

Ocean: Object-aware Anchor-free Tracking

Zhipeng Zhang^{1, 2} * Houwen Peng² † Jianlong Fu² Bing Li¹ and Weiming Hu¹

¹CASIA & AI School, UCAS & CEBSIT ²Microsoft Research

Abstract. Anchor-based Siamese trackers have achieved remarkable advancements in accuracy, yet the further improvement is restricted by the lagged tracking robustness. We find the underlying reason is that the regression network in anchor-based methods is only trained on the positive anchor boxes (*i.e.*, $IoU \geq 0.6$). This mechanism makes it difficult to refine the anchors whose overlap with the target objects are small. In this paper, we propose a novel object-aware anchor-free network to address this issue. First, instead of refining the reference anchor boxes, we directly predict the position and scale of target objects in an anchor-free fashion. Since each pixel in groundtruth boxes is well trained, the tracker is capable of rectifying inexact predictions of target objects during inference. Second, we introduce a feature alignment module to learn an object-aware feature from predicted bounding boxes. The object-aware feature can further contribute to the classification of target objects and background. Moreover, we present a novel tracking framework based on the anchor-free model. The experiments show that our anchor-free tracker achieves state-of-the-art performance on five benchmarks, including VOT-2018, VOT-2019, OTB-100, GOT-10k and LaSOT. The source code is available at <https://github.com/researchmm/TrackIt>.

Keywords: Visual tracking, Anchor-free, Object-aware


1 Introduction

Object tracking is a fundamental vision task. It aims to infer the location of an arbitrary target in a video sequence, given only its location in the first frame. The main challenge of tracking lies in that the target objects may undergo heavy occlusions, large deformation and illumination variations [44,49]. Tracking at real-time speeds has a variety of applications, such as surveillance, robotics, autonomous driving and human-computer interaction [16,25,33].

In recent years, Siamese tracker has drawn great attention because of its balanced speed and accuracy. The seminal works, *i.e.*, SINT [35] and SiamFC [1], employ Siamese networks to learn a similarity metric between the object target and candidate image patches, thus modeling the tracking as a search problem

* Work performed when Zhipeng was an intern of Microsoft Research. † Corresponding author. Z. Zhang, B. Li, W. Hu are with the Institution of Automation, Chinese Academy of Sciences (CASIA) and School of Artificial Intelligence, University of Chinese Academy of Sciences (UCAS) and CAS Center for Excellence in Brain Science and Intelligence Technology (CEBSIT). Email: {zhipeng.zhang,wmhu,bli}@nlpr.ia.ac.cn, {houwen.peng,jianf}@microsoft.com

Fig. 1. A comparison of the performance and speed of state-of-the-art tracking methods on VOT-2018. We visualize the Expected Average Overlap (EAO) with respect to the Frames-Per-Seconds (FPS). *Offline-1* and *Offline-2* indicate the proposed offline trackers with and without feature alignment module, respectively.



of the target over the entire image. A large amount of follow-up Siamese trackers have been proposed and achieved promising performances [9,11,21,22,50]. Among them, the Siamese region proposal networks, dubbed SiamRPN [22], is representative. It introduces region proposal networks [31], which consist of a classification network for foreground-background estimation and a regression network for anchor-box refinement, *i.e.*, learning 2D offsets to the predefined anchor boxes. This anchor-based trackers have shown tremendous potential in tracking accuracy. However, since the regression network is only trained on the positive anchor boxes (*i.e.*, $IoU \geq 0.6$), it is difficult to refine the anchors whose overlap with the target objects are small. This will cause tracking failures especially when the classification results are not reliable. For instance, due to the error accumulation in tracking, the predictions of target positions may become unreliable, *e.g.*, $IoU < 0.3$. The regression network is incapable of rectifying this weak prediction because it is previously unseen in the training set. As a consequence, the tracker gradually drifts in subsequent frames.

It is natural to throw a question: *can we design a bounding-box regressor with the capability of rectifying inaccurate predictions?* In this work, we show the answer is affirmative by proposing a novel object-aware anchor-free tracker. Instead of predicting the small offsets of anchor boxes, our object-aware anchor-free tracker directly regresses the positions of target objects in a video frame. More specifically, the proposed tracker consists of two components: an object-aware classification network and a bounding-box regression network. The classification is in charge of determining whether a region belongs to foreground or background, while the regression aims to predict the distances from each pixel within the target objects to the four sides of the groundtruth bounding boxes. Since each pixel in the groundtruth box is well trained, the regression network is able to localize the target object even when only a small region is identified as the foreground. Eventually, during inference, the tracker is capable of rectifying the weak predictions whose overlap with the target objects are small.

When the regression network predicts a more accurate bounding box (*e.g.*, rectifying weak predictions), the corresponding features can in turn help the classification of foreground and background. We use the predicted bounding box as a reference to learn an object-aware feature for classification. More concretely, we

introduce a feature alignment module, which contains a 2D spatial transformation to align the feature sampling locations with predicted bounding boxes (*i.e.*, regions of candidate objects). This module guarantees the sampling is specified within the predicted regions, accommodating to the changes of object scale and position. Consequently, the learned features are more discriminative and reliable for classification.

The effectiveness of the proposed framework is verified on five benchmarks: VOT-2018 [17], VOT-2019 [18], OTB-100 [44], GOT-10k [14] and LaSOT [8]. Our approach achieves state-of-the-art performance (an EAO of 0.467) on VOT-2018 [17], while running at 58 *fps*, as shown in Fig. 1. It obtains up to 92.2% and 12.8% relative improvements over the anchor-based methods, *i.e.*, SiamRPN [22] and SiamRPN++ [21], respectively. On other datasets, the performance of our tracker is also competitive, compared with recent state-of-the-arts. In addition, we further equip our anchor-free tracker with a plug-in online update module, and enable it to capture the appearance changes of objects during inference. The online module further enhances the tracking performance, which shows the scalability of the proposed anchor-free tracking approach.

