

Visual Object Tracking with Discriminative Filters and Siamese Networks: A Survey and Outlook

Sajid Javed, Martin Danelljan, Fahad Shahbaz Khan, *IEEE Senior Member*, Muhammad Haris Khan, Michael Felsberg, *IEEE Senior Member*, and Jiri Matas, *IEEE Senior Member*

Abstract—Accurate and robust visual object tracking is one of the most challenging and fundamental computer vision problems. It entails estimating the trajectory of the target in an image sequence, given only its initial location, and segmentation, or its rough approximation in the form of a bounding box. Discriminative Correlation Filters (DCF) and deep Siamese Networks (SNs) have emerged as dominating tracking paradigms, which have led to significant progress. Following the rapid evolution of visual object tracking in the last decade, this survey presents a systematic and thorough review of more than 90 DCFs and Siamese trackers, based on results in nine tracking benchmarks. First, we present the background theory of both the DCF and Siamese tracking core formulations. Then, we distinguish and comprehensively review the shared as well as specific open research challenges in both these tracking paradigms. Furthermore, we thoroughly analyze the performance of DCF and Siamese trackers on nine benchmarks, covering different experimental aspects of visual tracking: datasets, evaluation metrics, performance, and speed comparisons. We finish the survey by presenting recommendations and suggestions for distinguished open challenges based on our analysis.

Index Terms—Visual Object Tracking, Discriminative Correlation Filters, Siamese Networks.

1 INTRODUCTION

VISUAL Object Tracking (VOT) is one of the fundamental open problems in computer vision. The task is to estimate the trajectory and state of a target in an image sequence. VOT has a wide range of applications, including autonomous driving, robotics, intelligent video surveillance, sports analytic and medical imaging, where it typically plays an important role within large intelligent systems. Given the initial state of any arbitrary target object, the main challenge in VOT is to learn an appearance model to be used when searching for the target object in subsequent frames. In recent years, VOT has received considerable attention, much thanks to the introduction of a variety of tracking benchmarks such as, TrackingNet [101], VOT2018 [67], and GOT-10K [60]. Despite the recent progress, VOT is still an open research problem and is perhaps more active than ever [40], [77].

The core challenge in generic object tracking¹ is to learn an appearance model of an arbitrary target object online, given only its initial state. Several real-world factors, complicates learning an accurate appearance model. For instance, the target object may undergo a partial or full occlusion, scale variation, and deformation. Moreover, there are environmental factors, including illumination changes and motion blur, which influence the appearance of the target. Another factor is that scenes often include objects or background structures having similar appearance that can easily be




Fig. 1: Performance improvements from 2014 to 2020 on the popular VOT benchmark series. Tracking quality is measured in terms of Expected Average Overlap (EAO). We show DCF and Siamese trackers that significantly advanced the SOTA. The representative trackers are: SAMF [79], DSST [27], KCF [53], SRDCF [30], SiamFC [4], CCOT [23], CSR-DCF [92], ECO [25], DaSiamRPN [174], DRT [120], SiamMask [138], DiMP [5], D3S [91] and SiamAttn [158].

confused with the target itself [140]. To address these challenges, a plethora of trackers have been proposed in the literature that have contributed to advance the State-Of-The-Art (SOTA) in tracking.

In the past decade, Discriminative Correlation Filters (DCF) and deep Siamese Networks (SNs) have been the two most prominent paradigms for VOT [4], [40], [52], [127]. In DCF-based tracking, a correlation filter is trained online on the region of interest by minimizing a least squares loss. The target is then detected in consecutive frames by convolving the trained filter via the Fast Fourier Transform (FFT) [10], [53]. In the deep Siamese tracking framework, an embedding space is learned offline by maximizing

- S. Javed is with the EECS department, Khalifa University of Science and Technology, P.O Box : 127788, Abu Dhabi, UAE. (email: sajid.javed@ku.ac.ae).
- M. Danelljan is with the Computer Vision Lab, Dept. of Information Technology and Electrical Engineering, ETH Zürich, Switzerland.
- F. S. Khan and M. H. Khan are with computer vision department, MBZUAI, Abu Dhabi, UAE.
- M. Felsberg and F. S. Khan are with computer vision laboratory, Linköping University, Sweden.
- J. Matas is with Center for Machine Perception, Czech Technical University, Prague.

¹Generic object tracking refers to the problem, where the tracked object is not known a priori and it is not constrained to be from a specific class.

the distance between the target and background appearance, while minimizing the distance between the two patches from the target itself. The SNs consists of two identical sub-branches: one for the target template and the other for the target search region. The network takes both template and search regions as inputs and outputs the local similarity with the target for each location in the search region. The SNs were initially proposed for the signature verification task [12] and later explored for visual tracking by Tao *et al.* [127] and Bertinetto *et al.* [4]. With the design of DCFs [31], [54], [160] and SNs [4], [52], [127], the tracking community has been largely focusing on these two paradigms in recent years as these two frameworks have significantly improved the tracking performance on several datasets such as the performance improvements on VOT datasets [47], [67] as shown in Fig. 1.

In this work, we present a systematic review of popular DCF and Siamese-based tracking paradigms. Both paradigms share the same objective, that is to learn an accurate target appearance model that can effectively discriminate the target object from the background. While emerging from different underlying paradigms for addressing the aforementioned objective, the advent of deep learning has brought several important similarities and common challenges to these two paradigms. For instance, (i) *Feature Representation*: Both paradigms exploit different features representations to estimate target translations and scale variations. Leveraging deep feature representations extracted from pre-trained networks is a recent trend shared by both paradigms. However, the choice of deep architecture and feature hierarchies is still an open problem in both tracking paradigms. (ii) *Target State Estimation*: The core formulation of both DCF and Siamese trackers only addresses how to estimate the translation of the target object. Hence, neither paradigm provides an explicit method for estimating the full target state e.g., handling of elongated objects parameterized by a bounding box, which is crucial in most applications. (iii) *Offline Training*: While initially only Siamese trackers benefited from end-to-end offline training, recent DCF trackers [5], [24] also leverage large-scale offline learning, integrating it with efficient and differentiable online learning modules for robust tracking.

While having common attributes, these two popular paradigms also have specific issues. For instance, (i) *Boundary Artifacts*: DCFs trackers generally exploit the periodic assumption of training samples for learning an online classifier, which introduces undesirable boundary effects that severely degrade the quality of the target model. (ii) *Optimization*: The loss function minimization also introduces challenges in DCF trackers, especially when target-specific constraints, such as spatial or temporal, etc., are regularized within the regression loss. (iii) *Online Model Adaptability*: When the target appearance changes due to variation in lighting conditions or fast motion etc., the learned model is expected to cope with these variations. The DCF trackers possess the ability to update the appearance model over time through the loss function. On the first frame, the DCF builds a model of the foreground and background efficiently, via FFT, guaranteeing a Gaussian peak response around the target. On the other hand, Siamese trackers do not inherit such a mechanism for online model update. Siamese trackers either rely on the suitability of the pre-trained feature space, and the cross-correlation in it, or adapt to the current object and background via fine-tuning the deep network, which is computationally expensive operation. Therefore, online adaptability is an important issue in Siamese trackers.

Moreover, there are also some notions which can be cross-fertilized between both paradigms. For example, deep anchor-

based or anchor-free bounding box regression can be utilized within the end-to-end DCFs tracking pipeline. DCFs trackers have utilized the potential of deeper networks. ResNet-driven DCFs trackers could further be explored for robust appearance modeling. While only the DCFs trackers have the capabilities to update the appearance model online this component could also be employed in the Siamese tracking pipeline for temporal appearance modeling. Robust multi-channel features fusion is employed in Siamese trackers in which the weights are optimized offline. DCFs trackers lack such ability and so could potentially harness this technique. Siamese trackers could benefit from spatial or spatio-temporal regularization's to penalize the background pixels while computing the correlation response map within the cross-correlation layer. These insights open up new possibilities to develop powerful tracking framework based on the best of both paradigms. In the paper, we systematically review the popular DCF and Siamese tracking paradigms with the following contributions:

- We distinguish and thoroughly review the shared as well as specific open challenges in both tracking paradigms.
- We present an extensive background theory of both DCF and Siamese tracking core formulations.
- We present a comprehensive overview of more than 90 DCF and Siamese-based trackers in the literature, presenting a comparison on nine tracking benchmarks.
- Based on our analysis, we present our recommendations and suggestions for specific open challenges.

To the best of our knowledge, we are the first to present a comprehensive survey of two most popular tracking paradigms in the last decade, describing their respective background theory, specifying their shared as well as specific open research issues along with providing a set of recommendations and future directions.

The rest of this paper is organized as follows: Section 2 focuses on other surveys on tracking. In Sections 3, we describe the core DCF and Siamese tracking formulation. Then, we present a comprehensive overview of DCF and Siamese-based trackers, respectively, by distinguishing common and distinct open key research issues in Sections 4 & 5. Experimental evaluations are presented in Sections 6. We conclude our survey and provide further research directions in Section 7.

2 LITERATURE REVIEW

In the literature, numerous survey studies on VOT have been published in the past two decades [40], [76], [117], [155], [163]. In the first survey, Yilmaz *et al.* presented systematic analysis of the whole tracking procedure [155]. The tracking methods were categorized into: features correspondence, primitive geometric models, and contour methods. Smeulders *et al.* presented an experimental survey of 19 different online trackers that emerged from 1999 to 2012 with multiple evaluation metrics on ALOV++ dataset [117]. Zhang *et al.* presented sparse coding-based trackers survey in which the tracking methods were categorized into sparse coding and sparse representation-based trackers [163]. Recently, Li *et al.* presented a survey of deep trackers and provided an experimental comparison [76]. The trackers were classified and evaluated in terms of network structure, network function, and network training. Fiaz *et al.* reviewed SOTA DCFs and non DCF-based trackers, providing comparative study based on the feature extraction methods [40].

The main differences between this survey and previous VOT surveys are as follows. Unlike previous VOT surveys, our work focuses solely on the two best-performing tracking paradigms, DCFs

and SNs, in recent years. We present an extensive background theory of both DCF and Siamese tracking core formulations. We then provide an extensive overview of more than 90 DCF and Siamese trackers and the evolution of these two paradigms to segmentation-based tracking. While previous surveys are based on attributes-based taxonomy, deep network structure and training or introducing a new dataset, we present a comprehensive overview of DCF and Siamese trackers by distinguishing shared as well as specific *open research challenges* in these two popular tracking paradigms. Furthermore, to the best of our knowledge, we are the first to extensively compare the performance of DCF and Siamese trackers on *nine* popular visual tracking benchmarks.

3 DCFs AND SIAMESE TRACKING PARADIGMS

In this section, we discuss the two popular tracking paradigms and their core formulations which have attracted significant attention in the literature during the last decade.

3.1 Discriminative Correlation Filters

Discriminative correlation filters (DCF) is a supervised technique for learning a linear regressor. In recent years, DCF-based trackers have demonstrated excellent performance on multiple tracking benchmarks. The key to the DCF success is the computationally efficient approximation to dense sampling achieved by circularly shifting the training samples, which allows the Fast Fourier Transform (FFT) to be employed when learning and applying the correlation filter. By utilizing the properties of the Fourier transform, a DCF learns a correlation filter online to localize the target object in consecutive frames by efficiently minimizing a least-squares output error. In order to estimate the target location in the next frame, the learned filter is then applied to the region of interest in which the location of the maximum response estimates the target location. The filter is then updated in an iterative manner by annotating a new sample with this estimate.

3.1.1 Standard Single-Channel DCF Formulation

In the standard single-channel DCF formulation, the aim is to learn a single channel convolution or correlation filter w from a set of training samples $\{(x_j, y_j)\}_{j=1}^m$. In its most basic form [10], each training sample x_j consists of a grayscale feature map extracted from an image region. All samples are assumed to have the same spatial size $N_1 \times N_2$. At each spatial location $(n_1, n_2) \in \Omega := \{0, \dots, N_1 - 1\} \times \{0, \dots, N_2 - 1\}$, we thus have an intensity value $x_j(n_1, n_2) = x_j(n) \in \mathbb{R}$, where $n = (n_1, n_2)$. The desired output y_j is a scalar valued function over the domain Ω , which contains a label for each location in the sample x_j . Typically, y_j is set to a sampled Gaussian function with a narrow peak that is centered on the target.

The DCF aims to learn a filter w , of the same spatial size $N_1 \times N_2$, such that $x_j * w \approx y_j$ for all samples j in the dataset. Importantly, the standard DCF formulation employs circular convolution,²

$$x * w(n_1, n_2) = \sum_{l_1=0}^{N_1-1} \sum_{l_2=0}^{N_2-1} x((n_1 - l_1)_{N_1}, (n_2 - l_2)_{N_2}) w(l_1, l_2). \quad (1)$$

It is computed by cyclically shifting the filter w (or equivalently the sample x) as ensured by the modulo operation $(a)_N = a \bmod N$. Circular convolution (1) can also be seen as first padding the sample x by periodic repetition, followed by regular convolution with the filter w .

²Some authors instead employ circular correlation, obtained by simply reflecting the convolution filter w .




Fig. 2: The tracking pipeline of a classical discriminative correlation filter.

Since the convolution operation is linear, we can express it in matrix form $x * w = \mathcal{C}(x)w$. With slight abuse of notation in order to avoid clutter, w is interpreted as a vectorization of the original filter on the right-hand side of the equation. The matrix $\mathcal{C}(x)$ contains a highly regular structure, containing circulant block matrices, where the blocks themselves are also circulant [57]. Each row in $\mathcal{C}(x)$ consists of a cyclicly shifted version of the original patch x . The convolution theorem implies that a circulant matrix can be diagonalized as,

$$\mathcal{C}(x) = \mathcal{F} \text{diag}(\hat{x}) \mathcal{F}^H \text{ and } \mathcal{C}(x)^\top = \mathcal{F} \text{diag}(\hat{\bar{x}}) \mathcal{F}^H, \quad (2)$$

where \mathcal{F} represents the Discrete Fourier Transform (DFT) matrix. We denote the conjugate of x by \bar{x} and its Fourier transform $\mathcal{F}^H x$ by \hat{x} , where H denotes the conjugate transpose.

In order to find the optimal filter w , we formulate a linear least-squares regression problem. For convenience, we first concatenate all samples j in the dataset to get the data matrix X and target vector \mathbf{y} as,

$$X = \begin{pmatrix} \mathcal{C}(x_1) \\ \vdots \\ \mathcal{C}(x_m) \end{pmatrix}, \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_m \end{pmatrix}. \quad (3)$$

The DCF is learned by minimizing the squared error objective $L(w)$ over the samples of X and their regression targets \mathbf{y} as,

$$L(w) = \|Xw - \mathbf{y}\|^2 + \lambda \|w\|^2, \quad (4)$$

Here, λ is a regularization parameter. As this is a linear least squares problem, the solution is given by the following closed form expression

$$w = (X^\top X + \lambda I)^{-1} X^\top \mathbf{y}. \quad (5)$$

To compute weights of the filter w , we simply substitute (3) into (5)

$$w_m = \left(\sum_{j=1}^m \mathcal{C}(x_j)^H \mathcal{C}(x_j) + \lambda I \right)^{-1} \sum_{j=1}^m \mathcal{C}(x_j)^H y_j. \quad (6)$$

Since the matrix $\mathcal{C}(x_j)$ is circulant, it can be efficiently diagonalized using (2) as,

$$w_m = \mathcal{F} \left(\sum_{j=1}^m \text{diag}(\hat{x}_j)^H \text{diag}(\hat{x}_j) + \lambda I \right)^{-1} \sum_{j=1}^m \text{diag}(\hat{x}_j)^H \hat{y}_j. \quad (7)$$

Note that the matrix to be inverted in (7) is diagonal, which can be computed by a simple element-wise division. We can further simplify the expression in terms of its Fourier coefficients of the filter by multiplying with \mathcal{F}^H from the left, leading to a particularly simple and efficient formula

$$\hat{w}_m = \frac{\sum_{j=1}^m \hat{x}_j \hat{y}_j}{\sum_{j=1}^m \hat{x}_j \hat{x}_j + \lambda}. \quad (8)$$

The resulting formula contains only element-wise operations and DFTs. It is hence computed in $\mathcal{O}(N \log N)$ time by exploiting the FFT algorithm, where $N = N_1 N_2$ denotes total number of pixels within a sample.

3.1.2 Standard Multi-channel DCF Formulation

In DCFs, multi-channel features, such as the RGB image representation of sample x_j , offer a much richer representation compared

to a single-channel feature (e.g. grayscale intensity). Here, we describe the multi-channel DCFs formulation. In the literature, there are two general strategies to integrate multi-channel features, here termed early and late fusion.

