensions of the presented architecture. One such extension might be introduction of the recently proposed covariant backbone features from [75] to further improve invariance of the geometrically constrained Euclidean module under target rotation.

Tests on DAVIS16 and DAVIS17 segmentation benchmarks show performance close to top segmentation methods while running much faster than most. D3S<sub>2</sub> significantly outperforms recent top segmentation tracker SiamMask on all benchmarks in all performance measures and contributes towards narrowing the gap between two, currently separate, domains of short-term tracking and video object segmentation, thus blurring the boundary between the two.

## ACKNOWLEDGMENTS

This work was supported by Slovenian research agency program P2-0214 and projects J2-2506 and J2-9433. A. Lukežič was financed by the Young researcher program of the ARRS. J. Matas was supported by the Czech Science Foundation grant GA18-05360S. We would also like to thank Žiga Sajovic and Tilen Nedanovski for their technical contributions and Jer Pelhan for his contribution on the OTB sequences segmentation.

©20XX IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material

for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The original version of the paper is available at:

<https://ieeexplore.ieee.org/document/9662200>

## REFERENCES

- [1] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Čehovin Zajc, T. Vojir, G. Hager, A. Lukežič, A. Eldesokey, and G. Fernández, “The visual object tracking VOT2017 challenge results,” in *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [2] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Č. Zajc, T. Vojir, G. Bhat, A. Lukežič, A. Eldesokey, G. Fernández, and et al., “The sixth visual object tracking VOT2018 challenge results,” in *Computer Vision – ECCV 2018 Workshops*, 2018, pp. 3–53.
- [3] M. Danelljan, G. Bhat, F. Shahbaz Khan, and M. Felsberg, “ECO: Efficient convolution operators for tracking,” in *Comp. Vis. Patt. Recognition*, 2017, pp. 6638–6646.
- [4] G. Bhat, J. Johnander, M. Danelljan, F. Shahbaz Khan, and M. Felsberg, “Unveiling the power of deep tracking,” in *Proc. European Conf. Computer Vision*, 2018, pp. 493–509.
- [5] A. Lukežič, T. Vojir, L. Čehovin Zajc, J. Matas, and M. Kristan, “Discriminative correlation filter with channel and spatial reliability,” in *Comp. Vis. Patt. Recognition*, 2017, pp. 6309–6318.
- [6] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr, “Fully-convolutional siamese networks for object tracking,” in *ECCV VOT Workshop*, 2016.
- [7] Z. Zhu, Q. Wang, B. Li, W. Wu, J. Yan, and W. Hu, “Distractor-aware siamese networks for visual object tracking,” in *Proc. European Conf. Computer Vision*, 2018, pp. 103–119.
- [8] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu, “High performance visual tracking with siamese region proposal network,” in *Comp. Vis. Patt. Recognition*, 2018.
- [9] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg, “ATOM: Accurate tracking by overlap maximization,” in *Comp. Vis. Patt. Recognition*, June 2019.
- [10] B. Li, W. Wu, Q. Wang, F. Zhang, J. Xing, and J. Yan, “SiamRPN++: Evolution of siamese visual tracking with very deep networks,” in *Comp. Vis. Patt. Recognition*, 2019.
- [11] A. Lukežič, L. Č. Zajc, and M. Kristan, “Deformable parts correlation filters for robust visual tracking,” *IEEE Trans. on Cyber.*, vol. 48, no. 6, pp. 1849–1861, 2017.
- [12] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung, “A benchmark dataset and evaluation methodology for video object segmentation,” in *Comp. Vis. Patt. Recognition*, 2016.
- [13] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, “The 2017 DAVIS challenge on video object segmentation,” *arXiv:1704.00675*, 2017.
- [14] N. Xu, L. Yang, Y. Fan, D. Yue, Y. Liang, J. Yang, and T. Huang, “YouTube-VOS: A large-scale video object segmentation benchmark,” *arXiv:1809.03327*, 2018.
- [15] H. Possegger, T. Mauthner, and H. Bischof, “In defense of color-based model-free tracking,” in *Comp. Vis. Patt. Recognition*, June 2015.
- [16] L. Bertinetto, J. Valmadre, S. Golodetz, O. Miksik, and P. H. S. Torr, “Staple: Complementary learners for real-time tracking,” in *Comp. Vis. Patt. Recognition*, June 2016, pp. 1401–1409.
- [17] U. Kart, A. Lukežič, M. Kristan, J.-K. Kamarainen, and J. Matas, “Object tracking by reconstruction with view-specific discriminative correlation filters,” in *Comp. Vis. Patt. Recognition*, 2019.
- [18] Q. Wang, L. Zhang, L. Bertinetto, W. Hu, and P. H. Torr, “Fast online object tracking and segmentation: A unifying approach,” in *Comp. Vis. Patt. Recognition*, June 2019.
- [19] G. Bhat, M. Danelljan, L. Van Gool, and R. Timofte, “Learning discriminative model prediction for tracking,” in *Int. Conf. Computer Vision*, 2019, pp. 6181–6190.
- [20] A. Lukežič, J. Matas, and M. Kristan, “D3s - a discriminative single shot segmentation tracker,” in *CVPR*, 2020, pp. 7131–7140.