The main contributions of this work are two-fold. 1) We propose an object-aware anchor-free network based on the observation that the anchor-based method is difficult to refine the anchors whose overlap with the target object is small. The proposed algorithm can not only rectify the imprecise bounding-box predictions, but also learn an object-aware feature to enhance the matching accuracy. 2) We design a novel tracking framework by combining the proposed anchor-free network with an efficient feature combination module. The proposed tracking model achieves state-of-the-art performance on five benchmarks while running in real-time speeds.

2 Related Work

In this section, we review the related work on anchor-free mechanism and feature alignment in both tracking and detection, as well as briefly review recent Siamese trackers.

Siamese trackers. The pioneering works, *i.e.*, SINT [35] and SiamFC [1], employ Siamese networks to offline train a similarity metric between the object target and candidate image patches. SiamRPN [22] improves it with a region proposal network, which amounts to a target-specific anchor-based detector. With the predefined anchor boxes, SiamRPN [22] can capture the scale changes of objects effectively. The follow-up studies mainly fall into two camps: designing more powerful backbone networks [21,50] or proposing more effective proposal networks [9]. Although these offline Siamese trackers have achieved very promising results, their tracking robustness is still inferior to the recent state-of-the-art online trackers, such as ATOM [4] and DiMP [2].

Anchor-free mechanism. Anchor-free approaches recently became popular in object detection tasks, because of their simplicity in architectures and superiority in performance [7,19,36]. Different from anchor-based methods which

estimate the offsets of anchor boxes, anchor-free mechanisms predict the location of objects in a direct way. The early anchor-free work [47] predicts the intersection over union with objects, while recent works focus on estimating the keypoints of objects, *e.g.*, the object center [7] and corners [19]. Another branch of anchor-free detectors [30,36] predicts the object bounding box at each pixel, without using any references, *e.g.*, anchors or keypoints. The anchor-free mechanism in our method is inspired by, but different from that in the recent detection algorithm [36]. We will discuss the key differences in Sec. 3.4.

Feature alignment. The alignment between visual features and reference ROIs (Regions of Interests) is vital for localization tasks, such as detection and tracking [40]. For example, ROIAlign [12] are commonly recruited in object detection to align the features with the reference anchor boxes, leading to remarkable improvements on localization precision. In visual tracking, there are also several approaches [15,41] considering the correspondence between visual features and candidate bounding boxes. However, these approaches only take account of the bounding boxes with high classification scores. If the high scores indicate the background regions, then the corresponding features will mislead the detection of target objects. To address this, we propose a novel feature alignment method, in which the alignment is independent of the classification results. We sample the visual features from the predicted bounding boxes directly, without considering the classification score, generating object-aware features. This object-aware features, in turn, help the classification of foreground and background.

3 Object-aware Anchor-Free Networks

This section proposes the **Object-aware anchor-free networks** (Ocean) for visual tracking. The network architecture consists of two components: an object-aware classification network for foreground-background probability prediction and a regression network for target scale estimation. The input features to these two networks are generated by a shared backbone network (elaborated in Sec. 4.1). We introduce the regression network first, followed by the classification branch, because the regression branch provides object scale information to enhance the classification of the target object and background.

3.1 Anchor-free Regression Network

Revisiting recent anchor-based trackers [21,22], we observed that the trackers drift speedily when the predicted bounding box becomes unreliable. The underlying reason is that, during training, these approaches only consider the anchor boxes whose *IoU* with groundtruth are larger than a high threshold, *i.e.*, $IoU \geq 0.6$. Hence, these approaches lack the competence to amend the weak predictions, *e.g.*, the boxes whose overlap with the target are small.

To remedy this issue, we introduce a novel anchor-free regression for visual tracking. It considers all the pixels in the groundtruth bounding box as the training samples. The core idea is to estimate the distances from each pixel

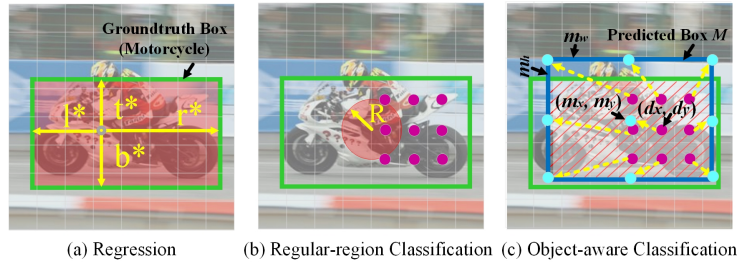


Fig. 2. (a) Regression: the pixels in groundtruth box, *i.e.* the red region, are labeled as the positive samples in training. (b) Regular-region classification: the pixels closing to the target’s center, *i.e.* the red region, are labeled as the positive samples. The purple points indicate the sampled positions of a location in the score map. (c) Object-aware classification: the IoU of predicted box and groundtruth box, *i.e.*, the region with red slash lines, is used as the label during training. The cyan points represent the sampling positions for extracting object-aware features. The yellow arrows indicate the offsets induced by spatial transformation. Best viewed in color.

within the target object to the four sides of the groundtruth bounding box. Specifically, let $B = (x_0, y_0, x_1, y_1) \in \mathbb{R}^4$ denote the top-left and bottom-right corners of the groundtruth bounding box of a target object. A pixel is considered as the regression sample if its coordinates (x, y) fall into the groundtruth box B . Hence, the labels $T^* = (l^*, t^*, r^*, b^*)$ of training samples are calculated as

$$\begin{aligned} l^* &= x - x_0, t^* = y - y_0, \\ r^* &= x_1 - x, b^* = y_1 - y, \end{aligned} \quad (1)$$

which represent the distances from the location (x, y) to the four sides of the bounding box B , as shown in Fig. 2(a). The learning of the regression network is through four 3×3 convolution layers with channel number of 256, followed by one 3×3 layer with channel number of 4 for predicting the distances. As shown in Fig. 3, the upper “Conv” block indicates the regression network.