Early Fusion: The first common scheme is known as early fusion in which the multiple channels are directly concatenated and then a multi-channel DCF classifier is trained. To generalize (8) to multi-channel DCFs, we let x be composed of several feature channels $x^d : \Omega \rightarrow \mathbb{R}$, where $d \in \{1, \dots, D\}$ is the channel index. In other words, at each spatial location $(n_1, n_2) \in \Omega := \{0, \dots, N_1 - 1\} \times \{0, \dots, N_2 - 1\}$ we have a D -dimensional feature vector $x_j(n_1, n_2) = x_j(n) \in \mathbb{R}^D$. Similarly, we extend the filter w to be D -dimensional, denoting the individual filter channels with w^d . The convolution between two multi-dimensional functions is defined by simply summing the convolution responses for each individual channel. We can learn the filter w by minimizing the generalized objective,

$$L(w) = \sum_{j=1}^m \alpha_j \left\| \sum_{d=1}^D x_j^d * w^d - y_j \right\|^2 + \lambda \sum_{d=1}^D \|w^d\|^2. \quad (9)$$

Here, we have additionally introduced the per-sample importance weights $\alpha_j \geq 0$.

The above multi-channel formulation of DCF is still a linear least squares problem. It cannot be fully diagonalized by the DFT in the general case. There are however two exceptions, where diagonal solutions are found. The first case is the original MOSSE model in (8) and the other case occurs when there is only one single training sample $m = 1$, resulting in the following optimal filter,

$$\hat{w}^d = \frac{\hat{x}^d \hat{y}}{\sum_{d=1}^D \hat{x}^d \hat{x}^d + \lambda}. \quad (10)$$

By again exploiting the FFT, the solution to the general multi-channel objective (9) can be obtained by solving N number of D -dimensional linear systems. This has a computational complexity of $O(DN \log N + ND^3)$, which is substantially smaller than the direct solution in the spatial domain, which requires $O(D^3 N^3)$ operations [55].

Late Fusion: Here, the aim is to compute the DCF on each single channel of the sample x_k and then aggregate all the filters to obtain the resulting classifier. Let \hat{w}^d be the estimated filter of the d -th feature layer using (8), then $w = \sum_{d=1}^D \beta_d w^d$ is aggregated filter obtained by summing up each DCF, where β is the scaling factor used to give relative importance to each of the feature layer.

3.1.3 Standard DCF Tracking Pipeline

For tracking, the DCF first learns the filter w online and then it performs tracking-by-detection. Once the target is tracked in the current frame, the model is then learnt recursively. The block diagram of the DCF tracking pipeline is shown in Fig. 2.

Target Detection: Let m be the number of the current image frame, in which we strive to localize the target. From the previous frame we are given the filter \hat{w}_{m-1} that has been recursively updated since the initial frame. We extract an image patch z centered at the predicted target location. Here, z has the same size $N_1 \times N_2$ as the training patches x_j . Using the convolution theorem, we then predict target scores $s(n)$ at each location $n \in \Omega$ in z by applying the learned filter,

$$s = \sum_{d=1}^D z^d * w_{m-1}^d = \mathcal{F}^{-1} \left\{ \sum_{d=1}^D \hat{z}^d \hat{w}_{m-1}^d \right\}. \quad (11)$$

This computes the score function at all locations $n \in \Omega$. We can then estimate the target location in frame m as the maximizer of

the target score function $n^* = \operatorname{argmax}_n s(n)$. In case we are confident about our estimate, we can construct a new training sample (x_m, y_m) by extracting a patch x_m centered at the estimated location n^* .

Model update: Another feature of (8) is that the filter can be easily updated with new training samples. Let us regard the sample index j as the frame number. That is, sample x_j was extracted from frame number j . In online learning, it is common to employ a learning rate parameter $\gamma \in [0, 1]$ which controls the speed at which the model is adapted to the new data. (8) suggests updating the numerator $\{w_m^{num}\}$ and denominator $\{w_m^{den}\}$ of the filter $\hat{w}_m = \frac{w_m^{num}}{w_m^{den}}$ recursively as,

$$\begin{aligned} \hat{w}_m^{num} &= (1 - \gamma) \hat{w}_{m-1}^{num} + \gamma \hat{x}_j \hat{y}_j, \\ \hat{w}_m^{den} &= (1 - \gamma) \hat{w}_{m-1}^{den} + \gamma (\hat{x}_j \hat{x}_j + \lambda). \end{aligned} \quad (12)$$

To achieve the solution given by the recursive formula (12), we need to choose the weights $\alpha_j = \gamma(1 - \gamma)^{m-j}$ for $j > 1$ and $\alpha_1 = (1 - \gamma)^{m-1}$ in (9). Fig. 4 shows influential DCFs-based trackers in literature.

3.2 Siamese Tracking

Deep learning models have revolutionized many machine learning applications. The key to the success of these models is the offline learning capabilities of features on a large volume of data. Such offline training models have a capability to learn complex and rich relationships from large amount of annotated data.

End-to-end offline training models have also been employed in the generic object tracking by posing it as a similarity learning problem [4], [52], [127]. Deep SNs have been widely used to learn a similarity between the target image and the search image region [4]. SNs were first used for signature verification task [12] and then adapted for other applications including fingerprint recognition [19], [125], stereo matching [159], ground-to-aerial image matching [83], and local patch descriptor learning [49].

In VOT, an offline deep network is trained on a large amount of pairs of target images to learn a matching function during training and then this network as a function is evaluated online during tracking. Bertinetto *et al.* unveiled the power of SNs and proposed Fully Convolutional SN (SiamFC) for VOT [4]. SiamFC used the backbone feature extractor and compared the similarities of the different positions in the search image to determine the position of the target object in each frame. It also utilized the fully-convolutional property that makes searching feasible and is adopted in most of the subsequent SNs-based trackers. SiamFC consists of two branches, the template branch and the detection branch. The template branch receives the target image patch in the previous frame as input while the detection branch receives the target image patch in the current frame as input. Both of these branches share CNN parameters so that the two image patches encode the same transformation which is suitable for tracking. The tracking pipeline of SiamFC is shown in Fig. 3 which is very similar to the generic SNs.

The main aim of SN is to overcome the limitations of pre-trained deep CNNs and take full advantage of end-to-end learning for real-time applications. The offline training videos are used to instruct the Siamese tracker to handle various tracking challenges including rotations, changes in viewpoint, and lighting variations etc., and locate the target object in the consecutive frames.

3.2.1 Training Pipeline

In SiamFC, we consider a pairs of training images (x, z) . The input x of size $127 \times 127 \times 3$ is the reference image which is derived from the ground truth bounding box of the object in the




Fig. 3: The Siamese tracking pipeline for generic object tracking.

first frame. The input z of size $255 \times 255 \times 3$ is the larger search region of each frame. We input these pairs (x, z) into a CNN to obtain two feature maps (e.g., using the final layer of the backbone architecture). By going through the same sub-networks of the same weights, the reference image is encoded as a feature vector of size $6 \times 6 \times 128$ and the search region is embedded as a feature vector of size $22 \times 22 \times 128$. Both feature maps are then matched using a cross-correlation,

$$g_\rho(x, z) = f_\rho(x) \star f_\rho(z) + b, \quad (13)$$

where \star denotes cross-correlation operation, $f_\rho(\cdot)$ is a deep CNN, e.g. AlexNet [1], with learnable parameters ρ , and $g_\rho(x, z)$ is a response map denoting the similarity between x and z , and $b \in \mathbb{R}$ denotes a scalar offset value. The cross-correlation similarity function (13) performs an exhaustive search of the features extracted from sample x over the features extracted from z .

The goal is for the maximum value of the response map $g_\rho(x, z)$ to correspond to the target location. To achieve it, the network is trained offline on millions of random pairs extracted from a collection of videos to track the generic objects. The mean of the logistic loss is typically employed to train the network as,

$$\ell(c, v) = \frac{1}{N} \sum_{i=1}^N \log(1 + \exp(-cv)), \quad (14)$$

where v represents real-valued score of a single target-test sample and $c \in \{-1, 1\}$, is its ground-truth label. Training proceeds by minimizing an element-wise logistic loss ℓ over the training set as:

$$\operatorname{argmin}_\rho \mathbb{E}_{(x, z, c)} \ell(c, g_\rho(x, z; \rho)). \quad (15)$$

3.2.2 Testing Pipeline

To prove the effectiveness of SiamFC and its generalization capability when trained on large-scale dataset, a simplistic tracking algorithm is employed by first extracting the feature representation of the x and z in the new frame. The feature representation of x is then compared to that of z , which is obtained in each new frame by extracting a window centred at the previously estimated position, with an area that is four times the size of the object. Cross-correlation is then performed on the two feature maps to get a score of size $17 \times 17 \times 1$. A cosine window is also employed to the correlation map to penalize large displacements. The position of the highest score in the score map is reverted back to its corresponding position, which is chosen as the predicted bounding box center of that frame. The original SiamFC tracker achieves surprisingly good results at a real-time speed of 140FPS on a GPU. However, it does not update a model and hence is not capable of handling large appearance variations. With the same spirit, Tao *et al.* proposed SINT in which Euclidean distance is employed as a similarity measure instead of cross-correlation [127]. Held *et al.* proposed the GOTURN in which a bounding box regression is employed [52]. Similarly, Valmadre *et al.* proposed CFNET in which correlation filter is added in the x as a separate block in the matching function (13) and makes this network shallower but more efficient [128].

With these tracking advantages, SiamFC also lacks the target scale estimation component. To address this issue, Region Proposal Network (RPN) [108] is used to predict the target scale within the traditional SiamFC tracker. RPN takes an input image and estimates a set of rectangular object proposals, each with an objectness score. To do so, a small network is slid over the convolutional feature map output by the last convolutional layer. Li *et al.* proposed the SiamRPN tracker [73] which incorporates an RPN component. The outputs of SiamRPN include one classification and one regression branches to regress the target bounding box for both position and scale estimation. SiamRPN showed improved accuracy compared to classical SiamFC tracker. Therefore, SNs are computationally efficient in both inference and in offline learning. SNs have demonstrated SOTA tracking performance and therefore receiving a lot of attention in the tracking community. Fig. 4 shows influential Siamese trackers in literature.

4 COMMON OPEN ISSUES OF DCFs AND SIAMESE TRACKING PARADIGMS

In this section, we discuss the common challenges of both paradigms including backbone architectures, target state estimation, the evolution of these trackers to segmentation-based trackers, and the integration these trackers in multiple object tracking pipeline. In the below sub-sections, we discuss these challenges in more details.

4.1 Backbone Architectures

In offline training, the backbone feature extraction network plays a dominant role to capture low-level fine-grained and high-level semantic information of the target. Here, the backbone networks, such as AlexNet [1], VGG-16 or VGG-19 [116] and ResNet18 or ResNet50 [51], are used to optimize deep features on tracking datasets [18], [127], [128], [151], [162]. Instead of relying on pre-trained networks, the task-specific deep feature learning promises improved representations for the tracking problem itself. Both tracking paradigms have demonstrated encouraging performance using powerful backbone networks.

For instance, Valmadre *et al.* proposed a CFNET that follows end-to-end learning of correlation filters in an offline manner [128]. CFNET employed a variant of the AlexNet model and conducted the ablation study by utilizing all the convolutional layers. A second convolutional layer demonstrated promising tracking results. Other trackers such as, CREST [118] and ACFN [18] also followed the same strategy in an online manner. CREST utilized VGG-16 model and extracted the feature maps from the conv4-3 layer. The feature channels were reduced to 64 using PCA dimensionality reduction. In these trackers, the objective is to improve the target regression. These methods demonstrated comparable performance, as compared to trackers that utilized deep features extracted from the pre-trained networks. Recently, an end-to-end target scale estimation component is incorporated in ATOM [24] while the discriminative strength of the classical DCF model is improved in DiMP [5] and PrDiMP [26]. ATOM utilized ResNet-18 pre-trained model as a backbone network. For target classification, it employed block 4 features, while the target estimation component utilized both block 3 and 4 features. DiMP and PrDiMP employed ResNet-18 and ResNet-50 backbone architectures. They utilized features extracted from the third block for the model prediction. *This recent trend of end-to-end features learning in DCF trackers [5], [24], [26] has resulted in excellent tracking performance on multiple benchmarks, paving the way to explore more sophisticated end-to-end feature learning in DCF.*

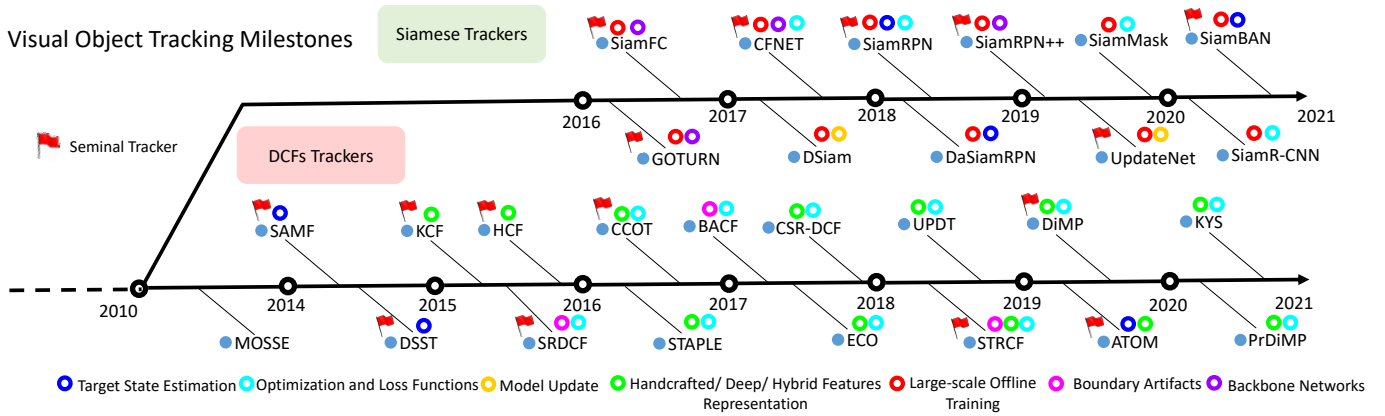


Fig. 4: Influential Siamese and DCF trackers proposed in the 2014-2020 period.

In early Siamese trackers (e.g., SiamFC [4], GOTURN [52], SINT [127], SiamRPN [73]), a modified pre-trained AlexNet is fine-tuned [1]. SiamFC extracted the convolutional features using the five layers and utilized last layer features for offline training. GOTURN extracted features using all layers and the features from the last fully connected layer were utilized to train the objective function. SINT showed significant performance improvement using a set of different features extracted from the backbone. SiamRPN fixed the first three convolutional layers and only fine-tuned the last two convolutional layers. A variety of trackers (FlowTrack [175], MemTrack [151], and EAST [58]) have also used AlexNet. However, it is observed that these trackers are still limited in performance because AlexNet is a relatively shallow network and does not produce very strong feature representations.

With the advent of wider and modern deeper networks, the tracking community has also designed SNs based on ResNet [51], VGG-19 [116], and Inception [124]. All these networks are pre-trained on the ImageNet dataset and fine-tuned to learn the generic matching function. Tao *et al.* also analysed VGG-16 backbone model in SINT [127]. Results show performance gap in tracking between these two pre-trained networks. The concatenated features of VGG-16 model demonstrated promising results. The deeper layer features encode rich semantic information useful for classification of target from the background however, it is also observed that direct replacement of powerful deep architectures does not bring performance improvement in the classical SNs-based tracking methods [72].

To address this issue, Li *et al.* investigated the main reason and proposed a ResNet-driven SiamRPN++ tracker [72]. In SNs, when modified AlexNet is employed without zero-padding [4], the learnt spatial feature representations of the target does not satisfy the spatial translational invariance constraint. Therefore, an effective sampling technique is developed to fulfill this spatial invariance constraint. The ResNet-50 architecture is adopted in SiamRPN++ with some refinements to tailor it for the tracking task. SiamRPN++ utilized the outputs of conv3, conv4 and conv5 blocks and fed them into three SiamRPN modules. For multi-layer features fusion, weighted sum and depth-wise cross-correlation modules were also introduced. A layer-wise aggregation is proposed which fuses the rich features representation of different convolutional blocks. The outputs of three RPN modules are of the same spatial size. The weighted sum of the three outputs is used for generating the final feature maps. The weights are end-to-end optimized offline together with the network. Due to the deep architecture, SiamRPN++ may have many more parameters which require more computation resources and decrease the tracking

speed, compared to its predecessors therefore, a tracker is also equipped with the depth-wise cross-correlation layer for efficiency. The depth-wise cross-correlation is performed between the search region and the template regions to obtain a multi-channel response map. The response is then convolved with a 1×1 kernel to reduce its dimension to fewer channels. Using this technique, the number of parameters are significantly reduced through the dimension reduction and stabilized the training procedure. The final dimension reduced response map is then adopted as the input to the both classification and regression branches.