- [21] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, J.-K. Kamarainen, L. Čehovin Zajc, M. Danelljan, A. Lukežič, O. Drbohlav, L. He, Y. Zhang, S. Yan, J. Yang, G. Fernández, and et al., “The eighth visual object tracking vot2020 challenge results,” in *European Conference on Computer Vision – ECCV 2020 Workshops*, 2020.
- [22] L. Huang, X. Zhao, and K. Huang, “GOT-10k: A large high-diversity benchmark for generic object tracking in the wild,” *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [23] M. Muller, A. Bibi, S. Giancola, S. Alsubaihi, and B. Ghanem, “TrackingNet: A large-scale dataset and benchmark for object tracking in the wild,” in *Proc. European Conf. Computer Vision*, September 2018.
- [24] Y. Wu, J. Lim, and M. H. Yang, “Object tracking benchmark,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 9, pp. 1834–1848, Sept 2015.
- [25] H. Fan, L. Lin, F. Yang, P. Chu, G. Deng, S. Yu, H. Bai, Y. Xu, C. Liao, and H. Ling, “Lasot: A high-quality benchmark for large-scale single object tracking,” in *Comp. Vis. Patt. Recognition*, 2019, pp. 5369–5378.
- [26] D. S. Bolme, J. R. Beveridge, B. A. Draper, and Y. M. Lui, “Visual object tracking using adaptive correlation filters,” in *Comp. Vis. Patt. Recognition*, 2010, pp. 2544–2550.
- [27] M. Danelljan, F. S. Khan, M. Felsberg, and J. van de Weijer, “Adaptive color attributes for real-time visual tracking,” in *Comp. Vis. Patt. Recognition*, 2014, pp. 1090–1097.
- [28] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, “High-speed tracking with kernelized correlation filters,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 3, pp. 583–596, 2015.
- [29] Y. Li and J. Zhu, “A scale adaptive kernel correlation filter tracker with feature integration,” in *Proc. European Conf. Computer Vision*, 2014, pp. 254–265.
- [30] M. Danelljan, A. Robinson, F. S. Khan, and M. Felsberg, “Beyond correlation filters: learning continuous convolution operators for visual tracking,” in *Proc. European Conf. Computer Vision*, 2016, pp. 472–488.
- [31] H. K. Galoogahi, T. Sim, and S. Lucey, “Multi-channel correlation filters,” in *Int. Conf. Computer Vision*, 2013, pp. 3072–3079.
- [32] M. Mueller, N. Smith, and B. Ghanem, “Context-aware correlation filter tracking,” in *Comp. Vis. Patt. Recognition*, 2017, pp. 1396–1404.
- [33] D. Comaniciu, V. Ramesh, and P. Meer, “Kernel-based object tracking,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 5, pp. 564–577, May 2003.
- [34] A. Lukežič, L. Čehovin Zajc, T. Vojir, J. Matas, and M. Kristan, “FuCoLoT - a fully-correlational long-term tracker,” in *ACCV*, 2018.
- [35] J. Valmadre, L. Bertinetto, J. Henriques, A. Vedaldi, and P. H. S. Torr, “End-to-end representation learning for correlation filter based tracking,” in *Comp. Vis. Patt. Recognition*, 2017, pp. 2805–2813.
- [36] R. Tao, E. Gavves, and A. W. M. Smeulders, “Siamese instance search for tracking,” in *Comp. Vis. Patt. Recognition*, 2016, pp. 1420–1429.
- [37] A. He, C. Luo, X. Tian, and W. Zeng, “A twofold siamese network for real-time object tracking,” in *Comp. Vis. Patt. Recognition*, 2018.
- [38] P. Voigtlaender, J. Luiten, P. H. Torr, and B. Leibe, “Siam r-cnn: Visual tracking by re-detection,” in *Comp. Vis. Patt. Recognition*, 2020, pp. 6578–6588.
- [39] P. Li, B. Chen, W. Ouyang, D. Wang, X. Yang, and H. Lu, “GradNet: Gradient-guided network for visual object tracking,” in *Int. Conf. Computer Vision*, 2019.
- [40] Z. Chen, B. Zhong, G. Li, S. Zhang, and R. Ji, “Siamese box adaptive network for visual tracking,” in *Comp. Vis. Patt. Recognition*, 2020, pp. 6668 – 6677.
- [41] D. Guo, J. Wang, Y. Cui, Z. Wang, and S. Chen, “Siamcar: Siamese fully convolutional classification and regression for visual tracking,” in *Comp. Vis. Patt. Recognition*, 2020, pp. 6269 – 6277.
- [42] Z. Tian, C. Shen, H. Chen, and T. He, “Fcos: Fully convolutional one-stage object detection,” in *Int. Conf. Computer Vision*, 2019, pp. 9627 – 9636.
- [43] P. Voigtlaender, Y. Chai, F. Schroff, H. Adam, B. Leibe, and L.-C. Chen, “FEELVOS: Fast end-to-end embedding learning for video object segmentation,” in *Comp. Vis. Patt. Recognition*, 2019.
- [44] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixe, D. Cremers, and L. Van Gool, “One-shot video object segmentation,” in *Comp. Vis. Patt. Recognition*, 2017.
- [45] P. Voigtlaender and B. Leibe, “Online adaptation of convolutional neural networks for video object segmentation,” in *Proc. British Machine Vision Conference*, 2017.
- [46] J. Cheng, Y.-H. Tsai, W.-C. Hung, S. Wang, and M.-H. Yang, “Fast and accurate online video object segmentation via tracking parts,” in *Comp. Vis. Patt. Recognition*, 2018.
- [47] L. Yang, Y. Wang, X. Xiong, J. Yang, and A. K. Katsaggelos, “Efficient video object segmentation via network modulation,” in *Comp. Vis. Patt. Recognition*, 2018.
- [48] Y.-T. Hu, J.-B. Huang, and A. G. Schwing, “VideoMatch: Matching based Video Object Segmentation,” in *Proc. European Conf. Computer Vision*, 2018.
- [49] Y. Chen, J. Pont-Tuset, A. Montes, and L. Van Gool, “Blazingly fast video object segmentation with pixel-wise metric learning,” in *Comp. Vis. Patt. Recognition*, 2018.