This anchor-free regression allows for all the pixels in the groundtruth box during training, thus it can predict the scale of target objects even when only a small region is identified as foreground. Consequently, the tracker is capable of rectifying weak predictions during inference to some extent.

3.2 Object-aware Classification Network

In prior Siamese tracking approaches [1,21,22], the classification confidence is estimated by the feature sampled from a fixed regular region in the feature map, *e.g.*, the purple points in Fig. 2(b). This sampled feature depicts a fixed local region of the image, and it is not scalable to the change of object scale. As a result, the classification confidence is not reliable in distinguishing the target object from complex background.

To address this issue, we propose a feature alignment module to learn an object-aware feature for classification. The alignment module transforms the fixed sampling positions of a convolution kernel to align with the predicted bounding box. Specifically, for each location (d_x, d_y) in the classification map, it has a corresponding object bounding box $M = (m_x, m_y, m_w, m_h)$ predicted by the regression network, where m_x and m_y denote the box center while m_w and m_h represent its width and height. Our goal is to estimate the classification confidence for each location (d_x, d_y) by sampling features from the corresponding candidate region M . The standard 2D convolution with kernel size of $k \times k$ samples features using a fixed regular grid $\mathcal{G} = \{(-\lfloor k/2 \rfloor, -\lfloor k/2 \rfloor), \dots, (\lfloor k/2 \rfloor, \lfloor k/2 \rfloor)\}$, where $\lfloor \cdot \rfloor$ denotes the floor function. The regular grid \mathcal{G} cannot guarantee the sampled features cover the whole content of region M .

Therefore, we propose to equip the regular sampling grid \mathcal{G} with a spatial transformation \mathcal{T} to convert the sampling positions from the fixed region to the predicted region M . As shown in Fig. 2(c), the transformation \mathcal{T} (the dashed yellow arrows) is obtained by measuring the relative direction and distance from the sampling positions in \mathcal{G} (the purple points) to the positions aligned with the predicted bounding box (the cyan points). With the new sampling positions, the object-aware feature is extracted by the feature alignment module, which is formulated as

$$\mathbf{f}[u] = \sum_{g \in \mathcal{G}, \Delta t \in \mathcal{T}} \mathbf{w}[g] \cdot \mathbf{x}[u + g + \Delta t], \quad (2)$$

where \mathbf{x} represents the input feature map, \mathbf{w} denotes the learned convolution weight, u indicates a location on the feature map, and \mathbf{f} represents the output object-aware feature map. The spatial transformation $\Delta t \in \mathcal{T}$ represents the distance vector from the original regular sampling points to the new points aligned with the predicted bounding box. The transformation is defined as

$$\mathcal{T} = \{(m_x, m_y) + \mathcal{B}\} - \{(d_x, d_y) + \mathcal{G}\}, \quad (3)$$

where $\{(m_x, m_y) + \mathcal{B}\}$ represents the sampling positions aligned with M , *e.g.*, the cyan points in Fig. 2(c), $\{(d_x, d_y) + \mathcal{G}\}$ indicates the regular sampling positions used in standard convolution, *e.g.*, the purple points in Fig. 2(c), and $\mathcal{B} = \{(-m_w/2, -m_h/2), \dots, (m_w/2, m_h/2)\}$ denotes the coordinates of the new sampling positions (*e.g.*, the cyan points in Fig. 2(c)) relative to the box center (*e.g.*, (m_x, m_y)). It is worth noting that when the transformation $\Delta t \in \mathcal{T}$ is set to 0 in Eq. (2), the feature sampling mechanism is degenerated to the fixed sampling on regular points, generating the regular-region feature. The transformations of the sampling positions are adaptive to the variations of the predicted bounding boxes in video frames. Thus, the extracted object-aware feature is robust to the changes of object scale, which is beneficial for feature matching during tracking. Moreover, the object-aware feature provides a global description of the candidate targets, which enables the distinguish of the object and background to be more reliable.

We exploit both the object-aware feature and the regular-region feature to predict whether a region belongs to target object or image background. For the classification based upon the object-aware feature, we apply a standard convolution with kernel size of 3×3 over \mathbf{f} to predict the confidence p_o (visualized as the ‘‘OA.Conv’’ block of the classification network in Fig. 3). For the classification based on the regular-region feature, four 3×3 standard convolution layers with channel number of 256, followed by one standard 3×3 layer with channel number of one are performed over the regular-region feature \mathbf{f}' to predict the confidence p_r (visualized as the ‘‘Conv’’ block of the classification network in Fig. 3). Calculating the summation of the confidence p_o and p_r obtains the final classification score. The object-aware feature provides a global description of the target, thus enhancing the matching accuracy of candidate regions. Meanwhile, the regular-region feature concentrates on local parts of images, which is robust to localize the center of target objects. The combination of the two features improves the reliability of the classification network.