Leveraging the powerful deep ResNet architecture, the performance of many Siamese trackers are improved. Zhang *et al.* also investigated the same problem and proposed SiamDW in which a shallow backbone AlexNet is replaced with deep networks including Inception, VGG-19, and ResNet [167]. It is investigated that apart from features padding, receptive fields of neurons and network strides are also main reasons why such a deeper network cannot directly replace shallow networks. Also, features from different layers and across various architectures were evaluated and the best performing candidate feature and architecture was used for the tracker. The results presented in both studies demonstrated an excellent performance compared to classical SNs-based trackers. With these foundations, recent trackers including SiamCAR [45], Ocean [168], and SiamBAN [15] etc., have also employed powerful deep architectures. These recent trackers extracted the features from the last three residual blocks of a ResNet-50 and fused to get the multi-channel response map. *The ResNet backbone has become the established and preferred alternative for Siamese tracking, thanks to its simplicity and strong performance. Also, depth-wise cross-correlation has attracted significant attention for obtaining multi-channel features map. However, recent advances in vision transformer networks [35], [88] are expected to have a substantial impact in the tracking community in the coming years.*

4.2 Target State Estimation

Both tracking paradigms have demonstrated promising results in terms of accuracy and robustness. However, when a target moves, its template size (also known as bounding box size) also changes. Both paradigms suffer from severe scale variations challenges. Therefore, accurate scale estimation poses a great challenge in these trackers. The problem of handling bounding box size for accurate target scale estimation is an established research direction. The tracking community has shown outstanding progress in this direction and proposed potential solutions to handle it. Here, we discuss the scale estimation methods proposed for both paradigms. **Multiple Resolution Scale Search Method:** One straight-forward strategy in DCFs is to apply learned translation filter w at different

image scales. That is, the image is first resized by different scale factors, followed by feature extraction. The feature map at each scale is then convolved with learned filter w to compute the target scores. The change in target location and scale can then be estimated by finding maximum score across all scales. This is a common strategy, often applied in both tracking and detection [39]. Li *et al.* proposed SAMF tracker in which the translation and scale filters were jointly trained using the standard DCF formulation [79]. The results demonstrated significant performance gain compared to the standard DCF. This scale adaptive component has been utilized in a number of DCF-based trackers such as, CACF [100], CFAT [9], and FD-KCF [169]. However, this approach suffers from a higher computational cost, since the translation filter has to be applied at several resolutions to achieve scale accuracy.

The Siamese trackers have also benefited from this scale estimation methods. In classical SiamFC, multiple scales are searched in a single forward-pass by assembling a mini-batch of scaled images and then the maximum response is computed. Early trackers including RASNET [137], SA-Siam [50], StructSiam [166], UDT [133], UDT++ [134], TADT [78], GradNet [75], RTINET [154], and FlowTrack [175] have employed this method. **Discriminative Scale Space Search Method:** Danelljan *et al.* proposed an alternative strategy for accurate scale estimation in DCFs tracking [27]. Unlike [79], the target estimation is performed in two steps in order to avoid an exhaustive search across translation and scales. Since the scale changes between two frames is usually small or moderate, the target translation is first found by applying the normal translation filter w at the current scale estimate. Then, a separate one-dimensional filter is applied in the scale dimension to update the target size. The scale filter is trained analogously to the translation filter, but operates in the scale dimension by extracting samples of the target appearance from a set of different scales.

The advantage of the aforementioned scale-filter approach [27] is two-fold. First, computational efficiency is gained by reducing the search space. Second, the scale filter is trained to discriminate between the appearance of the target at different scales, which can lead to a more accurate estimation. The proposed scale filter component has been utilized in a multitude of trackers, including STAPLE [3], MUSTer [56], ASRCF [21], CACF [100], BACF [42], CSR-DCF [92], MCCT [135], and LCT [96]. Moreover, the follow-up fDSST tracker reduces the computational cost of DSST by applying the PCA and sub-grid interpolation [28]. Siamese trackers could not explore this technique for scale estimation.

Bounding Box Regression Method: The aforementioned methods show improved performance. However, they depend on the scaling factor parameters and online accurate filter response. These methods do not exploit the powerful deep feature representations in an offline manner. As a result, these methods show performance degradation in the presence of sudden scale variations. Accurate target scale estimation is a complex task, requiring high-level a priori knowledge. The bounding box depends on the pose and viewpoint of the target, which cannot be modeled as a simple image transformation (e.g. uniform image scaling). It is therefore extremely difficult to learn accurate target estimation online.

In object detection methods, bounding box regression has been widely employed for precise localization of object bounding boxes [43], [59], [62], [106]–[108]. Conventional L_1 or L_2 loss is used between the predicted and ground-truth parameters of the bounding box for regression. To exploit the strength of end-to-end deep features learning for target scale estimation, this component has

recently been utilized in DCFs-based trackers [5], [24]. In ATOM [24], inspired by the IoU-Net [62], the targets-specific features are trained. Since the original IoU-Net is class-specific, and hence not suitable for generic tracking, a novel architecture is proposed for integrating target-specific information into the IoU prediction [24]. This is achieved by introducing a modulation-based network component that incorporates the target appearance in the reference image to obtain target-specific IoU estimates. This further enables the target estimation component to be trained offline on large-scale datasets. During tracking, the target bounding box is found by simply maximizing the predicted IoU overlap in each frame. Results demonstrated an excellent performance boost as compared to classical multi-scale search methods. Several recent DCFs trackers, including DiMP [5], PrDiMP [26], and KYS [6] have also utilized this strategy for state estimation. In PrDiMP, instead of predicting the IoU, it employs energy-based model that predicts the un-normalized probability density of the bounding box. This is trained by minimizing the KL-divergence to a Gaussian model of the label noise. Siamese trackers have also explored bounding box regression method for scale estimation. Two different types of methods are proposed in the literature including anchor-based and anchor-free bounding box regression.

Anchor-based Bounding Box Regression Method: To estimate the target bounding boxes more accurately, the object detection capabilities are also introduced within the SN [4]. Anchor-based bounding box regression using a Region Proposal Network (RPN) [108] efficiently predicts proposals with a wide variety of scales and aspect ratios. RPN produces the multiple region proposals which are simultaneously predicted at each sliding-window location. RPN is a fully convolutional network that simultaneously predicts object bounds and objectness scores at each position. The RPN is trained in offline manner to generate high-quality proposals [108].

RPN driven Siamese tracker is proposed by Li *et al.* [73] which outputs of a classification (\mathcal{L}_{cls}) and a regression (\mathcal{L}_{reg}) components for estimating both position and scale. Different from SiamFC which employs multi-scale search for scale estimation, SiamRPN first initializes top K proposals (i.e. the so called anchors) of different scales and aspect ratios at each convolutional features map location and applies some strategies to determine the best proposal among them. It discards the bounding boxes generated by the anchors too far away from the center and re-ranks the proposals by applying a cosine window and scale change penalty to scores. RPN shares features and thus enables highly efficient region proposal computation. This method advanced the Siamese tracking paradigm compared to classical multi-scale space search method. Many recent trackers such as DaSiamRPN [174], SiamRPN++ [72], SiamDW [167], SPLT [150], C-RPN [38], SiamAttn [158], CSA [148], and SPM [131], etc., are also built upon the same notion.

Anchor-free Bounding Box Regression Method: Anchor-based bounding box regression assists Siamese trackers [72], [73] to handle changes in scale and aspect ratio and shows encouraging results. However, it needs to carefully design anchor boxes based on heuristic knowledge, which introduces many hyper-parameters and computational complexity. Therefore, it is essential to carefully design and fix the parameters of the anchor boxes.

Anchor-free regression has also been proposed in object detection, which avoids hyper-parameters associated with the anchor boxes and is more flexible and general [69], [172]. Anchor-free detectors directly find objects without preset anchors using

keypoint-based methods [69] and center-based methods [172]. The keypoint-based methods first locate the pre-defined keypoints and then perform bounding box regression on the objects. In center-based methods, the four distances from positives to the object boundary are predicted using the center on the region of object. The anchor-free detectors are able to eliminate those hyperparameters related to anchors and have achieved similar performance with anchor-based detectors, making them more powerful in terms of generalization ability.

Chen *et al.* proposed the SiamBAN tracker [15] in which anchor-free box regression is employed to estimate the target scale. The tracker avoids hyper-parameters associated with target bounding boxes without any preset anchor boxes. The tracker exploits the expressive power of the fully convolutional network to classify objects and regresses their bounding boxes in a unified manner. Similar to SiamRPN [72], SiamBAN includes a classification module which performs foreground-background classification on each point of the correlation layer and a regression module performs bounding box prediction on the corresponding position [15]. While anchor-based bounding box regression can handle variations in scale and aspect ratio in the Siamese trackers, it mainly suffers with two drawbacks. Firstly, it typically needs a very large set of anchors to cover almost all objects in the image, and since the object in the image may only occupy a small area, making a huge gap between positive and negative samples, which can greatly reduce the performance of the prediction model. Secondly, the use of anchors introduces many hyper-parameters and choices, including the numbers, the sizes and the aspect ratios of anchors. To some degree, the better tracking results heavily depend on the preset anchors. On the other hand, anchor-free bounding box regression-based trackers need not choose boxes with prior knowledge. Its biggest advantage is to directly predicts the class and location information of the object according to the output of network, without cumbersome process of setting prior boxes. Due to its simplicity and convenience, anchor-free trackers further improve the tracking performance compared to anchor-based RPN scale estimation method. Ocean [168] and SiamCAR [45] trackers have also utilized the same method for scale estimation. *The object detection capabilities have demonstrated outstanding progress in the target state estimation component for both paradigms. The recent trends of using RPNs and anchor-free bounding box regression have revealed to further explore these techniques in an end-to-end paradigm. Deep DCFs-based trackers could also take advantage of these anchor-based and anchor-free regression methods for robust tracking.*

4.3 Offline Training

As discussed above, training data is crucial for VOT therefore it is very challenging to learn a robust discriminative model in DCFs (4) and a matching function in SNs (13).

To handle this issue, the tracking community has made outstanding progress by leveraging external images and videos dataset to learn the appearance model. Object detection, image classification, and object segmentation datasets including ImageNet ILSVRC2014 [110], ILSVRC2015 [111], COCO [84], YouTube-BB [105], and YouTube-VOS [142], have been widely used by both family of trackers. These datasets sufficiently cover a good amount of semantics and do not focus on particular objects, otherwise, the tuned network parameters will over-fit to particular object categories in offline training. The datasets are typically annotated with bounding boxes of a target object in every frame.

The tracking community has also compiled more diverse small and large-scale tracking benchmark datasets containing tens of thousands of annotated bounding boxes. These include the LaSOT [36], GOT-10K [60], and TrackingNet [101] which are also used for offline training by recent trackers [5], [15], [24], [26], [130], [148]. Interested readers can find more details about different training datasets utilized by different Siamese trackers in Table I of the supplementary material. It is evident from the results that more offline training data improves the tracking performance.

End-to-end DCFs trackers fully exploit large-scale training datasets to learn a robust discriminative model. For instance, CFNET [128], ATOM [24], DiMP [5], and PrDiMP [26] have employed TrackingNet, LaSOT, GOT10K, ILSVRC2014, and COCO datasets with some augmentation techniques. DiMP inputs a set of multiple training samples from the sequence and is trained on pairs of training and testing sets. Each set consists of images paired with their bounding boxes. The target model is then predicted using training samples and then evaluated on the testing frames. The same strategy is adopted in ATOM and PrDiMP trackers.

In Siamese tracking, the network is trained with image pairs, using one image to predict the target template and the other for evaluating the tracker. Unlike DCF paradigm, the standard Siamese formulation cannot exploit the appearance of known distractor objects during tracking. Siamese approaches therefore often struggle when objects similar to the target itself present. This occurs when, for instance, other objects of the same semantic class are in the view. Early Siamese trackers (SiamFC [4] and SiamRPN [73] particularly) only sample pairs of training images from the same video during training. This sampling strategy does not focus on challenging cases, with semantically similar distractor objects. To address this issue, hard negative mining techniques have been developed in the literature. For instance, Zhu *et al.* introduced hard negative mining technique in DaSiamRPN [174] to overcome data imbalance issue by including more semantic negative pairs into the training process. The constructed negative pairs consist of labelled targets both in the same and different categories. This technique assisted DaSiamRPN to overcome drifting by focusing more on fine grained representations. With the same spirit, Voigtlaender *et al.* proposed another hard negative mining technique using an embedding network and a nearest neighbor approximation [130]. For every ground truth target bounding box, the embedding vector is extracted for a similar target appearance using pre-trained network. The indexing structure is then employed to estimate the approximate nearest neighbors and use them to estimate the nearest neighbors of the target object in the embedding space. *This recent trend of utilizing more training data and designing data mining techniques have demonstrated excellent tracking performance on multiple benchmarks, opening many doors to explore more sophisticated techniques in offline training.*

4.4 Evolution of both Paradigms to Segmentation-based Trackers

Accurate object segmentation provides reliable object observations for tracking. The combination of segmentation and tracking may solve several tracking problems including rotated bounding box, occlusion, deformation, and scaling, etc., and fundamentally avoid tracking failures. Therefore, the segmentation branch can act as the complementary component of the trackers. In literature, segmentation-based approaches have been incorporated within both DCFs and Siamese-based trackers for improved filter learning in the presence of non-rectangular targets [3], [63], [92], [138].

For example, in DCFs tracking, Bertinetto *et al.*, used a color histogram-based segmentation method to improve tracking under varying illumination changes, motion blur, and target deformation [3]. Lukezic *et al.*, proposed a spatial reliability map using a color-based segmentation method to regularize filter learning [92]. A real-time tracker is proposed using handcrafted features and achieved comparable performance using deep features. Kart *et al.* [63], extended the CSR-DCF tracker for RGB-depth tracking based on both color and depth segmentation as depth cues provide more reliable segmentation map. Lukezic *et al.*, proposed a single shot segmentation tracker to address VOT and video object segmentation problems within a joint framework [91]. The target is encoded with two discriminative models for the joint tracking and segmentation task. The results are reported in many tracking and segmentation benchmarks and demonstrated the benefits. Recently, Robinson *et al.*, employed a powerful discriminative model using the fast optimization scheme borrowed from the ATOM [24] for video object segmentation task [109]. Bhat *et al.*, also used the target model discriminative capabilities for more robust video object segmentation [8].

Recently, SNs have also been extended to perform both video object segmentation and tracking [138]. Siamese trackers are quite fast and provide real-time performance while video segmentation approaches are slow and not real-times therefore, combining these two problems provide efficient solution for both tracking and segmentation. Wang *et al.*, present a SN to simultaneously estimate binary mask, bounding box, and the corresponding background-foreground scores [138]. This multi-stage deep network lacks the opportunity to process the visual tracking and target segmentation jointly to increase robustness. Lu *et al.*, employed unsupervised video object segmentation task in which a novel architecture is proposed based on the co-attention mechanism within the SN [90].

4.5 Integration of both Paradigms in Multiple Object Tracking Pipeline

Multiple Object tracking (MOT) is a task of estimating the trajectories of the multiple targets in a video sequence. MOT is challenging because a successful method needs to not only detect the objects of interest accurately in each frame, but also associate them throughout the video [70], [115], [144], [173]. While, having promising directions of DCFs and Siamese trackers towards single object tracking, both paradigms have also been integrated into MOT pipeline to localize multiple targets in each frame [144], [156], [170], [171].

For instance, Zhu *et al.*, integrated a ECO tracker in a unified framework to handle the noisy detections for robust targets association [173]. Chu *et al.* integrated the discriminative instance-aware KCF tracker into MOT framework [20]. Recently, Zhou *et al.*, integrated segmentation-based discriminative tracker for multi-object segmentation [171]. Zheng *et al.*, proposed center net loss and integrated variant of the DiMP tracker [170]. The tracker branch trains a separate single object tracking model per target online to distinguish the target from its surrounding targets assigning single object tracking model the novel discrimination.

Taixe *et al.* proposed one of the first SNs for target association in MOT [70]. Yin *et al.*, integrated classical SiamFC tracker and provided a unified model that estimates the object motion and affinity network [156]. Recent studies have also successfully integrated the SiamRPN and GOTURN trackers to improve the MOT performance in the presence of occlusions and similar distractors [115], [144]. The aforementioned MOT methods have

improved the robustness as well as relieved the dependency on external detector in the presence of single object tracker. *This recent trend of utilizing single object tracking paradigms as an integral components of the MOT models have demonstrated excellent performance on multiple benchmarks, opening new directions to further explore the inherent properties of DCFs and SNs.*

5 DISTINCT OPEN ISSUES OF DCFs AND SIAMESE TRACKERS

While having several common challenges, both paradigms also suffer from distinct issues which are rather specific to them. In the below sub-sections, we discuss distinct issues of both paradigms in more details.

5.1 Distinct Issues in the DCF Tracking Pipeline

While having important promising properties, the standard DCF framework presents several distinct challenges including feature representations, boundary artifacts, and optimization when applied for the task of generic object tracking. In the subsections below, we identify and discuss these important challenges developing a DCF based tracking pipeline.

5.1.1 Feature Representations

In object tracking, a variety of visual features have been investigated in the literature. Finding discriminative, yet invariant features is particularly important when applying linear discriminative models, such as DCF, which are restricted to finding a linear decision boundary. Handcrafted features [22], [39], [129], deep features [13], [51], [116], hybrid features, and end-to-end learning of features [24] have been explored within the DCF-based tracking framework. Next, we describe details about different types of features used in the DCF-based trackers.

Handcrafted Features: Early DCF trackers such as, MOSSE [10] and CSK [54] have exploited intensity features for object tracking. Other than intensity features, local color and intensity histograms features are also utilized in DCF trackers such as, RPAC [86], LCT+ [95], LCT [96], and CACF [100]. Simple color representations, including RGB and LAB have been used for DCFs-based trackers such as, STAPLE [3] (RGB), SCT [17] (RGB + LAB), and ACFN [18] (RGB + LAB).