- [50] J. Johnander, M. Danelljan, E. Brissman, F. S. Khan, and M. Felsberg, "A generative appearance model for end-to-end video object segmentation," in *Comp. Vis. Patt. Recognition*, 2019, pp. 8945–8954.
- [51] A. Robinson, F. J. Lawin, M. Danelljan, F. S. Khan, and M. Felsberg, "Learning fast and robust target models for video object segmentation," in *Comp. Vis. Patt. Recognition*, 2020, pp. 7404–7413.
- [52] G. Bhat, F. J. Lawin, M. Danelljan, A. Robinson, M. Felsberg, L. V. Gool, and R. Timofte, "Learning what to learn for video object segmentation," in *Proc. European Conf. Computer Vision*, 2020, pp. 777–794.
- [53] L. Trottier, P. Gigu, B. Chaib-draa *et al.*, "Parametric exponential linear unit for deep convolutional neural networks," in *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2017, pp. 207–214.
- [54] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, 2015, pp. 234–241.
- [55] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár, "Learning to refine object segments," in *Proc. European Conf. Computer Vision*, 2016, pp. 75–91.
- [56] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang, "Learning a discriminative feature network for semantic segmentation," in *Comp. Vis. Patt. Recognition*, 2018, pp. 1857–1866.
- [57] C. Rother, V. Kolmogorov, and A. Blake, "Grabcut-interactive foreground extraction using iterated graph cuts," *ACM Transactions on Graphics (SIGGRAPH)*, August 2004.
- [58] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Comp. Vis. Patt. Recognition*, 2016, pp. 770–778.
- [59] Y. Chen, H. Fan, B. Xu, Z. Yan, Y. Kalantidis, M. Rohrbach, S. Yan, and J. Feng, "Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution," in *Int. Conf. Computer Vision*, 2019, pp. 3434–3443.
- [60] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Proc. European Conf. Computer Vision*, 2014, pp. 740–755.
- [61] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [62] M. Kristan, J. Matas, A. Leonardis, M. Felsberg, L. Cehovin, G. Fernandez, T. Vojir, G. Hager, G. Nebehay, and R. et al. Pflugfelder, "The visual object tracking VOT2015 challenge results," in *Int. Conf. Computer Vision*, 2015.
- [63] Z. Yang, S. Liu, H. Hu, L. Wang, and S. Lin, "Reppoints: Point set representation for object detection," in *Int. Conf. Computer Vision*, 2019, pp. 9656–9665.
- [64] Y. Xu, Z. Wang, Z. Li, Y. Yuan, and G. Yu, "Siamfc++: Towards robust and accurate visual tracking with target estimation guidelines," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 12 549–12 556.
- [65] Z. Zhang, H. Peng, J. Fu, B. Li, and W. Hu, "Ocean: Object-aware anchor-free tracking," in *Proc. European Conf. Computer Vision*, 2020.
- [66] A. Lukežič, T. Vojř, L. Čehovin Zajc, J. Matas, and M. Kristan, "Discriminative correlation filter tracker with channel and spatial reliability," *Int. J. Comput. Vision*, vol. 126, no. 7, p. 671–688, 2018.
- [67] S. W. Oh, J. Lee, N. Xu, and S. J. Kim, "Video object segmentation using space-time memory networks," in *Int. Conf. Computer Vision*, 2019, pp. 9225–9234.
- [68] D. Held, S. Thrun, and S. Savarese, "Learning to track at 100 fps with deep regression networks," in *Proc. European Conf. Computer Vision*, 2016, pp. 749–765.
- [69] H. Nam and B. Han, "Learning multi-domain convolutional neural networks for visual tracking," in *Comp. Vis. Patt. Recognition*, June 2016, pp. 4293–4302.
- [70] G. Wang, C. Luo, X. Sun, Z. Xiong, and W. Zeng, "Tracking by instance detection: A meta-learning approach," in *Comp. Vis. Patt. Recognition*, 2020, pp. 6288–6297.
- [71] Y. Zhang, L. Wang, J. Qi, D. Wang, M. Feng, and H. Lu, "Structured siamese network for real-time visual tracking," in *Proc. European Conf. Computer Vision*, 2018, pp. 355–370.
- [72] Y. Song, C. Ma, X. Wu, L. Gong, L. Bao, W. Zuo, C. Shen, R. Lau, and M.-H. Yang, "Vital: Visual tracking via adversarial learning," in *Comp. Vis. Patt. Recognition*, 2018, pp. 8990–8999.
- [73] Q. Guo, W. Feng, C. Zhou, R. Huang, L. Wan, and S. Wang, "Learning dynamic siamese network for visual object tracking," in *Int. Conf. Computer Vision*, 2017, pp. 1763–1771.
- [74] D. R. Martin, C. C. Fowlkes, and J. Malik, "Learning to detect natural image boundaries using local brightness, color, and texture cues," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 5, pp. 530–549, 2004.
- [75] D. K. Gupta, D. Arya, and E. Gavves, "Rotation equivariant siamese networks for tracking," in *Comp. Vis. Patt. Recognition*, 2021, pp. 12 362–12 371.