3.3 Loss Function

To optimize the proposed anchor-free networks, we employ *IoU* loss [47] and binary cross-entropy (BCE) loss [6] to train the regression and classification networks jointly. In regression, the loss is defined as

$$\mathcal{L}_{reg} = - \sum_i \ln(IoU(p_{reg}, T^*)), \quad (4)$$

where p_{reg} denotes the prediction, and i indexes the training samples. In classification, the loss \mathcal{L}_o based upon the object-aware feature \mathbf{f} is formulated as

$$\mathcal{L}_o = - \sum_j p_o^* \log(p_o) + (1 - p_o^*) \log(1 - p_o), \quad (5)$$

while the loss \mathcal{L}_r based upon the regular-region feature \mathbf{f}' is formulated as

$$\mathcal{L}_r = - \sum_j p_r^* \log(p_r) + (1 - p_r^*) \log(1 - p_r), \quad (6)$$

where p_o and p_r are the classification score maps computed over the object-aware feature and regular-region feature respectively, j indexes the training samples for classification, and p_o^* and p_r^* denote the groundtruth labels. More concretely, p_o^* is a probabilistic label, in which each value indicates the *IoU* between the predicted bounding box and groundtruth, *i.e.*, the region with red slash lines in Fig. 2(c). p_r^* is a binary label, where the pixels closing to the center of the target are labeled as 1, *i.e.*, the red region in Fig. 2(b), which is formulated as

$$p_r^*[v] = \begin{cases} 1, & \text{if } \|v - c\| \leq R, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

The joint training of the entire object-aware anchor-free networks is to optimize the following objective function:

$$\mathcal{L} = \mathcal{L}_{reg} + \lambda_1 \mathcal{L}_o + \lambda_2 \mathcal{L}_r, \quad (8)$$

where λ_1 and λ_2 are the tradeoff hyperparameters.

3.4 Relation to Prior Anchor-Free Work

Our anchor-free mechanism shares similar spirit with recent detection methods [7,19,36] (discussed in Sec. 2). In this section, we further discuss the differences to the most related work, *i.e.*, FCOS [36]. Both FCOS and our method predict the object locations directly on the image plane at pixel level. However, our work differs from FCOS [36] in two fundamental ways. 1) In FCOS [36], the training samples for the classification and regression networks are identical. Both are sampled from the positions within the groundtruth boxes. Differently, in our method, the data sampling strategies for classification and regression are asymmetric which is tailored for tracking tasks. More specifically, the classification network only considers the pixels closing to the target as positive samples (*i.e.*, $R \leq 16$ pixels), while the regression network considers all the pixels in the ground-truth box as training samples. This fine-grained sampling strategy guarantees the classification network can learn a robust similarity metric for region matching, which is important for tracking. 2) In FCOS [36], the objectness score is calculated with the feature extracted from a fixed regular-region, similar to the purple points in Fig. 2(b). By contrast, our method additionally introduce an object-aware feature, which captures the global appearance of target objects. The object-aware feature aligns the sampling regions with the predicted bounding box (*e.g.*, cyan points in Fig. 2(c)), thus it is adaptive to the scale change of objects. The combination of the regular-region feature and the object-aware feature allows the classification to be more reliable, as verified in Sec. 5.3.

4 Object-aware Anchor-Free Tracking

This section depicts the tracking algorithm building upon the proposed object-aware anchor-free networks (Ocean). It contains two parts: an offline anchor-free model and an online update model, as illustrated in Fig. 3.

4.1 Framework

The offline tracking is built on the object-aware anchor-free networks, consisting of three steps: feature extraction, combination and target localization.

Feature extraction. Following the architecture of Siamese tracker [1], our approach takes an image pair as input, *i.e.*, an exemplar image and a candidate search image. The exemplar image represents the object of interest, *i.e.*, an image patch centered on the target object in the first frame, while the search image is typically larger and represents the search area in subsequent video frames. Both inputs are processed by a modified ResNet-50 [13] backbone and then yield two feature maps. More specifically, we cut off the last stage of the standard ResNet-50 [13], and only retain the first fourth stages as the backbone. The first three stages share the same structure as the original ResNet-50. In the fourth stage, the convolution stride of down-sampling unit [13] is modified from 2 to 1 to increase the spatial size of feature maps, meanwhile, all the 3×3 convolutions

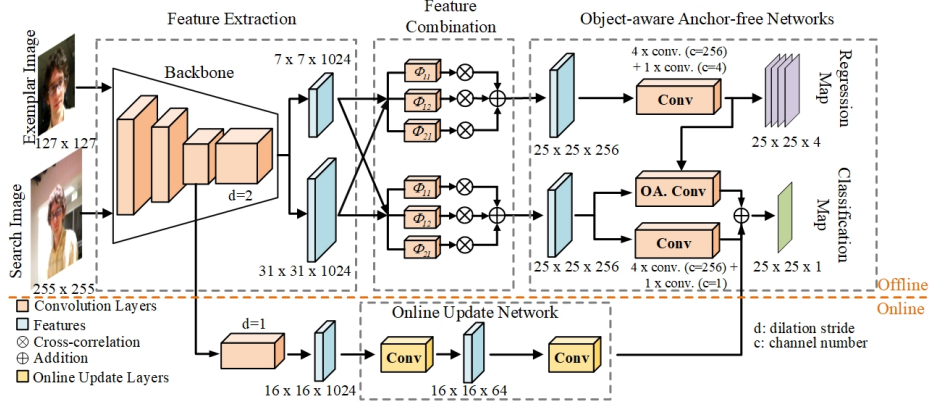


Fig. 3. Overview of the proposed tracking framework, consisting of an offline anchor-free part (top) and an online model update part (bottom). The offline tracking includes feature extraction, feature combination and target localization with object-aware anchor-free networks, as elaborated in Sec. 4.1. The plug-in online update network models the appearance changes of target objects, as detailed in Sec. 4.2. Φ_{ab} indicates a 3×3 convolution layer with dilation stride of a along the X -axis and b along the Y -axis.

are augmented with a dilation with stride of 2 to increase the receptive fields. These modifications increase the resolution of output features, thus improving the feature capability on object localization [3,21].

Feature combination. This step exploits a depth-wise cross-correlation operation [21] to combine the extracted features of the exemplar and search images, and generates the corresponding similarity features for the subsequent target localization. Different from the previous works performing the cross-correlation on multi-scale features [21], our method only performs over a single scale, *i.e.*, the last stage of the backbone. We pass the single-scale features through three parallel dilated convolution layers [48], and then fuse the correlation features through point-wise summation, as presented in Fig. 3 (feature combination).