To achieve a more discriminative image representation, ACA [31] investigated different color descriptors and proposed to use the Color Names (CN) features [129] along with the intensity channel. The ACA tracker further introduced an adaptive dimensionality reduction technique to compress the CN features, thereby providing a tradeoff between speed and tracking performance. CN features have also been employed in several subsequent DCFs-based trackers such as, MCCT [135], MKCF [126], MUSTer [56], CSR-DCF [92], CCOT [23], ECO [25], UPDT [7], AutoTrack [81], ARCF [61], GFS-DCF [143], RPCF [123], and DRT [120].

Another popular handcrafted feature employed in DCF-based trackers is Histogram of Oriented Gradients (HOG) [22]. HOG captures shape information by collecting statistics of the image gradients. HOGs are formed in a dense image grid of cells. Within the DCF paradigm, KCF [53] was the first tracker to utilize HOG features for tracking. Several DCF trackers such as, MCCF [64], CFLB [65], BACF [42], SRDCF [30], STRCF [74], RPCF [123], GFS-DCF [143], RPT [80], RCF [119], LMCF [132], PTAV [37], StruckCF [85], CFAT [9], and LSART [121] have utilized HOG features. *These features have been the preferred alternative among handcrafted methods due to its speed and effectiveness. Further, HOG features have also been effectively combined with CN features to utilize both shape and color information.*

Deep Features: In recent years, deep learning has revolutionized many areas in computer vision. Deep Convolutional Neural Networks (CNNs) have shown to be particularly well suited for image related tasks [44]. They apply a sequence of learnable convolutions and non-linear operations onto the image. However, employing deep features for object tracking has proved challenging. This was primarily due to the scarcity of training data for tracking in the initial years of deep learning, as well as the high dimensionality of the features. Therefore, many DCF-based trackers such as, HCF [93], HDT [104], CCOT [23], ECO [25], ASRCF [21], and RPCF [123], employ deep CNNs that are pre-trained on the ImageNet dataset [32] for image classification. Despite being trained for classification, such deep representations are applicable to a wide range of vision tasks [114].

Within the DCF framework, the deep features are extracted from the convolutional layers. Some popular pre-trained deep networks, including VGG-19 [116], imagenet-vgg-m-2048 [13], VGG-16 [116], ResNet50 [51], and GoogleNet [124] are used to extract deep feature representations. For instance, Ma *et al.*, incorporated the hierarchical deep convolutional features for visual tracking [93]. The correlation response maps were computed from each convolutional layer. The smaller weights were then assigned to earlier layers and higher weights were given to later layers for more accurate tracking. The HDT tracker also employed deep features from six convolutional layers of the same network [104]. The DeepSRDCF tracker used the imagenet-vgg-m-2048 network and performed an analysis of the convolutional feature maps for tracking [29], indicating the importance of the shallow layer. Shallow layers contain low-level information at a high spatial resolution, important for accurate target localization.

On the other hand, deeper-layer feature maps possess high-level invariance to complex appearance changes, such as deformations and out-of-plane rotations. Thus, deeper layers have the potential of improving tracking robustness, while largely invariant to small translation and scale changes. Therefore, an accurate strategy of fusing shallow and deeper convolutional layers within the DCF framework has been a topic of interest. In CCOT, a continuous-domain formulation of the DCF framework is proposed that enables the integration of multi-resolution features [23]. The ECO investigates strategies to reduce the computational cost of the CCOT and mitigates the risk of overfitting [25]. Other trackers, such as HDT [104], HCFs [94], MCCT [135], MCPF [164], MCPFs [165], LMCF [132], STRCF [74], TRACA [16], DRT [120], UPDT [7], and GFS-DCF [143] integrate deep features using late fusion strategy. This strategy is to train a classifier on each individual feature representation and then aggregate feature response map.

5.1.2 Boundary Artifacts

The standard convolution is effectively replaced by circular convolution (1) in the DCF formulation to ensure the applicability of the DFT, resulting in the formula (11) for evaluating the target predictions. The circular convolution improves the computational efficiency, but as a disadvantage it introduces unwanted boundary artifacts.

The most severe consequence of the circular convolution assumption arises in the learning formulation (9). The fundamental notion of the DCF paradigm is to train a filter w that can discriminate the target from background image regions. Due to the periodic effects, most of the original background content is replaced by synthetic repetition of a smaller image patch. The model thus sees fewer background samples during training, severely limiting its

discriminative power. Furthermore, due to the distortions caused by the periodic repetitions, the predicted target scores are only accurate near the center of the image patch. The size of the search area is therefore limited. As traditionally performed in signal processing, DCF methods typically pre-process the samples x by multiplying them with a window function [10]. However, this technique does not attempt to solve the aforementioned problems and only serves to smooth out the discontinuities at border regions. Several solutions have been proposed in the literature to overcome the aforementioned boundary artifacts problems. To this end, several approaches are proposed that incorporate target-specific spatial, spatiotemporal, and smoothness constraints within the DCF objective function [21], [30], [42], [65], [74]. Below, we summarize some of these major developments.

Spatial regularization: In the SRDCF, Danelljan *et al.* proposed a spatially regularized framework to control the spatial extent of the filter in order to alleviate the boundary problem [30]. A spatial regularization component is integrated into the multi-channel DCF formulation (9) as,

$$L(w) = \sum_{j=1}^m \alpha_j \left\| \sum_{d=1}^D x_j^d * w^d - y_j \right\|^2 + \lambda \sum_{d=1}^D \|f w^d\|^2, \quad (16)$$

The spatial weight function $f : \Omega \rightarrow \mathbb{R}$ takes positive values $f(n) > 0$, penalizes the filter coefficients $w^d(n)$ based on their spatial location n . By letting f takes large values at background pixels and small values inside the target region, background filter coefficients are penalized. As a result, a compact filter w focusing on the target region can be learned, even for large image sample sizes. The spatial regularization strategy has been employed in a variety of trackers, including ARCF [61], ASRCF [21], and AutoTrack [81]. In order to improve the efficiency of the SRDCF formulation, Li *et al.* proposed the STRCF [74] which only employs a single training sample, and instead introduces a temporal regularization term to integrate historic information. It has also become a popular baseline for numerous works [81].

Constraint optimization: While the SRDCF [30] aims to penalize the filter coefficients outside the target region, Kiani *et al.* propose to introduce hard constraints [42], [65]. This strategy enforces that the filter coefficients $w(n)$ are zero outside the target region. The resulting DCF formulation can be expressed as follows by introducing a binary mask P ,

$$L(w) = \frac{1}{2} \sum_{j=1}^m \|y_j - (P \cdot w) * x_j\|_2^2 + \frac{\lambda}{2} \|w\|_2^2, \quad (17)$$

Here, the element-wise product $P \cdot w$ effectively masks out the influence of filter weights corresponding to background features. The resulting optimization problem can be efficiently solved through iterative techniques such as ADMM [42], [65]. The aforementioned target-specific constraints investigated in (16) and (17) are usually fixed for different objects and they do not vary during the tracking process. Recently, Dai *et al.* extended BACF and SRDCF by introducing an adaptive regularization term [21].

Implicit methods: With the GFS-DCF, Xu *et al.* proposed a joint group feature selection model that simultaneously learns three regularization terms including spatial regularization for feature selection, channel regularization for feature channel selection, and low-rank temporal regularization term to enforce smoothness on the filter weights [143]. Mueller *et al.* proposed to regularize the contextual information of each target patch [100]. In every frame, CACF samples several context patches, which serve as negative samples.

Spatial formulation: Danelljan *et al.* and Bhat *et al.* proposed

ATOM [24] and DiMP [5] trackers. Both these trackers employ deep features in low resolution (stride 16) in order to first coarsely, yet robustly, localize the target object. Due to the coarse resolution and the resulting small target filter size (4×4), Danelljan *et al.* found that the filter can be directly learned in the spatial domain using dedicated and efficient iterative solvers [24]. This approach allows both ATOM and DiMP to completely circumvent the boundary artifacts problem, since no periodic extension of the training samples are performed.

Both the regularization-based (SRDCF [30], STRCF [74]) and constraint-based (CFLB [65]/BACF [42]) formulations have seen large success and been employed in a wide range of trackers. However, more recent deep learning methods (ATOM [24]/DiMP [5]), have completely circumvented the boundary artifact problem by directly optimizing the filter in the spatial domain. Thus, while the Fourier domain is computationally attractive for high-resolution feature maps, efficient spatial-domain optimization methods prevail for online learning when using powerful low-resolution deep features. *Recent works constituting the current SOTA in DCF-based tracking [5], [26] therefore employ a purely spatial formulation, which does not require additional strategies for alleviating boundary effects. By further extending the filter to a multi-channel output, the latter strategy has also demonstrated its use for segmentation [8].*

5.1.3 Optimization

In the standard DCF formulation, inference is performed by computing the DFT coefficients \hat{w} using the least squares solution (8). However, when the model grows more complex and advanced e.g., by introducing multiresolution feature maps and target-specific constraints, such as spatial regularization, and temporal regularizations, the model inference cannot be performed using the simple least-square solution, as provided in (17). Since computational efficiency is a crucial factor in most applications, these modifications require alternative inference methods. Therefore, finding efficient and robust inference schemes is a key problem in DCF-based tracking.

Model inference is performed by minimizing the multi-channel loss (9), which forms the basis of the DCF framework. However, it does not allow for any efficient closed-form solution. Therefore, many DCF trackers such as CACF [100], CSK [54], KCF [53], MUSTer [56], SP-KCF [122], and CFAT [9] employ the diagonalizable cases in both primal or dual domains to derive approximate model inference schemes. These loss functions rely on the very restrictive assumptions of a single feature channel $D = 1$ and a single training sample $m = 1$, respectively. Moreover, these solutions cannot benefit from the aforementioned additional regularizations. Several efficient optimization methods for model inference have been introduced to minimize the loss functions in the literature.

Gauss-Seidel Method [48]: Minimizing the DCF loss functions with its spatially regularized variants online is a highly challenging problem, since the filter w contains tens or hundreds of thousand parameters to be optimized. In [30], an optimization approach based on the iterative Gauss-Seidel method is proposed to minimize the spatially regularized loss function (16). The same strategy is also considered in its variant DeepSRDCF [29] that employs deep features. By employing the Gauss-Seidel-based optimization, the tracker achieves a tracking speed of few frames per second. While not yet real-time, it demonstrated superior robustness and accuracy compared to previous approaches, yet faster than many of its competitors.

Conjugate Gradient Based Method [102]: To pave the way for the use of deep features and to further improve the computational efficiency, a Conjugate Gradient (CG)-based strategy is utilized in CCOT [23]. CG can be applied to any set of normal equations $A\tilde{w} = b$ of full rank. It is based on finding a set of conjugate directions $p^{(i)}$ and optimal step lengths $\beta^{(i)}$ used for updating the filter $\tilde{w}^{(i)} = \tilde{w}^{(i-1)} + \beta^{(i)}p^{(i)}$. Theoretically, the algorithm converges to the solution in a finite number of iterations i . In practice however, the algorithm is stopped after a fixed number of iterations or when the error has decreased to a satisfactory level.

The computational bottleneck of CG is the evaluation of the matrix-vector product $Ap^{(i)}$ in each iteration i . This is in fact the key advantage of CG. The Gauss-Seidel method requires solving a triangular system in each iteration. On the contrary, CG can exploit the particular sparsity structure of A in the computation $Ap^{(i)}$. Moreover, $Ap^{(i)}$ can be implemented as a series of simple block-wise dense matrix-vector products, convolutions and point-wise multiplications. This reduces the quadratic $O(D^2)$ complexity in the feature dimension to linear $O(D)$, a crucial improvement enabling tractable integration of high-dimensional deep features.

To address non-linear least squares problems, the Gauss-Newton optimization is also used in many trackers, including ECO [25], ATOM [24], and UPDT [7]. The method linearizes the error residuals using a Taylor series expansion around the current estimate to find a quadratic approximation of the objective. The resulting quadratic problem can then be tackled with iterative methods, for instance the CG approach described above. In the ECO [25] and ATOM [24], the Gauss Newton combined with CG is employed to jointly optimize the filter w and the dimensionality reduction matrix. DiMP employs Gauss Newton together with Steepest Descent iterations to learn the filter w using a non-linear robust loss function [5]. The optimization steps are themselves differentiable, which further enables end-to-end learning of the underlying deep features. The PrDiMP further employs the more general Newton approximation to address the convex and non-linear KL-divergence objective [26].

Alternating Direction Method of Multipliers (ADMM) Method [11]: When a classical DCF formulation grows in terms of additional regularizations such as those presented in models (16)-(17), the formulation becomes a constraint optimization problem which can be solved using efficient convex or non-convex solvers such as an ADMM-based optimization method. The ADMM method has recently been used in many DCF-based trackers to efficiently solve the DCF loss functions specially when additional regularizations are introduced. The ADMM-based optimization approach provides closed-form solutions for each sub-problem and empirically converges within a very few iterations. In ADMM, the model is solved by breaking it into smaller pieces, each of which is then easier to handle. The augmented Lagrangian formulation is always used to convert the constrained optimization model into unconstrained model with a Lagrangian penalty as an additional variable. Each sub-problem of the unconstrained model is then solved in an iterative manners. Trackers such as BACF [42], DRT [120], AutoTrack [81], ARCF [61], and RPCF [123] have employed ADMM for efficient solutions. Other optimization methods such as Gradient descent is used by ACFN [18], CREST [118], and DSLT+ [89] trackers for network optimization.

Current Choices: Both Conjugate Gradient [7], [23]–[25] and ADMM (STRCF [74], BACF [42], RPCF [123], ASRCF [21], GFS-DCF [143], and MCCT [135]) have been popular choices for DCF trackers. Both support the more complex DCF formulations

required to mitigate the boundary artifacts, using *e.g.* spatial regularization or constraints. The most recent approaches DiMP employed a steepest descent in a spatial-domain formulation [5], while PrDiMP combined it with a Newton approximation to minimize the non-linear KL-divergence objective [26]. In addition to efficiency, the adoption of end-to-end learning have brought the importance of differentiability and simplicity of the optimizer, motivating the latter choice. It has also opened the opportunity to train components and hyper-parameters of the optimizer in a meta-learning fashion, setting the stage for future directions.

5.2 Distinct Issues in the Siamese Tracking Pipeline

The classical SN outperforms DCFs trackers in both accuracy and efficiency. However, SN also suffers from several limitations in terms of a large annotated image pairs in offline training, lack of online adaptability, and loss function formulation. In the subsections below, we identify and discuss these important challenges for developing a robust Siamese-based tracker. We also briefly describe the details of these problems in Siamese tracking and their potential solutions developed in recent years.

5.2.1 Online Model Update

In SiamFC [4], the target template is initialized in the first frame and then kept fixed during the remainder of the video. The tracker does not perform any model update and therefore, the performance totally relies on the general matching ability of the SN. However, appearance changes in the presence of tracking challenges are often large and failing to update the model leads to failure of the tracker. In such scenarios, it is important to adapt the model to the current target appearance. In the literature, the tracking community has also proposed potential solutions in this direction.

Moving Average Update Method: Many recent SOTA trackers including GOTURN [52], SINT [127] and SiamAttn [158] etc., employ a simple linear update strategy using a running average with a fixed learning rate. While it provides a simple means of integrating new information, the trackers cannot recover from drift due to constant update rate and simple linear combination of previous appearance templates.

Learning Dynamic SN Method: Guo *et al.* proposed the DSiam tracker and designed dynamic transformation matrices [46]. Two distinct online transformation matrices including target appearance variation and background suppression are incorporated within the classical SN. Both matrices are solved in the Fourier domain with a closed-form solutions. DSiam provides effective online learning however it ignores the historical target variations which is important for a smoother adaptation of the exemplar template.

Dynamic Memory Network Method: Yang *et al.* proposed the MemTrack that dynamically writes and reads previous templates to cope with target appearance variations [151]. A long term short term memory is used as a memory controller. The input to this network is the search feature map and the network outputs the control signals for the reading and writing process of the memory block. An attention mechanism is also applied with a gated residual template to control the amount of retrieved memory that is used to combine with the initial template. This method enables the tracker to memorize long-term target appearance. However, it only focuses on combining the previous target features, ignoring the discriminative information in background clutter, which leads to an accuracy gap in the presence of drastic target variations. To address this lack, Yang *et al.* [152] proposed a negative memory unit that stores templates for distractors, which is used to cancel out wrong responses from the object template.

Gradient-Guided Method: Li *et al.* proposed GradNet in which gradient information is encoded for updating the target template via feedforward and backward operations [75]. The tracker utilizes the information from the gradient to update the template in the current frame and then incorporates the adaptation process to simplify the process of gradient-based optimization. Unlike aforementioned methods, this method makes full use of the discriminative information in backward gradients instead of just integrating previous templates. This results in a performance improvement as compared to other methods however, computing gradient in a backpropagation manner introduces a computational burden.