**Alan Lukežič** received the Ph.D. degree from the Faculty of Computer and Information Science, University of Ljubljana, Slovenia in 2021. He is currently with the Visual Cognitive Systems Laboratory, Faculty of Computer and Information Science, University of Ljubljana, as a Teaching Assistant and a Researcher. His research interests include computer vision, data mining and machine learning.



**Jiří Matas** is a professor at the Center for Machine Perception, Czech Technical University in Prague. He holds a PhD degree from the University of Surrey, UK (1995). He has published more than 200 papers in refereed journals and conferences. His publications have approximately 34000 citations in Google Scholar and 13000 in the Web of Science. His h-index is 65 (Google scholar) and 43 (Clarivate Analytics Web of Science) respectively. He received the best paper prize at the British Machine Vision Conferences in 2002 and 2005, at the Asian Conference on Computer Vision in 2007 and at Int. Conf. on Document Analysis and Recognition in 2015. He is on the editorial board of IJCV and was an Associate Editor-in-Chief of IEEE T. PAMI. His research interests include visual tracking, object recognition, image matching and retrieval, sequential pattern recognition, and RANSAC-type optimization methods.



**Matej Kristan** received a Ph.D from the Faculty of Electrical Engineering, University of Ljubljana in 2008. He is a full professor and a vice chair of the department of artificial intelligence at the Faculty of Computer and Information Science. He is president of the IAPR Slovenian pattern recognition society and an Associate Editor of IJCV. His research interests include visual object tracking, anomaly detection, object detection and segmentation, perception methods for autonomous boats and physics-informed machine learning.