For concreteness, the feature combination process can be formulated as

$$S = \sum_{ab} \Phi_{ab}(\mathbf{f}_e) * \Phi_{ab}(\mathbf{f}_s) \quad (9)$$

where \mathbf{f}_e and \mathbf{f}_s represent the features of the exemplar and search images respectively, Φ_{ab} indicates a single dilated convolution layer, and $*$ denotes the cross-correlation operation [1]. The kernel size of the dilated convolution Φ_{ab} is set to 3×3 , while the dilation strides are set to a along the X -axis and b along the Y -axis. Φ_{ab} also reduces the feature channels from 1024 to 256 to save computation cost. In experiments, we found that increasing the diversity of dilations can improve the representability of features, thereby we empirically choose three different dilations, whose strides are set to $(a, b) \in \{(1, 1), (1, 2), (2, 1)\}$. The

convolutions with different dilations can capture the features of regions with different scales, improving the scale invariance of the final combined features.

Target localization. This step employs the proposed object-aware anchor-free networks to localize the target from search images. The probabilities p_o and p_r predicted by the classification network are averaged with a weight ω as

$$p_{cls} = \omega p_o + (1 - \omega) p_r. \quad (10)$$

Similar to [1,21], we impose a penalty on scale change to suppress the large variation of object size and aspect ratio. We provide more details in the supplementary materials.

4.2 Integrating Online Update

We further equip the offline algorithm with an online update model. Inspired by [2,4], we introduce an online branch to capture the appearance changes of target object during tracking. As shown in Fig. 3 (bottom part), the online branch inherits the structure and parameters from the first three stages of the backbone network, *i.e.*, modified ResNet-50 [13]. The fourth stage keep the same structure as the backbone, but its initial parameters are obtained through the pretraining strategy proposed in [2]. For model update, we employ the fast conjugate gradient algorithm [2] to train online branch during inference. The foreground score maps estimated by the online branch and the classification branch are weighted as

$$p = \omega' p_{onl} + (1 - \omega') \hat{p}_{cls}, \quad (11)$$

where ω' represents the weights between the classification score \hat{p}_{cls} and the online estimation score p_{onl} . Note that the IoUNet in [2,4] is not used in our model. We refer readers to [2,4] for more details.

5 Experiments

This section presents the results of our Ocean tracker on five tracking benchmark datasets, with comparisons to the state-of-the-art algorithms. Experimental analysis is provided to evaluate the effects of each component in our model.

5.1 Implementation Details

Training. The backbone network is initialized with the parameters pre-trained on ImageNet [32]. The proposed trackers are trained on the datasets of Youtube-BB [29], ImageNet VID [32], ImageNet DET [32], GOT-10k [14] and COCO [26]. The size of input exemplar image is 127×127 pixels, while the search image is 255×255 pixels. We use synchronized SGD [20] on 8 GPUs, with each GPU hosting 32 images, hence the mini-batch size is 256 images per iteration. There are 50 epochs in total. Each epoch uses 6×10^5 training pairs. For the first 5 epochs, we start with a warmup learning rate of 10^{-3} to train the object-aware anchor-free networks, while freezing the parameters of the backbone. For

the remaining epochs, the backbone network is unfrozen, and the whole network is trained end-to-end with a learning rate exponentially decayed from 5×10^{-3} to 10^{-5} . The weight decay and momentum are set to 10^{-3} and 0.9, respectively. The threshold R of the classification label in Eq. (7) is set to 16 pixels. The weight parameters λ_1 and λ_2 in Eq. (8) are set to 1 and 1.2, respectively.

We noticed that the training settings (data selection, iterations, *etc.*) are often different in recent trackers, *e.g.*, SiamRPN [22], SiamRPN++ [21], ATOM [4] and DiMP [4]. It is difficult to compare different models under a unified training schedule. But for a fair comparison, we additionally evaluate our method and SiamRPN++ [21] under the same training setting, as discussed in Sec. 5.3.

Testing. For the offline model, tracking follows the same protocols as in [1,22]. The feature of the target object is computed once at the first frame, and then is continuously matched to subsequent search images. The fusion weight ω of the object-aware classification score in Eq. (10) is set to 0.07, while the weight ω' in Eq. (11) is set to 0.5. These hyper-parameters in testing are selected with the tracking toolkit [50], which contains an automated parameter tuning algorithm. Our trackers are implemented using Python 3.6 and PyTorch 1.1.0. The experiments are conducted on a server with 8 Tesla V100 GPUs and a Xeon E5-2690 2.60GHz CPU. Note that we run the proposed tracker three times, the standard deviation of the performance is $\pm 0.5\%$, demonstrating the stability of our model. We report the average performance of the three-time runs in the following comparisons.

Evaluation datasets and metrics. We use five benchmark datasets including VOT-2018 [17], VOT-2019 [18], OTB-100 [44], GOT-10k [14] and LaSOT [8] for tracking performance evaluation. In particular, VOT-2018 [17] contains 60 sequences. VOT-2019 [18] is developed by replacing the 20% least challenging videos in VOT-2018 [17]. We adopt the Expected Average Overlap (EAO) [18] which takes both accuracy (A) and robustness (R) into account to evaluate overall performance. The standardized OTB-100 [44] benchmark consists of 100 videos. Two metrics, *i.e.*, precision (Prec.) and area under curve (AUC) are used to rank the trackers. GOT-10k [14] is a large-scale dataset containing over 10 thousand videos. The trackers are evaluated using an online server on a test set of 180 videos. It employs the widely used average overlap (AO) and success rate (SR) as performance indicators. Compared to these benchmark datasets, LaSOT [8] has longer sequences, with an average of 2,500 frames per sequence. Success (SUC) and precision (Prec.) are used to evaluate tracking performance.