UpdateNet Method: Zhang *et al.*, proposed the UpdateNet [161] which learns a generic function θ according to $\tilde{z}_t = \theta(\tilde{z}_{gt}, \tilde{z}_{t-1}, z_t)$. The learned function θ computes the updated template based on initial ground-truth template z_{gt} , the last accumulated template \tilde{z}_{t-1} and the template z_t extracted from the predicted target location in the current frame. In general, the function updates the previous accumulated template \tilde{z}_{t-1} by integrating the new information given by the current frame z_t . Therefore, θ can be adapted to the specific updating requirements of the current frame, based on the difference between the current and accumulated templates. Moreover, it also considers the initial template z_{gt} in every frame, which provides highly reliable information and increases robustness against model drift. The function θ is implemented as a CNN, which grants great expressive power and the ability to learn from large amounts of data. The results demonstrate excellent performance compared to baseline trackers and it has also recently been adopted by CSA tracker [148].

While a number of other techniques for model update have been proposed [103], [153], simply using no update has remain a surprisingly robust and popular alternative [72], [130]. Without a model update, the tracker may better recover from tracking failures (if the object appearance reverts to the first frame), and is not susceptible to model drift. Further research in this direction is required to develop simple, general, and end-to-end trainable techniques, which could further improve the robustness of Siamese tracking.

5.2.2 Loss Functions

The tracking performance also relies on the loss functions employed within the SNs. Different loss functions have been used in the SNs either for regression, classification or for both tasks. Below, we summarize these developments in more detail.

Logistic Loss: The classical SiamFC employed logistic loss defined in (14)-(15) [4]. A variety of other trackers including DSiam [46], RASNET [137], SA-SIAM [50], CFNET [128], SiamDW [167], and GradNet [75] etc., have used logistic loss to train their models built upon SiamFC. This training method utilizes the pairwise relationship on image pairs by maximizing the similarity scores on target-positive pairs and minimizing them on target-negative pairs.

Contrastive Loss: The margin contrastive loss is defined as [19]:

$$\mathcal{L}(x_i, z_i, y_{xz}) = \frac{1}{2}D^2 + \frac{1}{2}(1 - y_{xz}) \max(0, \epsilon - D^2), \quad (18)$$

$$D = \|f_\rho(x_i) - f_\rho(z_i)\|_2,$$

where ϵ is the minimum distance margin that pairs depicting different objects should satisfy, D is the Euclidean distance of l_2 -normalized feature representations, and $y_{xz} \in \{0, 1\}$ indicates whether x_i and z_i are the same object or not. SINT [127] employed the contrastive loss while GOTURN [52] employed L_1 loss between the predicted and the ground-truth bounding boxes.

Triplet Loss: The above loss exploits the pairwise relationship between images only and ignores the underlying structural connections between the positive and negative instances of the target. Yan *et al.* proposed SPLT tracker [150] in which the triplet loss [113], [139] is employed during training. The triplet loss is defined as

$$\mathcal{L}(x, x_i^p, x_i^n) = \sum_i^M \left[\|f_\rho(x_i) - f_\rho(x_i^p)\|_2^2 - \|f_\rho(x_i) - f_\rho(x_i^n)\|_2^2 + \epsilon \right]_+ \quad (19)$$

where x_i^p denotes the positive patch of the target image x i.e., one of other images of the target and x_i^n is a negative patch of any other target or background. \mathcal{T} is the set of all possible triplets in the training set and has cardinality M . The triplet loss not only can further mine the potential relationship among target, positive and negative instances, but also contains more robust similarity structure. Dong *et al.*, proposed SiamFC-Tri in which the probabilistic triplet loss is used [33]. The Siam R-CNN tracker has also been trained using the same loss [130].

Cross Entropy Loss: The classification component in the SNs are normally borrowed from object detection methods [108]. To incorporate this branch, a cross-entropy loss (\mathcal{L}_{cls}) is used which is equivalent to the logistic loss defined in (14) over two classes (target vs. not target).

Li *et al.*, proposed the SiamRPN tracker in which cross entropy loss is employed [73]. Other trackers such as SiamRPN++ [72], SiamAttn [158], Ocean [168], CLNET [34], SPM [131], C-RPN [38] etc., have also been built upon the SiamRPN tracker by training classification branch using the cross entropy loss.

Regression Loss: To train a regression network, three types of loss functions are employed including the smooth L_1 norm [108], the Intersection over Union (IoU) loss [157], and regularized linear regression [112]. The smooth L_1 loss is used in Faster R-CNN for bounding box regression [108]. In the SiamRPN tracker [73], this norm is used to train the regression branch. Following this study, other trackers including SiamRPN++ [72], SiamAttn [158], CLNET [34], SPM [131], and C-RPN [38] have also trained the regression branch of the tracker using the $smooth_{L_1}$ loss. The IoU loss for regression \mathcal{L}_{reg}^{IoU} is defined as [157]:

$$\mathcal{L}_{reg}^{IoU} = - \sum_j \ln(\text{IoU}(p, r)), \quad (20)$$

where p and r denote the predicted and groundtruth bounding box coordinates. Chen *et al.* proposed the SiamBAN tracker in which the regression branch of the network is trained using \mathcal{L}_{reg}^{IoU} [15]. The Ocean [168] and SiamFC++ [145] trackers have also utilized this loss during training.

Multi-Task Loss: For the joint training of classification and regression branches, multi-task loss has also been used in SNs. For instance, a sum of the cross-entropy loss (14) and the regression loss (20) (i.e., $\mathcal{L}_{cls} + \mathcal{L}_{reg}$) is used in SiamRPN [73]. Trackers such as SiamRPN++ [72], SiamAttn [158], CLNET [34], and SPM [131] have also employed the multi-task training method. Moreover, a sum of cross entropy loss (14) and IoU loss for regression (20) (i.e., $\mathcal{L}_{cls} + \mathcal{L}_{reg}^{IoU}$) has also been utilized by Ocean [168], SiamBAN [15], and SiamFC++ [145].

Regularized Linear Regression: To regularize SNs with a correlation filter as a separate layer, a linear regression loss defined in (9)-(10) is employed in many Siamese trackers including CFNET [128], TADT [78], RTINET [154], DSiam [46], FlowTrack [175], UDT [133], and UDT++ [134] etc. The ridge regression problem is then solved by a closed-form solution and the filter is trained in an




Fig. 5: Sample images from different tracking benchmark datasets. The ground-truth bounding box annotation is overlaid.

end-to-end fashion. In a broad sense, the training of these trackers are somewhat similar to end-to-end DCFs training. The classical SiamFC locate the target object using powerful embedding without online model update. However, using a fixed target template to compare appearance prevents the learning from exploiting any video-specific cues that could be helpful for discrimination. In contrast to the classical Siamese tracking, using correlation filter as a separate layer in these trackers helps to discriminate target from the background regions, thereby exploiting target-specific information in the video [128]. *Currently, there is no general consensus in the literature regarding the employed loss function. Instead, recent SOTA methods adopt different alternatives. Among the aforementioned approaches, the Cross Entropy loss have remained a popular choice, also for recent trackers [72].*

6 EXPERIMENTAL COMPARISON

Here, we thoroughly analyze the performance of 59 DCFs and 33 Siamese-based trackers. The performance of these trackers have been quantitatively compared on nine tracking benchmarks: On-line Tracking Benchmark 100 (OTB100) [140], Temple Color 128 (TC128) [82], Unmanned Aerial Vehicle 123 (UAV) [99], Visual Object Tracking 2014 (VOT) [66], VOT2016 [47], VOT2018 [67], TrackingNet [101], Large-scale Single Object Tracking (LaSOT) [36], and Generic Object Tracking 10,000 (GOT-10K) [60]. Fig. 5 shows example frames from different tracking benchmarks. The quantitative results of the compared trackers are either taken directly from respective papers or from other papers. We have also evaluated representative trackers and reproduced the tracking performance for fair comparison.

6.1 Tracking Datasets

To provide a standard and fair performance evaluation of visual trackers, a number of benchmarks have been proposed with the passage of time. In addition to short-term tracking, several recent datasets provide both short and long-term tracking sequences. The publicly available benchmark datasets contain a variety of tracking challenges, including Scale Variation (SV), Out-of-View (OV), DEformation (DEF), Low Resolution (LR), Illumination Variation (IV), Out-of-Plane rotation (OPR), OCclusion (OCC), Background Clutter (BC), Fast Motion (FM), In-Plane Rotation (IPR), Motion Blur (MB), Partial OCclusion (POC), abrupt Camera Motion (CM), Aspect Ratio Change (ARC), Full OCclusion (FOC), Viewpoint Change (VC), Similar OBject (SOB), Object Color Change (OCC), Absolute Motion (AM), target ROTation (ROT), Scene COmplexity (SCO), Fast Camera Motion (FCM), Low Resolution Objects (LRO), and MOTion Change (MOC). Table 1 presents the description of each dataset employed in our experimental comparison. Please see supplementary material for more brief description of each tracking dataset.

TABLE 1: Details of different tracking benchmarks used in our experimental comparison.

Description	OTB100	TC128	UAV123	UAV20L	VOT2014	VOT2016	VOT2018	VOT2018-LT	LaSOT	GOT-10K	TrackingNet
Publication	PAMI2015 [140]	TIP2015 [82]	ECCV2016 [99]	ECCV2016 [99]	ECCV-W2014 [66]	ECCV-W2016 [47]	ICCV-W2017 [67]	ECCV-W2018 [67]	CVPR2019 [36]	PAMI2019 [60]	ECCV2018 [101]
Sequences	100	128	123	20	25	60	60	35	Total: 1400 Training: 1120 Testing: 280	Total: 10k Training: 9,34k Testing: 420	Total: 30.643k Training: 30.132k Testing: 511
Boxes	58.61k	55.652k	113.476k	58.67k	10k	21.455k	21.356k	146.847k	Total: 3.52M Training: 2.8M Testing: 685K	Total: 1.5M Training: 1.4M Testing: 56k	Total: 14M Training: 14M Testing: 226k
Attributes	11	11	12	12	12	12	12	12	14	6	15
Classes	22	27	9	9	11	24	24	21	70	Total: 563 Training: 480 Testing: 84	21
Minimum Frames	71	71	109	1717	164	48	41	1389	1000	51	96
Maximum Frames	3872	3872	3085	5527	1210	1507	1500	29700	11397	920	2368
Average Resolution	356 × 530	461 × 737	1231 × 699	1231 × 699	448 × 304	757 × 480	758 × 465	896 × 468	632 × 1089	929 × 1638	591 × 1013
Average Duration	19.68s	15.6s	30.48s	97.8s	13.68s	357.6s	355.9s	139.85s	83.57s	15s	16.7s
Attributes Name	IV, SV, OCC, DEF, MB, FM, IPR, OPR, OV, BC, LR, SOB, SV, VC, 30 FPS	IV, SV, OCC, DEF, MB, FM, IPR, OPR, OV, BC, LR, SOB, SV, VC, 30 FPS	ARC, BC, CM, FM, FOC, IV, LR, OV, POC, SOB, SV, VC, 30 FPS	ARC, BC, CM, FM, FOC, IV, LR, OV, POC, SOB, SV, VC, 30 FPS	IV, SV, OCC, DEF, MB, BC, ARC, CM, MOC, DEF, MB, BC, IV, SV, OCC, 30 FPS	OCO, SCO, AM, ARC, CM, MOC, DEF, MB, BC, IV, SV, OCC, 30 FPS	OCO, SCO, AM, ARC, CM, MOC, DEF, MB, BC, IV, SV, OCC, 30 FPS	IV, SV, OCC, DEF, MB, BC, ARC, CM, MOC, OCO, SCO, AM, 30 FPS	IV, SV, DEF, MB, FM, OV, BC, LR, ARC, CM, POC, VC, ROT, 30 FPS	IV, SV, DEF, MB, FM, IPR, OPR, OV, BC, LR, ARC, CM, FOC, POC, SOB, 30 FPS	IV, SV, DEF, MB, FM, IPR, OPR, OV, BC, LR, ARC, CM, FOC, POC, SOB, 30 FPS
Frame Rate	30 FPS	30 FPS	30 FPS	30 FPS	30 FPS	30 FPS	30 FPS	30 FPS	30 FPS	30 FPS	30 FPS
Overlapped Datasets	OTB50 VOT TC128	OTB50 VOT OTB100	VOT	VOT	OTB50 TC128 ALOVS++ [117]	OTB100 TC128 ALOVS++ [117]	OTB100, NUS-PRO [71] ALOVS++ [117] TC128, UAV123	OTB100, NUS-PRO [71] ALOVS++ [117] TC128, UAV123	YouTube ImageNet	VOT WordNet [98] ImageNet	YouTube-BB

6.2 Performance Evaluation Measures

To compare the performance of trackers, different evaluation metrics have been proposed in the literature that evaluate the effectiveness in terms of robustness, accuracy and speed.

Precision Plot: The precision plot is based on the central location error which is defined as the average Euclidean distance between the predicted centers of the target object and the ground truth centers in a frame [140]. However, this error does not compute the tracking performance accurately. Therefore distance precision is employed, which is defined as the percentage of frames where the target object is located within a center location error of T pixels [2]. The trackers are ranked using this metric with a threshold of $T = 20$ pixels. The precision plot is generated by plotting the distance precision over a range of thresholds.

Success Plot: The precision metric only measures the localization performance of the tracker which is not accurate to measure the target scale variations [140]. Instead of the center location error, the IoU is employed to measure the prediction error. Given the estimated bounding box p and ground-truth bounding box g , the IoU is defined as $\frac{p \cap g}{p \cup g}$. The success rate is thus the percentage of frames where the IoU is smaller than a T . The success plot is generated by varying the overlap threshold from 0 to 1. The trackers are ranked using area under the curve of the success plot.

Normalized Precision Plot: As distance precision is sensitive to the target scale, Muller *et al.*, employed normalized precision metric to evaluate trackers based on the relative error [101]. It computes errors relative to the target size instead of considering the absolute distance i.e., $\|W(p_x, p_y) - (g_x, g_y)\|$, where $W = \text{diag}(g_x, g_y)$. This relative error is then plotted in the range of 0 to 0.5. The area under this curve is called normalized precision which is used to rank trackers.

Average Overlap: This metric estimates the average of overlaps between the ground-truth and estimated bounding boxes, as in the success plot [60].

SR_{0.50} and SR_{0.75}: These metrics denote the success rate that measures the percentage of successfully tracked frames, where the overlap precision exceeds a threshold of 0.50 and 0.75

The One pass evaluation criteria is used as defined in [140] to measure the tracking performance in terms of precision and success plots on OTB100, TC128, UAV123, and LaSOT datasets. The trackers on these datasets are evaluated by initializing bounding box on the first frame and letting it run until the end of the sequence. In the VOT series, a tracker is reset once it drifts off the target. Following the VOT evaluation protocols [47], [66], [68], the trackers are compared in terms of Accuracy (A), Robustness (R), and Expected Average Overlap (EAO) metrics. A is the average overlap between the predicted and ground truth bounding

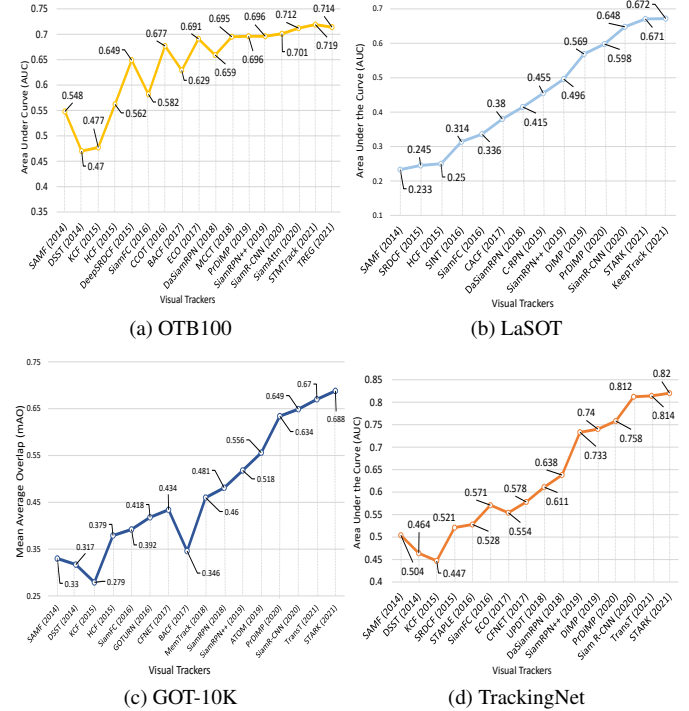


Fig. 6: Tracking performance trends on popular benchmarks: (a) OTB100, (b) LaSOT, (c) GOT-10k and (d) TrackingNet, during 2014 - 2020. The representative trackers are: SAMF [79], DSST [27], KCF [53], HCF [93], SRDCF [30], DeepSRDCF [29], STAPLE [3], SiamFC [4], GOTURN [52], SINT [127], CACF [100], ECO [25], CCOT [23], BACF [42], CSR-DCF [92], CFNET [128], MemTrack [151], UPDT [7], SiamRPN [73], DaSiamRPN [174], DRT [120], C-RPN [38], MCCT [135], DiMP [5], SiamMask [138], D3S [91], PrDiMP [26], ATOM [24], SiamRPN++ [72], SiamR-CNN [130], SiamAttn [158], TransT [14], STARK [146], and STMTrack [41]

boxes during successful tracking periods. R measures how many times the tracker loses the target (fails) during tracking. A reset mechanism starts after some frames once tracker losses the target object. EAO is an estimator of the average overlap a tracker is expected to attain on a large collection of short-term sequences with the same visual properties as the given dataset.

6.3 Quantitative Comparison

Tables 2, 3, 4, and 5 present the performance comparison of representative DCF-based trackers on nine tracking benchmarks.³ While earlier DCF-based trackers employing deep features achieve promising performance on OTB100, they provide inferior results on recent more challenging large-scale datasets such as LaSOT.

³For a more detailed comparison, please see supplementary material.

TABLE 2: Performance comparison of representative DCF and Siamese trackers, in terms of Precision Rate (PR) at a threshold of 20 pixels, Normalized Precision Rate (NPR), and Area Under the Curve (AUC), on OTB100, TC128, UAV123, UAV20L, and LaSOT datasets. For both DCF and Siamese trackers, the best two results are in red and blue font, respectively.

DCF Trackers	Publication	OTB100		TC128		UAV123		UAV20L		LaSOT			Speed (fps)	
		PR	AUC	PR	AUC	PR	AUC	PR	AUC	PR	NPR	AUC	CPU	GPU
HCF [93]	ICCV2015	0.837	0.562	0.703	0.488	-	-	0.500	0.299	0.241	0.286	0.250	-	10.4
ECO [25]	CVPR2017	0.910	0.691	0.800	0.605	0.741	0.537	0.604	0.435	0.301	0.338	0.324	60	8
UPDT [7]	ECCV2018	0.931	0.702	-	0.622	-	0.550	-	-	-	-	-	-	-
GFS-DCF [143]	ICCV2019	0.932	0.693	-	-	-	-	-	-	-	-	-	-	8
ASRCF [21]	CVPR2019	0.922	0.692	0.825	0.603	-	-	-	-	0.337	0.391	0.359	-	28
ATOM [24]	CVPR2019	0.879	0.667	-	-	0.856	0.650	-	-	0.505	0.576	0.515	-	30
DiMP [5]	ICCV2019	0.902	0.684	-	-	0.858	0.654	-	-	0.563	0.650	0.569	-	40
Siamese Trackers	Publication	OTB100		TC128		UAV123		UAV20L		LaSOT			Speed (fps)	
		PR	AUC	PR	AUC	PR	AUC	PR	AUC	PR	NPR	AUC	CPU	GPU
SiamFC [4]	ECCV2016	0.771	0.582	0.688	0.503	0.691	0.461	0.613	0.399	0.339	0.420	0.336	-	86
DaSiamRPN [174]	ECCV2018	0.880	0.659	-	-	0.796	0.586	0.838	0.617	-	0.496	0.415	-	160
SiamRPN [73]	CVPR2018	0.851	0.637	-	-	0.796	0.527	0.617	0.454	-	-	0.433	-	160
SiamRPN++ [72]	CVPR2019	0.914	0.696	-	-	0.807	0.613	-	-	0.491	0.569	0.496	-	35
Ocean [168]	ECCV2020	0.920	0.684	-	-	-	-	-	-	0.566	-	0.560	-	26
SiamBAN [15]	CVPR2020	0.910	0.696	-	-	0.833	0.631	-	-	0.521	0.598	0.514	-	40
SiamR-CNN [130]	CVPR2020	0.891	0.701	-	0.649	0.834	0.649	-	-	0.684	0.722	0.648	-	4.7
SiamAttn [158]	CVPR2020	0.926	0.712	-	-	0.845	0.650	-	-	-	0.648	0.560	-	45

For instance, ECO achieves an impressive PR score of 91.0% on OTB but obtains a PR score of only 30.1% on LaSOT. In contrast, recent end-to-end DCF frameworks such as, DiMP and its successor PrDiMP achieve impressive performance on OTB100 as well as on LaSOT. For instance, PrDiMP achieves PR scores of 90.3%, 87.8% and 60.9% on OTB100, UAV123, and LaSOT, respectively. Among existing DCF-based trackers, DiMP and PrDiMP achieve superior results on most benchmarks. PrDiMP achieves top performance on UAV123, LaSOT, and GOT-10K, while also achieving competitive results (among top-three) on OTB100, VOT2016, and VOT2018-ST. The success of these modern DCF trackers (DiMP and [6], [26]) is due to their efficient end-to-end trainable architectures that are capable of learning a discriminative target model prediction by fully utilizing both target and background appearance information. These trackers employ a dedicated optimization process to learn a powerful model in few iterations. For instance, PrDiMP utilizes a more general Newton approximation for addressing KL-divergence objective. Further, these modern trackers comprise a dedicated target estimation component to perform deep bounding box regression and also circumvent the problem of boundary artifacts.

Tables 2, 3, 4, and 5 also present the performance comparison of representative Siamese trackers on nine benchmarks³. Among recent Siamese methods, we observe trackers to focus on different fundamental issues such as, online model update, re-detection components, improved region refinement, effective box regression, and bridging the gap between object tracking and object segmentation. For instance, SiamAttn introduces an attention mechanism to adaptively update the target template and obtains the best performance on OTB100, UAV123, VOT2016, while also achieving competitive results (among top-three) on LaSOT, VOT2018-ST, and TrackingNet. SiamAttn obtains AUC scores of 71.2%, 65.0%, 56.0%, and 75.2% on OTB100, UAV123, LaSOT, and TrackingNet, respectively. Further, it obtains EAO scores of 53.7% and 47.0% on VOT2016 and VOT2018-ST, respectively. SiamRPN introduces anchor-based RPN component for accurate scale estimation and obtains better performance (EAO of 34.4% on VOT2016) than the SiamFC. SiamBAN introduces anchor-free bounding box regression and obtains the top performance of 50.5% over the SiamRPN on VOT2016 dataset. SiamR-CNN introduces a re-detection architecture combined with a tracklet-based dynamic programming scheme and obtains top performance

on TC128 (64.9% AUC score), LaSOT (64.8% AUC score), GOT-10K (64.9% mAO score) and TrackingNet (81.2% AUC score), while also achieving competitive results (among top-three) on other datasets. Ocean introduces an approach to refine the imprecise bounding-box predictions along with learning object-aware features and obtains the top performance on VOT2018-ST (48.9% EAO score). D3S is a single-shot segmentation tracker employing two target models with complementary properties and obtains the best results on VOT2018-ST and VOT2020-ST (48.9% and 43.9% EAO score).

Figure 6 shows the tracking performance improvement trend on different benchmarks (OTB100, LaSOT, GOT-10k, and TrackingNet) in recent years. We can observe that the performance on OTB100 has saturated in recent years with several visual trackers obtaining over 90% PR score (Table 2), likely due to numerous relatively easy videos. While, the recently introduced LaSOT, GOT10K, and TrackingNet all show a similar trend with consistent improvements obtained by recent trackers on these datasets. We also observed a similar trend on the VOT dataset in Figure 1. For instance, the best reported AUC score on LaSOT is still around 65%. Similarly, there is still a significant room to further improve the tracking performance in the VOT dataset, despite witnessing an impressive leap in performance in recent years. This suggests that these new challenging benchmarks are still very challenging for SOTA trackers and their introduction is significantly contributing to pushing the boundaries of visual tracking research.

We have also evaluated some representative SOTA trackers using same training datasets, hyper-parameters tuning on same dataset, and speed comparison on same machine. Please refer to the supplementary material (Section II) for more implementation details. Table 6 shows the performance comparison of these representative trackers. Overall, the offline trackers demonstrated consistent performance improvement on VOT datasets. For instance, DiMP is the top performer among the DCFs-based trackers on both VOT datasets achieving an EAO scores of 49.1% and 45.1%, respectively. Similarly, SiamCAR demonstrated best performance among the Siamese trackers. The anchor-free bounding box regression component for scale estimation has also demonstrated promising performance improvements in both SiamBAN and SiamCAR trackers. For instance, compared to anchor-based scale estimation used in SiamRPN, anchor-free bounding box regression in SiamCAR achieved 12.0% improvement and 21.9%

TABLE 3: Performance comparison of representative DCF and Siamese trackers, in terms of EAO, A, and R, on VOT2016 and VOT2018-ST. Best two results are shown in red and blue fonts for both paradigms, respectively.

DCF Trackers	Publication	VOT2016			VOT2018-ST			Speed
		EAO \uparrow	A \uparrow	R \downarrow	EAO \uparrow	A \uparrow	R \downarrow	
HCF [93]	ICCV2015	0.220	0.450	0.396	-	-	-	1.057
ECO [25]	CVPR2017	0.374	0.546	11.67	0.280	0.484	0.276	-
UPDT [7]	ECCV2018	-	-	-	0.378	0.536	0.182	-
GFS-DCF [143]	ICCV2019	-	-	-	0.397	0.511	0.143	-
ASRCF [21]	CVPR2019	0.391	0.563	0.187	0.328	0.494	0.234	-
ATOM [24]	CVPR2019	0.424	0.617	0.190	0.401	0.590	0.204	-
DiMP [5]	ICCV2019	0.479	0.617	0.190	0.440	0.597	0.153	-
Siamese Trackers		EAO \uparrow	A \uparrow	R \downarrow	EAO \uparrow	A \uparrow	R \downarrow	Speed
SiamFC [4]	ECCV2016	0.235	0.532	0.461	0.188	0.503	0.585	9.213
DaSiamRPN [174]	ECCV2018	0.411	0.610	0.22	0.326	0.570	0.337	-
SiamRPN [73]	CVPR2018	0.344	0.560	1.12	0.244	0.490	0.464	-
SiamRPN++ [72]	CVPR2019	0.370	0.580	0.240	0.414	0.600	0.234	-
Ocean [168]	ECCV2020	-	-	-	0.489	0.592	0.117	-
SiamBAN [15]	CVPR2020	0.505	0.632	0.150	0.452	0.597	0.178	-
SiamR-CNN [130]	CVPR2020	0.461	0.645	0.173	0.408	0.609	0.220	-
SiamAttn [158]	CVPR2020	0.537	0.680	0.140	0.470	0.630	0.160	-

TABLE 4: Performance comparison of representative DCF and Siamese trackers, in terms of EAO, A, and R on VOT2020-ST. In case of DCF and Siamese trackers, the best two results are in red and blue fonts, respectively.

DCF Trackers	Publication	VOT2020-ST		
		EAO \uparrow	A \uparrow	R \uparrow
KCF [53]	PAMI2015	0.154	0.407	0.432
CSR-DCF [92]	CVPR2017	0.193	0.406	0.582
UPDT [7]	ECCV2018	0.278	0.465	0.755
DiMP [5]	ICCV2019	0.274	0.457	0.740
ATOM [24]	CVPR2019	0.271	0.462	0.734
Siamese Trackers		EAO \uparrow	A \uparrow	R \uparrow
SiamFC [4]	ECCV2016	0.179	0.418	0.502
SiamMask [138]	CVPR2019	0.321	0.624	0.648
D3S [91]	CVPR2020	0.439	0.699	0.769
Ocean [168]	ECCV2020	0.430	0.693	0.754

improvement compared to classical SiamFC tracker.

The aforementioned datasets also have radically different properties and characteristics. LaSOT and UAV123 contain long sequences and multiple distractors. Trackers achieving high performance here demonstrate substantial robustness and re-detection capabilities. We observe that the recent trackers DiMP and PrDiMP achieve strong results, and that the distractor-aware track generation in SiamR-CNN improves the robustness in such scenarios. In contrast to LaSOT, TrackingNet, and GOT10k contain short sequences where robustness and re-detection capability is of much lesser importance. Instead, these datasets reward trackers with highly accurate bounding box prediction, such as SiamR-CNN and PrDiMP. Among Siamese trackers, we observe that SiamR-CNN and SiamAttn achieve the most consistently good results across several datasets. The exception being LaSOT in case of SiamAttn, while SiamR-CNN struggles on VOT. Among DCF-based methods, PrDiMP consistently achieves SOTA results across all evaluated datasets.

6.4 Speed Comparison

The tracking speed is another very important metric to evaluate the trackers especially to meet the real-time requirements. However, evaluating the tracking speed is not straightforward since a number of key-factors influence including feature extraction, model update method, programming language, and most importantly the hardware that the trackers are implemented on. We evaluated the tracking speed of representative trackers in terms of Frames Per Second (FPS) shown in Table 6 using Tesla V100 GPUs. Please see the machine specifications in the supplementary material. SPM and SiamRPN trackers clearly show the best FPS while DiMP tracker shows competitive performance but worst speed.

TABLE 5: Performance comparison of representative DCF and Siamese trackers, in terms of mean Average Overlap (mAO), mSR_{0.50}, and mSR_{0.75}, PR, NPR, and AUC, on GOT-10K and TrackingNet. In case of DCF and Siamese-trackers, the best two results are in red and blue fonts, respectively.

DCF Trackers	Publication	GOT-10K			TrackingNet		
		mAO	mSR _{0.50}	mSR _{0.75}	PR	NPR	AUC
HCF [93]	ICCV2015	0.379	0.380	0.134	-	-	-
ECO [25]	CVPR2017	0.395	0.407	0.170	0.492	0.618	0.554
UPDT [7]	ECCV2018	-	-	-	0.557	0.703	0.611
GFS-DCF [143]	ICCV2019	-	-	-	0.565	0.717	0.609
ATOM [24]	CVPR2019	0.556	0.634	0.402	0.648	0.771	0.703
DiMP [5]	ICCV2019	0.611	0.717	0.492	0.800	0.801	0.740
Siamese Trackers		mAO	mSR _{0.50}	mSR _{0.75}	PR	NPR	AUC
SiamFC [4]	ECCV2016	0.392	0.426	0.135	0.533	0.654	0.571
SiamRPN [73]	CVPR2018	0.481	0.581	0.270	-	-	-
SiamRPN++ [72]	CVPR2019	0.518	0.618	0.325	0.694	0.800	0.733
Ocean [168]	ECCV2020	0.611	0.721	-	-	-	-
SiamR-CNN [130]	CVPR2020	0.649	0.728	0.597	0.800	0.854	0.812
SiamAttn [158]	CVPR2020	-	-	-	0.715	0.817	0.752

Overall, the DCF tracking paradigm needs to show significant improvements for tracking speed.

7 DISCUSSION AND CONCLUSIONS

Here, we summarize the common lessons learned and a set of recommendations for future work in generic object tracking.

Importance of end-to-end tracking framework: These frameworks have demonstrated excellent performance recently. While end-to-end offline learning is a prerequisite for Siamese tracking, recent DCF methods have also adopted this paradigm with success. Thus, learning the underlying features, along with prediction heads, have proven crucial for optimal performance. This has only been possible in the last few years, with the introduction of large scale training datasets.

Importance of robust target modeling: While Siamese-based approaches have excelled in many areas, end-to-end DCF-based methods still demonstrate an advantage in challenging long-term tracking scenarios, such as LaSOT. This shows the importance of robust online target appearance modelling, achieved by embedding discriminative learning modules into the network architecture. Such methods effectively integrate background appearance cues and can easily be updated during the tracking procedure using online learning.

Target state estimation: Siamese-based approaches have driven the advancement of more accurate bounding box regression by leveraging progress in the neighboring field of object detection. Recent one-stage (anchor free) based approaches, such as Ocean, achieve simple, accurate, and efficient bounding box regression. Furthermore, these strategies are generic and can easily be integrated in any visual tracking architecture.

Role of segmentation: Although the task of bounding box regression have seen substantial progress in tracking, such a target state model is inherently limited. Instead, segmentation promises a pixel-precise estimation of the target, which is highly desired in many applications. Moreover, segmentation offers the potential of improving the tracking itself, by for example aiding the target model update. Furthermore, as demonstrated in for instance [149], segmentation further aids the regression of accurate bounding boxes and helps estimating the scale at which the object is tracked. Future efforts should therefore be aimed towards integrating accurate segmentation into robust tracking frameworks [109], [142].

Backbone architectures: The ResNet architecture has withstood the test of time in several computer vision applications. In visual tracking, it remains the most popular choice for feature extraction. The architecture is simple, effective, and allows for the extraction of features at multiple resolutions. While effective when pushing the boundaries of SOTA, it is still remains computationally costly

TABLE 6: Performance comparison of evaluated representative DCFs and Siamese trackers, in terms of EAO, A, and R, on VOT2016 and VOT2018-ST. The tracker with (H) represents the handcrafted features and the tracker with (H+D) represents the hybrid features representation used within the tracker. For more details please see supplementary material. In case of DCF and Siamese trackers, the best two results are shown in red and blue fonts, respectively.