5.2 State-of-the-art Comparison

To extensively evaluate the proposed method, we compare it with 22 state-of-the-art trackers, which cover most of current representative methods. There are 9 anchor-based Siamese framework based methods (SiamFC [1], GradNet [24], DSiam [11], MemDTC [46], SiamRPN [22], C-RPN [9], SiamMASK [43], SiamRPN++ [21] and SiamRCNN [39]), 8 discriminative correlation filter based methods (CFNet [38], ECO [5], STRCF [23], LADCF [45], UDT [42], STN [37], ATOM [4] and DiMP [2]), 3 multi-domain learning based methods (MDNet [27],

	CRPN [9]	ECO [5]	STRCF [23]	LADCF [45]	ATOM [4]	SRCNN [39]	SiamRPN++ [21]	DiMP [2]	Ocean offline	Ocean online
EAO \uparrow	0.273	0.280	0.345	0.389	0.401	0.408	0.414	0.440	0.467	0.489
A \uparrow	0.550	0.484	0.523	0.503	0.590	0.609	0.600	0.597	0.598	0.592
R \downarrow	0.320	0.276	0.215	0.159	0.204	0.220	0.234	0.153	0.169	0.117

Table 1. Performance comparisons on VOT-2018 benchmark. Red, green and blue fonts indicate the top-3 trackers. “Ocean” denotes our propose model.

	MemDTC [46]	SiamMASK [43]	SiamRPN++ [21]	ATOM [4]	STN [37]	DiMP ^r [2]	DiMP ^b [2]	Ocean offline	Ocean online
EAO \uparrow	0.228	0.287	0.292	0.301	0.314	0.321	0.379	0.327	0.350
A \uparrow	0.485	0.594	0.580	0.603	0.589	0.582	0.594	0.590	0.594
R \downarrow	0.587	0.461	0.446	0.411	0.349	0.371	0.278	0.376	0.316

Table 2. Performance comparisons on VOT-2019. The “DiMP^r” and “DiMP^b” indicate realtime and baseline performances of DiMP, as reported in [18].

	CFNet [38]	MDNet [27]	SiamFC [1]	ECO [5]	DSiam [11]	SiamRPN++ [21]	ATOM [4]	DiMP [2]	Ocean offline	Ocean online
AO \uparrow	0.261	0.299	0.392	0.395	0.417	0.518	0.556	0.611	0.592	0.611
SR _{0.5} \uparrow	0.243	0.303	0.406	0.407	0.461	0.618	0.634	0.712	0.695	0.721

Table 3. Performance comparisons on GOT-10k test set.

RT-MDNet [15] and VITAL [34], 1 graph network based method (GCT [10]) and 1 meta-learning based tracker (MetaCREST [28]). The results are summarized in Tab. 1 - 3 and Fig. 4.

VOT-2018. The evaluation on VOT-2018 is performed by the official toolkit [17]. As shown in Tab. 1, Our offline Ocean tracker outperforms the champion method of VOT-2018, *i.e.*, (LADCF [45]), by 7.8 points. Compared to the state-of-the-art offline tracker SiamRPN++ [21], our offline model achieves EAO improvements of 5.3 points, while running faster, as shown in Fig. 1. It is worth noting that the improvements mainly come from the robustness score, which obtains 27.8% relative increases over SiamRPN++. Moreover, our offline model is superior to the recent online trackers ATOM [4] and DiMP [2]. The online augmented model further improves our tracker by 2.2 points in terms of EAO.

VOT-2019. Tab. 2 reports the evaluation results with the comparisons to recent prevailing trackers on VOT-2019. We can see that the recent proposed DiMP [2] achieves the best performance, while our method ranks second. However, in real-time testing scenarios, our offline Ocean tracker achieves the best performance, surpassing DiMP^r by 0.6 points in terms of EAO. Moreover, the EAO of our offline model surpasses SiamRPN++ [21] by 3.5 points.

GOT-10k. The evaluation on GOT-10k follows the protocols in [14]. The proposed offline Ocean tracker model achieves the state-of-the-art AO score of 0.592, outperforming SiamRPN++ [21], as shown in Tab. 3. Our online model improves the AO by 4.5 points over ATOM [4], while outperforming DiMP [2] by 0.9 points in terms of success rate.

OTB-100. The last evaluation in short-term tracking is performed on the classical OTB-100 benchmark. As reported in Fig. 4, among the compared methods,

our online tracker achieves the best precision score of 0.920, while DiMP [21] achieves best AUC score of 0.686.

LaSOT. To further evaluate the proposed models, we report the results on LaSOT, which is larger and more challenging than previous benchmarks. The results on the 280 videos test set is presented in Fig. 4. Our offline Ocean tracker achieves SUC score of 0.527, outperforming SiamRPN++ with score of 0.496. Compared to ATOM [4], our online tracker improves the SUC score by 4.6 points, giving comparable results to top-ranked tracker DiMP-50 [2]. Moreover, the proposed online tracker achieves the best precision score of 0.566.

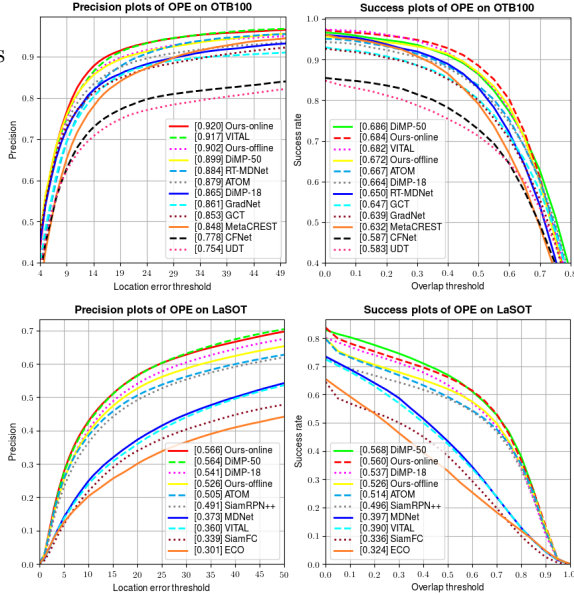


Fig. 4. Success and precision plots on OTB-100 [44] (top) and LaSOT [8] (bottom).