DCF Trackers	VOT2016			VOT2018-ST			Speed
	EAO \uparrow	A \uparrow	R \downarrow	EAO \uparrow	A \uparrow	R \downarrow	
DSST (H) [27]	0.211	0.548	0.508	0.111	0.421	1.120	33.60
HCF (ResNet) [93]	0.225	0.419	0.391	0.197	0.346	0.455	10.61
SRDCF (H) [30]	0.260	0.542	0.402	0.126	0.509	0.773	6.66
BACF (H) [42]	0.245	0.583	1.510	0.161	0.492	0.721	40.21
UPDT (H) [7]	0.381	0.501	0.201	0.301	0.455	0.288	15.55
UPDT (H+D) [7]	0.433	0.577	0.101	0.362	0.521	0.195	10.11
GFS-DCF (H) [143]	0.401	0.511	0.161	0.333	0.462	0.211	17.20
GFS-DCF (H+D) [143]	0.451	0.566	0.090	0.382	0.499	0.161	12.23
ASRCF (H) [21]	0.382	0.541	0.201	0.314	0.480	0.261	32.89
ASRCF (H+D) [21]	0.415	0.581	0.140	0.349	0.519	0.228	27.55
DiMP [5]	0.491	0.611	0.170	0.451	0.612	0.141	48.21
Siamese Trackers	EAO \uparrow	A \uparrow	R \downarrow	EAO \uparrow	A \uparrow	R \downarrow	Speed
SiamFC [4]	0.282	0.556	0.293	0.315	0.577	0.312	98.90
DSiam [46]	0.301	0.572	0.212	0.305	0.566	0.221	58.30
DaSiamRPN [174]	0.435	0.625	0.151	0.331	0.588	0.301	166.60
SiamRPN [73]	0.381	0.601	0.771	0.332	0.551	0.211	170.50
SiamRPN++ [72]	0.391	0.581	0.744	0.341	0.540	0.203	82.30
SPM [131]	0.414	0.601	0.256	0.321	0.561	0.311	155.55
SiamBAN [15]	0.491	0.651	0.091	0.467	0.611	0.157	51.20
SiamCAR [45]	0.501	0.633	0.113	0.455	0.591	0.171	56.33
SiamR-CNN [130]	0.472	0.662	0.141	0.421	0.618	0.188	22.30

for real-time applications on platforms with harder computational constraints, such as CPUs. A highly interesting future direction is therefore to develop efficient backbone networks tailored for the tracking task [147]. The fusion strategies to incorporate the fine-grained and semantic information can also be proposed to enhance the discriminative ability of the trackers. Moreover, vision transformer-based backbone architectures can be utilized for tracking task either as a feature extractor module or as an end-to-end training of transformer backbone.

Estimating geometry: In some applications, e.g. in augmented reality, a precise geometric transformation between frames is necessary for the added graphics to appear as attached to objects. For planar objects, at least an affine transformation, but preferably a homography, between a reference and current view is required. For non-planar objects, the problems is linked to online reconstruction of the 3D shape of the object [141]. Neither DCF nor Siamese methods have been equipped to provide precise geometric correspondence, which remains an open research issue.

Role of Transformers: Transformers have recently shown success in a variety of vision tasks [35]. Very recent tracking approaches employ transformers in different ways. [14], [136] utilize transformers for feature enhancement, in combination with either DCF or Siamese trackers. [97] employs a transformer to associate the target object between frames in the presence of distractors. In particular, STARK employs a transformer module for target detection and bounding box regression [146]. In this work, the transformer thus takes on the role of the DCF or Siamese correlation component. The transformer, with its embedded attention module, bares interesting commonalities with DCF. Most importantly, it allows for integration of background appearance information through global operations. Moreover, the transformer employed in STARK predicts a correlation filter. It therefore can be seen as a replacement of the optimization based filter prediction in DCF. Much future effort is needed to further analyze the effectiveness of transformers, as well as their relation to the DCF and Siamese paradigms.

Future directions: In the recent past, with the introduction of precise segmentation capabilities to visual trackers, the connection to video segmentation has become apparent. Some top visual tracking methods perform well on video segmentation benchmarks

[91], [138], despite the restriction to causal processing of the input.

In the near future, we expect convergence with areas as SLAM, which estimates the relative position of the camera model of the scene that is build online, assuming a rigid scene. As soon as the rigidity assumption is dropped, as in multi-body SLAM, the output for the object not corresponding to the background can be seen as tracking of the object, with an estimated 3D shape model added.

The steadily improving performance of single-object tracking, including robustness to nearly identical distractors, demonstrated e.g., in sequences of groups of insects, seem likely to lead to convergence with the research in the area of multi-target tracking. We are probably likely to see strong progress in the Open-world Tracking Problem, as described in [87], i. e., of methods capable of tracking, segmenting, and shape modelling of multiple objects from a-priori unknown classes.

ACKNOWLEDGMENTS

This publication acknowledges the support provided by the Khalifa University of Science and Technology under Faculty Start Up grants FSU-2022-003 Award No. 84740 0 0401. JM was supported by the Research Center for Informatics project CZ.02.1.01/0.0/0.0/16_019/0000765.

REFERENCES

- [1] K. Alex, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional networks," in *NIPS*, 2012.
- [2] B. Babenko, M.-H. Yang, and S. Belongie, "Robust object tracking with online multiple instance learning," in *T-PAMI*, vol. 33, no. 8, pp. 1619–1632, 2010.
- [3] L. Bertinetto, J. Valmadre, S. Golodetz, O. Miksik, and P. H. Torr, "Staple: Complementary learners for real-time tracking," in *CVPR*, 2016.
- [4] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr, "Fully-convolutional siamese networks for object tracking," in *ECCV*. Springer, 2016.
- [5] G. Bhat, M. Danelljan, L. V. Gool, and R. Timofte, "Learning discriminative model prediction for tracking," in *ICCV*, 2019.
- [6] G. Bhat, M. Danelljan, L. Van Gool, and R. Timofte, "Know your surroundings: Exploiting scene information for object tracking," in *ECCV*, 2020.
- [7] G. Bhat, J. Johnander, M. Danelljan, F. Shahbaz Khan, and M. Felsberg, "Unveiling the power of deep tracking," in *ECCV*, 2018.
- [8] G. Bhat, F. J. Lawin, M. Danelljan, A. Robinson, M. Felsberg, L. Van Gool, and R. Timofte, "Learning what to learn for video object segmentation," in *ECCV*, 2020.
- [9] A. Bibi, M. Mueller, and B. Ghanem, "Target response adaptation for correlation filter tracking," in *ECCV*, 2016.
- [10] D. S. Bolme, J. R. Beveridge, B. A. Draper, and Y. M. Lui, "Visual object tracking using adaptive correlation filters," in *CVPR*, 2010.
- [11] S. Boyd, N. Parikh, and E. Chu, *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- [12] J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. LeCun, C. Moore, E. Säckinger, and R. Shah, "Signature verification using a "siamese" time delay neural network," in *IJPRAI*, vol. 7, no. 04, pp. 669–688, 1993.
- [13] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional nets," in *BMVC*, 2014.
- [14] X. Chen, B. Yan, J. Zhu, D. Wang, X. Yang, and H. Lu, "Transformer tracking," in *arXiv preprint arXiv:2103.15436*, 2021.
- [15] Z. Chen, B. Zhong, G. Li, S. Zhang, and R. Ji, "Siamese box adaptive network for visual tracking," in *CVPR*, 2020.
- [16] J. Choi, H. Jin Chang, T. Fischer, S. Yun, K. Lee, J. Jeong, Y. Demiris, and J. Young Choi, "Context-aware deep feature compression for high-speed visual tracking," in *CVPR*, 2018.
- [17] J. Choi, H. Jin Chang, J. Jeong, Y. Demiris, and J. Young Choi, "Visual tracking using attention-modulated disintegration and integration," in *CVPR*, 2016.

- [18] J. Choi, H. Jin Chang, S. Yun, T. Fischer, Y. Demiris, and J. Young Choi, "Attentional correlation filter network for adaptive visual tracking," in *CVPR*, 2017.
- [19] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *CVPR*, 2005.
- [20] P. Chu, H. Fan, C. C. Tan, and H. Ling, "Online multi-object tracking with instance-aware tracker and dynamic model refreshment," in *IEEE WACV*, 2019.
- [21] K. Dai, D. Wang, H. Lu, C. Sun, and J. Li, "Visual tracking via adaptive spatially-regularized correlation filters," in *CVPR*, 2019.
- [22] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *CVPR*, 2005.
- [23] M. Danelljan, A. Robinson, F. K. Khan, and M. Felsberg, "Beyond correlation filters: Learning continuous convolution operators for visual tracking," in *ECCV*, 2016.
- [24] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg, "Atom: Accurate tracking by overlap maximization," in *CVPR*, 2019.
- [25] M. Danelljan, G. Bhat, F. Shahbaz Khan, and M. Felsberg, "Eco: Efficient convolution operators for tracking," in *CVPR*, 2017.
- [26] M. Danelljan, L. V. Gool, and R. Timofte, "Probabilistic regression for visual tracking," in *CVPR*, 2020.
- [27] M. Danelljan, G. Häger, F. Khan, and M. Felsberg, "Accurate scale estimation for robust visual tracking," in *BMVC*, 2014.
- [28] M. Danelljan, G. Häger, F. S. Khan, and M. Felsberg, "Discriminative scale space tracking," in *T-PAMI*, vol. 39, no. 8, pp. 1561–1575, 2016.
- [29] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg, "Convolutional features for correlation filter based visual tracking," in *ICCV-W*, 2015.
- [30] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg, "Learning spatially regularized correlation filters for visual tracking," in *ICCV*, 2015.
- [31] M. Danelljan, F. Shahbaz Khan, M. Felsberg, and J. Van de Weijer, "Adaptive color attributes for real-time visual tracking," in *CVPR*, 2014.
- [32] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *CVPR*, 2009.
- [33] X. Dong and J. Shen, "Triplet loss in siamese network for object tracking," in *ECCV*, 2018.
- [34] X. Dong, J. Shen, L. Shao, and F. Porikli, "Clnet: A compact latent network for fast adjusting siamese trackers," in *ECCV*, 2020.
- [35] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *arXiv preprint arXiv:2010.11929*, 2020.
- [36] 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 *CVPR*, 2019.
- [37] H. Fan and H. Ling, "Parallel tracking and verifying: A framework for real-time and high accuracy visual tracking," in *ICCV*, 2017.
- [38] H. Fan and H. Ling, "Siamese cascaded region proposal networks for real-time visual tracking," in *CVPR*, 2019.
- [39] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," in *T-PAMI*, vol. 32, no. 9, pp. 1627–1645, 2009.
- [40] M. Fiaz, A. Mahmood, S. Javed, and S. K. Jung, "Handcrafted and deep trackers: Recent visual object tracking approaches and trends," in *CSUR*, vol. 52, no. 2, pp. 1–44, 2019.
- [41] Z. Fu, Q. Liu, Z. Fu, and Y. Wang, "Stmtrack: Template-free visual tracking with space-time memory networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13 774–13 783.
- [42] H. K. Galoogahi, A. Fagg, and S. Lucey, "Learning background-aware correlation filters for visual tracking," in *CVPR*, 2017.
- [43] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *CVPR*, 2014.
- [44] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*. in MIT press Cambridge, 2016, vol. 1, no. 2.
- [45] D. Guo, J. Wang, Y. Cui, Z. Wang, and S. Chen, "Siamcar: Siamese fully convolutional classification and regression for visual tracking," in *CVPR*, 2020.
- [46] Q. Guo, W. Feng, C. Zhou, R. Huang, L. Wan, and S. Wang, "Learning dynamic siamese network for visual object tracking," in *ICCV*, 2017.
- [47] S. Hadfield, R. Bowden, and K. Lebeda, "The visual object tracking vot2016 challenge results," in *ECCV-W*, 2016.
- [48] A. Hadjidimos, "Successive overrelaxation (sor) and related methods," in *JCAM*, vol. 123, no. 1-2, pp. 177–199, 2000.
- [49] X. Han, T. Leung, Y. Jia, R. Sukthankar, and A. C. Berg, "Matchnet: Unifying feature and metric learning for patch-based matching," in *CVPR*, 2015.
- [50] A. He, C. Luo, X. Tian, and W. Zeng, "A twofold siamese network for real-time object tracking," in *CVPR*, 2018.
- [51] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [52] D. Held, S. Thrun, and S. Savarese, "Learning to track at 100 fps with deep regression networks," in *ECCV*, 2016.
- [53] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "High-speed tracking with kernelized correlation filters," in *T-PAMI*, vol. 37, no. 3, pp. 583–596, 2015.
- [54] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "Exploiting the circulant structure of tracking-by-detection with kernels," in *ECCV*, 2012.
- [55] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "High-speed tracking with kernelized correlation filters," in *T-PAMI*, vol. 37, no. 3, pp. 583–596, 2014.
- [56] Z. Hong, Z. Chen, C. Wang, X. Mei, D. Prokhorov, and D. Tao, "Multi-store tracker (muster): A cognitive psychology inspired approach to object tracking," in *CVPR*, 2015, pp. 749–758.
- [57] R. A. Horn and C. R. Johnson, *Matrix analysis*. in CUP, 2012.
- [58] C. Huang, S. Lucey, and D. Ramanan, "Learning policies for adaptive tracking with deep feature cascades," in *ICCV*, 2017.
- [59] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama *et al.*, "Speed/accuracy trade-offs for modern convolutional object detectors," in *CVPR*, 2017.
- [60] L. Huang, X. Zhao, and K. Huang, "Got-10k: A large high-diversity benchmark for generic object tracking in the wild," in *T-PAMI*, 2019.
- [61] Z. Huang, C. Fu, Y. Li, F. Lin, and P. Lu, "Learning aberrance repressed correlation filters for real-time uav tracking," in *ICCV*, 2019.
- [62] B. Jiang, R. Luo, J. Mao, T. Xiao, and Y. Jiang, "Acquisition of localization confidence for accurate object detection," in *ECCV*, 2018.
- [63] U. Kart, A. Lukezic, M. Kristan, J.-K. Kamarainen, and J. Matas, "Object tracking by reconstruction with view-specific discriminative correlation filters," in *CVPR*, 2019.
- [64] H. Kiani Galoogahi, T. Sim, and S. Lucey, "Multi-channel correlation filters," in *ICCV*, 2013.
- [65] H. Kiani Galoogahi, T. Sim, and S. Lucey, "Correlation filters with limited boundaries," in *CVPR*, 2015.
- [66] M. Kristan *et al.*, "The visual object tracking vot2014 challenge results," in *ECCV-W*, 2014, pp. 191–217.
- [67] M. Kristan *et al.*, "The sixth visual object tracking vot2018 challenge results," in *ECCV-W*, 2018.
- [68] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Čehovin Zajc, T. Vojir, G. Hager, A. Lukezic, A. Eldesokey *et al.*, "The visual object tracking vot2017 challenge results," in *ICCV-W*, 2017.
- [69] H. Law and J. Deng, "Cornernet: Detecting objects as paired keypoints," in *ECCV*, 2018.
- [70] L. Leal-Taixé, C. Canton-Ferrer, and K. Schindler, "Learning by tracking: Siamese cnn for robust target association," in *IEEE CVPR-W*, 2016.
- [71] A. Li, M. Lin, Y. Wu, M. Yang, and S. Yan, "NUS-PRO: A New Visual Tracking Challenge," in *T-PAMI*, vol. 38, no. 2, pp. 335–349, 2016.
- [72] B. Li, W. Wu, Q. Wang, F. Zhang, J. Xing, and J. Yan, "Siamrpn++: Evolution of siamese visual tracking with very deep networks," in *CVPR*, 2019.
- [73] B. Li, J. Yan, W. Wu, Z. Zhu, and X. Hu, "High performance visual tracking with siamese region proposal network," in *CVPR*, 2018.
- [74] F. Li, C. Tian, W. Zuo, L. Zhang, and M.-H. Yang, "Learning spatial-temporal regularized correlation filters for visual tracking," in *CVPR*, 2018.
- [75] P. Li, B. Chen, W. Ouyang, D. Wang, X. Yang, and H. Lu, "Gradnet: Gradient-guided network for visual object tracking," in *ICCV*, 2019.
- [76] P. Li, D. Wang, L. Wang, and H. Lu, "Deep visual tracking: Review and experimental comparison," in *PR*, vol. 76, pp. 323–338, 2018.
- [77] X. Li, W. Hu, C. Shen, Z. Zhang, A. Dick, and A. V. D. Hengel, "A survey of appearance models in visual object tracking," in *T-IST*, vol. 4, no. 4, pp. 1–48, 2013.
- [78] X. Li, C. Ma, B. Wu, Z. He, and M.-H. Yang, "Target-aware deep tracking," in *CVPR*, 2019.
- [79] Y. Li and J. Zhu, "A scale adaptive kernel correlation filter tracker with feature integration," in *ECCV-W*, 2014.
- [80] Y. Li, J. Zhu, and S. C. Hoi, "Reliable patch trackers: Robust visual tracking by exploiting reliable patches," in *CVPR*, 2015.