5.3 Analysis of the Proposed Method

Component-wise analysis. To verify the efficacy of the proposed method, we perform a component-wise analysis on the VOT-2018 benchmark, as presented in Tab. 4. The baseline model consists of a backbone network (detailed in Sec. 4.1), an anchor-free regression network (detailed in Sec. 3.1) and a classification network using regular-region feature (detailed in Sec. 3.2). In the training of baseline model, all pixels in the groundtruth box are considered as positive samples. The baseline model obtains an EAO of 0.358. In ②, the “centralized sampling” indicates that we only consider the pixels closing to the target’s center as positive samples in the training of classification (formulated as Eq. (7)). It brings significant gains, *i.e.*, 3.8 points on EAO (② *vs.* ①). This verifies that the sampling helps to learn a robust similarity metric for region matching. Adding the feature combination module (detailed in Sec. 4.1) can bring a large improvement of 4.2 points in terms of EAO (③ *vs.* ②). This demonstrates the effectiveness of the proposed irregular dilated convolution module. It introduces a multi-scale modeling of target objects, without increasing much computation overhead. Furthermore, the object-aware classification (detailed in Sec. 3.2) can also bring an improvement of 2.9 points in terms of EAO (④ *vs.* ③). This shows that the object-aware features generated by the proposed feature alignment module contribute significantly to the tracker. Finally, the tracker equipped with the

#Num	Components	EAO	#Num	Settings	EAO
①	baseline	0.358	①	SiamRPN++ [21]	0.414
②	+ centralized sampling	0.396	②	Ocean tracker (ours)	0.455
③	+ feature combination	0.438	③	+ GOT-10k in training	0.467
④	+ object-aware classification	0.467	④	+ LaSOT in training	0.462
⑤	+ online update	0.489			

Table 4. Component-wise analysis.**Table 5.** Analysis of training settings on VOT2018.

plug-in online update module (detailed in Sec. 4.2) yields another improvement of 2.2 points (⑤ *vs.* ④), showing the scalability of the proposed framework.

Training setting. We conduct another ablation study to evaluate the impact of training settings. For a fair comparison, we follow the same setting as the well-performing SiamRPN++ [21], *i.e.* training on YTB [29], VID [32], DET [32] and COCO [26] datasets for 20 epochs and using 6×10^5 image pairs in each epoch. As the results presented in Tab. 5 (② *v.s.* ①), our model surpasses SiamRPN++ by 4.1 points in terms of EAO under the same training settings. Moreover, we further add GOT-10k [14] images into training, and observe that it brings an improvement of 1.2 points (③ *v.s.* ②). This demonstrates that the main performance gains are induced by the proposed model. If we continue to add LaSOT [8] into training, the performance cannot improve further (④ *v.s.* ③). One possible reason is that the object categories in LaSOT [8] have been covered by other datasets, thus it cannot further elevate model capacities.

We provide the ablation experiments on feature combination module and alignment module in supplementary material due to space limit.

6 Conclusion

In this work, we propose a novel object-aware anchor-free tracking framework (Ocean) based upon the observation that the anchor-based method is difficult to refine the anchors whose overlap with the target objects are small. Our model directly regresses the positions of target objects in a video frame instead of predicting offsets for predefined anchors. Moreover, the learned object-aware feature by the alignment module provides a global description of the target, contributing to the reliable matching of objects. The experiments demonstrate that the proposed tracker achieves state-of-the-art performance on five benchmark datasets. In the future work, we will study applying our framework to other online video tasks, *e.g.*, video object detection and segmentation.

Acknowledgement. This work is supported by the NSFC (Grant No. U1636218, and Grant No. 61751212, 61721004, 61902401, 61972071), NSF of Beijing (Grant No. L172051), the KRPfS, CAS (Grant No. QYZDJ-SSW-JSC040), and the NSF of Guangdong (No. 2018B030311046).

References

1. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: ECCV. pp. 850–865 (2016)
2. Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: ICCV. pp. 6182–6191 (2019)
3. Chen, L.C., et al.: DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. TPAMI **40**(4), 834–848 (2017)
4. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: Atom: Accurate tracking by overlap maximization. In: CVPR. pp. 4660–4669 (2019)
5. Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M.: ECO: Efficient convolution operators for tracking. In: CVPR. pp. 6931–6939 (2017)
6. De Boer, P.T., Kroese, D.P., Mannor, S., Rubinstein, R.Y.: A tutorial on the cross-entropy method. *Annals of operations research* **134**(1), 19–67 (2005)
7. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet: Keypoint triplets for object detection. In: ICCV. pp. 6569–6578 (2019)
8. Fan, H., Lin, L., Yang, F., et al.: LaSOT: A high-quality benchmark for large-scale single object tracking. In: CVPR. pp. 5374–5383 (2019)
9. Fan, H., Ling, H.: Siamese cascaded region proposal networks for real-time visual tracking. In: CVPR. pp. 7952–7961 (2019)
10. Gao, J., et al.: Graph convolutional tracking. In: CVPR. pp. 4649–4659 (2019)
11. Guo, Q., Feng, W., Zhou, C., Huang, R., Wan, L., Wang, S.: Learning dynamic siamese network for visual object tracking. In: ICCV. pp. 1763–1771 (2017)
12. He, K., Gkioxari, G., et al.: Mask R-CNN. In: ICCV. pp. 2961–2969 (2017)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
14. Huang, L., Zhao, X., Huang, K.: Got-10k: A large high-diversity benchmark for generic object tracking in the wild. TPAMI **1**(1), 1–17 (2019)
15. Jung, I., Son, J., et al.: Real-Time MDNet. In: ECCV. pp. 83–98 (2018)
16. Kiani Galoogahi, H., Fagg, A., Lucey, S.: Learning background-aware correlation filters for visual tracking. In: ICCV. pp. 1135–1143 (2017)
17. Kristan, M., Leonardis, A., Matas, et al.: The sixth visual object tracking vot2018 challenge results. In: ECCVW. pp. 1–52 (2018)
18. Kristan, M., Matas, et al.: The seventh visual object tracking vot2019 challenge results. In: ICCVW. pp. 1–36 (2019)
19. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: ECCV. pp. 734–750 (2018)
20. LeCun, Y., Boser, B., et al.: Backpropagation applied to handwritten zip code recognition. *Neural computation* **1**(4), 541–551 (1989)
21. Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., Yan, J.: SiamRPN++: Evolution of siamese visual tracking with very deep networks. In: CVPR. pp. 4282–4291 (2019)
22. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: CVPR. pp. 8971–8980 (2018)
23. Li, F., Tian, C., Zuo, W., Zhang, L., Yang, M.H.: Learning spatial-temporal regularized correlation filters for visual tracking. In: CVPR. pp. 4904–4913 (2018)
24. Li, P., Chen, B., Ouyang, W., Wang, D., Yang, X., Lu, H.: GradNet: Gradient-guided network for visual object tracking. In: ICCV. pp. 6162–6171 (2019)
25. Li, X., Hu, W., Shen, C., Zhang, Z., Dick, A., Hengel, A.V.D.: A survey of appearance models in visual object tracking. *TIST* **4**(4), 1–48 (2013)

26. Lin, T.Y., Maire, M., Belongie, S., et al.: Microsoft COCO: Common objects in context. In: ECCV. pp. 740–755 (2014)
27. Nam, H., Han, B.: Learning multi-domain convolutional neural networks for visual tracking. In: CVPR. pp. 4293–4302 (2016)
28. Park, E., Berg, A.C.: Meta-tracker: Fast and robust online adaptation for visual object trackers. In: ECCV. pp. 569–585 (2018)
29. Real, E., Shlens, J., et al.: Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In: CVPR. pp. 7464–7473 (2017)
30. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR. pp. 779–788 (2016)
31. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: NIPS. pp. 91–99 (2015)
32. Russakovsky, O., Deng, et al.: ImageNet large scale visual recognition challenge. *IJCV* **115**(3), 211–252 (2015)
33. Smeulders, A.W., Chu, D.M., Cucchiara, R., Calderara, S., Dehghan, A., Shah, M.: Visual tracking: An experimental survey. *TPAMI* **36**(7), 1442–1468 (2013)
34. Song, Y., Ma, C., Wu, X., Gong, L., Bao, L., Zuo, W., Shen, C., Yang, M.H.: VITAL: Visual tracking via adversarial learning. In: CVPR. pp. 8990–8999 (2018)
35. Tao, R., Gavves, E., Smeulders, A.W.: Siamese instance search for tracking. In: CVPR. pp. 1420–1429 (2016)
36. Tian, Z., Shen, C., Chen, H., He, T.: FCOS: Fully convolutional one-stage object detection. In: ICCV. pp. 9627–9636 (2019)
37. Tripathi, A.S., Danelljan, M., Van Gool, L., Timofte, R.: Tracking the known and the unknown by leveraging semantic information. In: BMVC. pp. 1192–1198 (2019)
38. Valmadre, J., Bertinetto, L., Henriques, et al.: End-to-end representation learning for correlation filter based tracking. In: CVPR. pp. 2805–2813 (2017)
39. Voigtlaender, P., Luiten, J., Torr, P.H., Leibe, B.: Siam R-CNN: Visual tracking by re-detection. arXiv preprint arXiv:1911.12836 (2019)
40. Vu, T., Jang, H., Pham, T.X., Yoo, C.: Cascade RPN: Delving into high-quality region proposal network with adaptive convolution. In: NIPS. pp. 1430–1440 (2019)
41. Wang, G., et al.: SPM-tracker: Series-parallel matching for real-time visual object tracking. In: CVPR. pp. 3643–3652 (2019)
42. Wang, N., Song, Y., Ma, C., Zhou, W., Liu, W., Li, H.: Unsupervised deep tracking. In: CVPR. pp. 1308–1317 (2019)
43. Wang, Q., Zhang, L., Bertinetto, L., Hu, W., Torr, P.H.: Fast online object tracking and segmentation: A unifying approach. In: CVPR. pp. 1328–1338 (2019)
44. Wu, Y., Lim, et al.: Object tracking benchmark. *TPAMI* **37**(9), 1834–1848 (2015)
45. Xu, T., Feng, Z.H., Wu, X.J., Kittler, J.: Learning adaptive discriminative correlation filters via temporal consistency preserving spatial feature selection for robust visual object tracking. *TIP* **28**(11), 5596–5609 (2019)
46. Yang, T., Chan, A.B.: Learning dynamic memory networks for object tracking. In: ECCV. pp. 152–167 (2018)
47. Yu, J., Jiang, Y., Wang, Z., Cao, Z., Huang, T.: Unitbox: An advanced object detection network. In: ACM MM. pp. 516–520 (2016)
48. Zhang, H., Dana, K., Shi, J., Zhang, Z., Wang, X., Tyagi, A., Agrawal, A.: Context encoding for semantic segmentation. In: CVPR. pp. 7151–7160 (2018)
49. Zhang, K., Zhang, L., Yang, M.H.: Fast compressive tracking. *TPAMI* **36**(10), 2002–2015 (2014)
50. Zhang, Z., Peng, H.: Deeper and wider siamese networks for real-time visual tracking. In: CVPR. pp. 4591–4600 (2019)