- [81] Y. Li, C. Fu, F. Ding, Z. Huang, and G. Lu, "Autotrack: Towards high-performance visual tracking for uav with automatic spatio-temporal regularization," in *CVPR*, 2020.
- [82] P. Liang, E. Blasch, and H. Ling, "Encoding color information for visual tracking: Algorithms and benchmark," in *T-IP*, vol. 24, no. 12, pp. 5630–5644, 2015.
- [83] T.-Y. Lin, Y. Cui, S. Belongie, and J. Hays, "Learning deep representations for ground-to-aerial geolocalization," in *CVPR*, 2015.
- [84] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *ECCV*. Springer, 2014.
- [85] S. Liu, T. Zhang, X. Cao, and C. Xu, "Structural correlation filter for robust visual tracking," in *CVPR*, 2016.
- [86] T. Liu, G. Wang, and Q. Yang, "Real-time part-based visual tracking via adaptive correlation filters," in *CVPR*, 2015.
- [87] Y. Liu, I. E. Zulfikar, J. Luiten, A. Dave, A. Ošep, D. Ramanan, B. Leibe, and L. Leal-Taixé, "Opening up open-world tracking," *arXiv preprint arXiv:2104.11221*, 2021.
- [88] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," *arXiv preprint arXiv:2103.14030*, 2021.
- [89] X. Lu, C. Ma, J. Shen, X. Yang, I. Reid, and M.-H. Yang, "Deep object tracking with shrinkage loss," in *T-PAMI*, 2020.
- [90] X. Lu, W. Wang, C. Ma, J. Shen, L. Shao, and F. Porikli, "See more, know more: Unsupervised video object segmentation with co-attention siamese networks," in *CVPR*, 2019.
- [91] A. Lukezic, J. Matas, and M. Kristan, "D3s-a discriminative single shot segmentation tracker," in *CVPR*, 2020.
- [92] A. Lukezic, T. Vojir, L. Čehovin Zajc, J. Matas, and M. Kristan, "Discriminative correlation filter with channel and spatial reliability," in *CVPR*, 2017.
- [93] C. Ma, J. B. Huang, X. Yang, and M. H. Yang, "Hierarchical convolutional features for visual tracking," in *ICCV*, 2015.
- [94] C. Ma, J. Huang, X. Yang, and M. Yang, "Robust visual tracking via hierarchical convolutional features," in *T-PAMI*, vol. 41, no. 11, pp. 2709–2723, 2019.
- [95] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang, "Adaptive correlation filters with long-term and short-term memory for object tracking," in *IJCV*, vol. 126, no. 8, pp. 771–796, 2018.
- [96] C. Ma, X. Yang, C. Zhang, and M.-H. Yang, "Long-term correlation tracking," in *CVPR*, 2015.
- [97] C. Mayer, M. Danelljan, D. P. Paudel, and L. Van Gool, "Learning target candidate association to keep track of what not to track," in *arXiv preprint arXiv:2103.16556*, 2021.
- [98] G. A. Miller, "Wordnet: a lexical database for english," in *Communications of the ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [99] M. Mueller, N. Smith, and B. Ghanem, "A benchmark and simulator for uav tracking," in *ECCV*, 2016, pp. 445–461.
- [100] M. Mueller, N. Smith, and B. Ghanem, "Context-aware correlation filter tracking," in *CVPR*, 2017.
- [101] 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 *ECCV*, 2018.
- [102] J. Nocedal and S. Wright, *Numerical optimization*. Springer Science & Business Media, 2006.
- [103] E. Park and A. C. Berg, "Meta-tracker: Fast and robust online adaptation for visual object trackers," in *ECCV*, 2018.
- [104] Y. Qi, S. Zhang, L. Qin, H. Yao, Q. Huang, J. Lim, and M.-H. Yang, "Hedged deep tracking," in *CVPR*, 2016.
- [105] E. Real, J. Shlens, S. Mazzocchi, X. Pan, and V. Vanhoucke, "Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video," in *CVPR*, 2017.
- [106] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *CVPR*, 2016.
- [107] J. Redmon and A. Farhadi, "Yolo9000: better, faster, stronger," in *CVPR*, 2017.
- [108] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Adv. in NIPS*, 2015.
- [109] A. Robinson, F. J. Lawin, M. Danelljan, F. S. Khan, and M. Felsberg, "Learning fast and robust target models for video object segmentation," in *CVPR*, 2020.
- [110] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," in *IJCV*, vol. 115, no. 3, pp. 211–252, 2014.
- [111] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," in *IJCV*, vol. 115, no. 3, pp. 211–252, 2015.
- [112] B. Schölkopf, A. J. Smola, F. Bach *et al.*, *Learning with kernels: support vector machines, regularization, optimization, and beyond*, 2002.
- [113] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *CVPR*, 2015.
- [114] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "Cnn features off-the-shelf: an astounding baseline for recognition," in *CVPR-W*, 2014.
- [115] B. Shuai, A. Berneshawi, X. Li, D. Modolo, and J. Tighe, "Siammot: Siamese multi-object tracking," in *IEEE CVPR*, 2021.
- [116] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [117] A. W. Smeulders, D. M. Chu, R. Cucchiara, S. Calderara, A. Dehghan, and M. Shah, "Visual tracking: An experimental survey," in *T-PAMI*, vol. 36, no. 7, pp. 1442–1468, 2013.
- [118] Y. Song, C. Ma, L. Gong, J. Zhang, R. W. Lau, and M.-H. Yang, "Crest: Convolutional residual learning for visual tracking," in *CVPR*, 2017.
- [119] Y. Sui, Z. Zhang, G. Wang, Y. Tang, and L. Zhang, "Real-time visual tracking: Promoting the robustness of correlation filter learning," in *ECCV*, 2016.
- [120] C. Sun, D. Wang, H. Lu, and M.-H. Yang, "Correlation tracking via joint discrimination and reliability learning," in *CVPR*, 2018.
- [121] C. Sun, D. Wang, H. Lu, and M.-H. Yang, "Learning spatial-aware regressions for visual tracking," in *CVPR*, 2018.
- [122] X. Sun, N.-M. Cheung, H. Yao, and Y. Guo, "Non-rigid object tracking via deformable patches using shape-preserved kcf and level sets," in *ICCV*, 2017.
- [123] Y. Sun, C. Sun, D. Wang, Y. He, and H. Lu, "Roi pooled correlation filters for visual tracking," in *CVPR*, 2019.
- [124] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *CVPR*, 2015.
- [125] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *CVPR*, 2014.
- [126] M. Tang and J. Feng, "Multi-kernel correlation filter for visual tracking," in *ICCV*, 2015.
- [127] R. Tao, E. Gavves, and A. W. Smeulders, "Siamese instance search for tracking," in *CVPR*, 2016.
- [128] J. Valmadre, L. Bertinetto, J. Henriques, A. Vedaldi, and P. H. Torr, "End-to-end representation learning for correlation filter based tracking," in *CVPR*, 2017.
- [129] J. Van De Weijer, C. Schmid, J. Verbeek, and D. Larlus, "Learning color names for real-world applications," in *T-IP*, vol. 18, no. 7, pp. 1512–1523, 2009.
- [130] P. Voigtlaender, J. Luiten, P. H. Torr, and B. Leibe, "Siam r-cnn: Visual tracking by re-detection," in *CVPR*, 2020.
- [131] G. Wang, C. Luo, Z. Xiong, and W. Zeng, "Spm-tracker: Series-parallel matching for real-time visual object tracking," in *CVPR*, 2019.
- [132] M. Wang, Y. Liu, and Z. Huang, "Large margin object tracking with circulant feature maps," in *CVPR*, 2017.
- [133] N. Wang, Y. Song, C. Ma, W. Zhou, W. Liu, and H. Li, "Unsupervised deep tracking," in *CVPR*, 2019.
- [134] N. Wang, W. Zhou, Y. Song, C. Ma, W. Liu, and H. Li, "Unsupervised deep representation learning for real-time tracking," in *IJCV*, pp. 1–19, 2020.
- [135] N. Wang, W. Zhou, Q. Tian, R. Hong, M. Wang, and H. Li, "Multi-cue correlation filters for robust visual tracking," in *CVPR*, 2018.
- [136] N. Wang, W. Zhou, J. Wang, and H. Li, "Transformer meets tracker: Exploiting temporal context for robust visual tracking," in *arXiv preprint arXiv:2103.11681*, 2021.
- [137] Q. Wang, Z. Teng, J. Xing, J. Gao, W. Hu, and S. Maybank, "Learning attentions: residual attentional siamese network for high performance online visual tracking," in *CVPR*, 2018.
- [138] Q. Wang, L. Zhang, L. Bertinetto, W. Hu, and P. H. Torr, "Fast online object tracking and segmentation: A unifying approach," in *CVPR*, 2019.
- [139] K. Q. Weinberger and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," in *JMLR*, vol. 10, no. 2, 2009.
- [140] Y. Wu, J. Lim, and M. Yang, "Object tracking benchmark," in *T-PAMI*, vol. 37, no. 9, pp. 1834–1848, 2015.
- [141] Y. Wu and Q. Ji, "Facial landmark detection: A literature survey," *IJCV*, vol. 127, no. 2, pp. 115–142, 2019.
- [142] N. Xu, L. Yang, Y. Fan, J. Yang, D. Yue, Y. Liang, B. Price, S. Cohen, and T. Huang, "Youtube-vos: Sequence-to-sequence video object segmentation," in *ECCV*, 2018.

- [143] T. Xu, Z.-H. Feng, X.-J. Wu, and J. Kittler, "Joint group feature selection and discriminative filter learning for robust visual object tracking," in *ICCV*, 2019.
- [144] Y. Xu, A. Osep, Y. Ban, R. Horaud, L. Leal-Taixé, and X. Alameda-Pineda, "How to train your deep multi-object tracker," in *IEEE CVPR*, 2020.
- [145] Y. Xu, Z. Wang, Z. Li, Y. Yuan, and G. Yu, "Siamfc++: Towards robust and accurate visual tracking with target estimation guidelines," in *AAAI*, 2020.
- [146] B. Yan, H. Peng, J. Fu, D. Wang, and H. Lu, "Learning spatio-temporal transformer for visual tracking," in *arXiv preprint arXiv:2103.17154*, 2021.
- [147] B. Yan, H. Peng, K. Wu, D. Wang, J. Fu, and H. Lu, "Lighttrack: Finding lightweight neural networks for object tracking via one-shot architecture search," in *IEEE CVPR*, 2021.
- [148] B. Yan, D. Wang, H. Lu, and X. Yang, "Cooling-shrinking attack: Blinding the tracker with imperceptible noises," in *CVPR*, 2020.
- [149] B. Yan, X. Zhang, D. Wang, H. Lu, and X. Yang, "Alpha-refine: Boosting tracking performance by precise bounding box estimation," in *CVPR*, 2021.
- [150] B. Yan, H. Zhao, D. Wang, H. Lu, and X. Yang, "'skimming-perusal' tracking: A framework for real-time and robust long-term tracking," in *ICCV*, 2019.
- [151] T. Yang and A. B. Chan, "Learning dynamic memory networks for object tracking," in *ECCV*, 2018.
- [152] T. Yang and A. B. Chan, "Visual tracking via dynamic memory networks," *IEEE TPAMI*, vol. 43, no. 1, pp. 360–374, 2019.
- [153] T. Yang, P. Xu, R. Hu, H. Chai, and A. B. Chan, "Roam: Recurrently optimizing tracking model," in *CVPR*, 2020.
- [154] Y. Yao, X. Wu, L. Zhang, S. Shan, and W. Zuo, "Joint representation and truncated inference learning for correlation filter based tracking," in *ECCV*, 2018.
- [155] A. Yilmaz, O. Javed, and M. Shah, "Object tracking: A survey," in *CSUR*, vol. 38, no. 4, pp. 13–es, 2006.
- [156] J. Yin, W. Wang, Q. Meng, R. Yang, and J. Shen, "A unified object motion and affinity model for online multi-object tracking," in *IEEE CVPR*, 2020.
- [157] J. Yu, Y. Jiang, Z. Wang, Z. Cao, and T. Huang, "Unitbox: An advanced object detection network," in *ACM Multimedia*, 2016.
- [158] Y. Yu, Y. Xiong, W. Huang, and M. R. Scott, "Deformable siamese attention networks for visual object tracking," in *CVPR*, 2020.
- [159] J. Zbontar and Y. LeCun, "Computing the stereo matching cost with a convolutional neural network," in *CVPR*, 2015.
- [160] K. Zhang, L. Zhang, Q. Liu, D. Zhang, and M.-H. Yang, "Fast visual tracking via dense spatio-temporal context learning," in *ECCV*, 2014.
- [161] L. Zhang, A. Gonzalez-Garcia, J. v. d. Weijer, M. Danelljan, and F. S. Khan, "Learning the model update for siamese trackers," in *ICCV*, 2019.
- [162] M. Zhang, Q. Wang, J. Xing, J. Gao, P. Peng, W. Hu, and S. Maybank, "Visual tracking via spatially aligned correlation filters network," in *ECCV*, 2018.
- [163] S. Zhang, H. Yao, X. Sun, and X. Lu, "Sparse coding based visual tracking: Review and experimental comparison," in *PR*, vol. 46, no. 7, pp. 1772–1788, 2013.
- [164] T. Zhang, C. Xu, and M.-H. Yang, "Multi-task correlation particle filter for robust object tracking," in *CVPR*, 2017.
- [165] T. Zhang, C. Xu, and M.-H. Yang, "Learning multi-task correlation particle filters for visual tracking," in *T-PAMI*, vol. 41, no. 2, pp. 365–378, 2018.
- [166] Y. Zhang, L. Wang, J. Qi, D. Wang, M. Feng, and H. Lu, "Structured siamese network for real-time visual tracking," in *ECCV*, 2018.
- [167] Z. Zhang and H. Peng, "Deeper and wider siamese networks for real-time visual tracking," in *CVPR*, 2019.
- [168] Z. Zhang, H. Peng, J. Fu, B. Li, and W. Hu, "Ocean: Object-aware anchor-free tracking," in *ECCV*, 2020.
- [169] L. Zheng, M. Tang, Y. Chen, J. Wang, and H. Lu, "Fast-deepkcf without boundary effect," in *ICCV*, 2019.
- [170] L. Zheng, M. Tang, Y. Chen, G. Zhu, J. Wang, and H. Lu, "Improving multiple object tracking with single object tracking," in *IEEE CVPR*, 2021.
- [171] T. Zhou, J. Li, X. Li, and L. Shao, "Target-aware object discovery and association for unsupervised video multi-object segmentation," in *IEEE CVPR*, 2021.
- [172] C. Zhu, Y. He, and M. Savvides, "Feature selective anchor-free module for single-shot object detection," in *CVPR*, 2019.

- [173] J. Zhu, H. Yang, N. Liu, M. Kim, W. Zhang, and M.-H. Yang, "Online multi-object tracking with dual matching attention networks," in *ECCV*, 2018.
- [174] Z. Zhu, Q. Wang, B. Li, W. Wu, J. Yan, and W. Hu, "Distractor-aware siamese networks for visual object tracking," in *ECCV*, 2018.
- [175] Z. Zhu, W. Wu, W. Zou, and J. Yan, "End-to-end flow correlation tracking with spatial-temporal attention," in *CVPR*, 2018.



Sajid Javed is a faculty member at Khalifa University (KU), UAE. Prior to that, he was a research fellow at KU from 2019 to 2021 and at University of Warwick, U.K, from 2017-2018. He received his B.Sc. degree in computer science from University of Hertfordshire, U.K, in 2010. He completed his combined Master's and Ph.D. degree in computer science from Kyungpook National University, Republic of Korea, in 2017.



Martin Danelljan is a group leader and lecturer at ETH Zürich, Switzerland. He received his Ph.D. degree from Linköping University, Sweden in 2018. His Ph.D. thesis was awarded the biennial Best Nordic Thesis Prize at SCIA 2019. His main research interests are meta and online learning, deep probabilistic models, and conditional generative models. His research includes applications to visual tracking, video object segmentation, dense correspondence estimation.



Fahad Khan is a faculty member at MBZUAI, UAE and Linköping University, Sweden. He received the M.Sc. degree in Intelligent Systems Design from Chalmers University of Technology, Sweden and a Ph.D. degree in Computer Vision from Autonomous University of Barcelona, Spain. His research interests include a wide range of topics within computer vision, including object recognition, detection, segmentation, and tracking.



Muhammad Haris Khan is a faculty member at the MBZUAI, UAE. Prior to MBZUAI, he was Research Scientist at the Inception Institute of Artificial Intelligence, UAE. He obtained his PhD in Computer Vision from University of Nottingham, UK. He has published several papers in top computer vision venues. His research interests span active topics in computer vision



Michael Felsberg is a full professor at Linköping University, Sweden. He holds a Ph.D. degree from Kiel University, Germany (2002) and a docent degree from Linköping University (2005). His research interests include, besides visual object tracking, video object and instance segmentation, point cloud processing, and efficient machine learning techniques.



Jiri Matas is a full professor at Czech Technical University in Prague. He holds a PhD degree from the University of Surrey, UK (1995). He has published more than 250 papers in refereed journals and conferences. His publications have about 53000 citations registered in Google Scholar; his h-index is 85. J. Matas has served in various roles at major international computer vision conferences. He is an EiC of IJCV and was an AE of IEEE T. PAMI. His research interests include visual tracking, object recognition, image

matching and retrieval, sequential pattern recognition, and RANSAC-type optimization methods.