

Visual Tracking by TridentAlign and Context Embedding

Janghoon Choi¹, Junseok Kwon², and Kyoung Mu Lee¹

¹ ASRI, Department of ECE, Seoul National University

² School of CSE, Chung-Ang University

ultio791@snu.ac.kr, jskwon@cau.ac.kr, kyoungmu@snu.ac.kr

Abstract. Recent advances in Siamese network-based visual tracking methods have enabled high performance on numerous tracking benchmarks. However, extensive scale variations of the target object and distractor objects with similar categories have consistently posed challenges in visual tracking. To address these persisting issues, we propose novel TridentAlign and context embedding modules for Siamese network-based visual tracking methods. The TridentAlign module facilitates adaptability to extensive scale variations and large deformations of the target, where it pools the feature representation of the target object into multiple spatial dimensions to form a feature pyramid, which is then utilized in the region proposal stage. Meanwhile, context embedding module aims to discriminate the target from distractor objects by accounting for the global context information among objects. The context embedding module extracts and embeds the global context information of a given frame into a local feature representation such that the information can be utilized in the final classification stage. Experimental results obtained on multiple benchmark datasets show that the performance of the proposed tracker is comparable to that of state-of-the-art trackers, while the proposed tracker runs at real-time speed.[†]

1 Introduction

Visual tracking, which has practical applications such as automated surveillance, robotics, and image stabilization, is one of the fundamental problems among the fields under computer vision research. Given initial target bounding box coordinates along with the first frame of a video sequence, visual tracking algorithms aim to precisely locate the target object in the subsequent frames of the video. Tracking algorithms are designed to successfully track targets under various circumstances such as scale change, illumination change, occlusion, deformation, and motion blur.

Along with the wide application of convolutional neural networks (CNNs) to various computer vision tasks [1,2,3], recent advances in Siamese network-based visual tracking methods [4,5,6,7] have advantages in performance and speed owing to the scalability of the network and elimination of online updates.

[†]Code available on <https://github.com/JanghoonChoi/TACT>

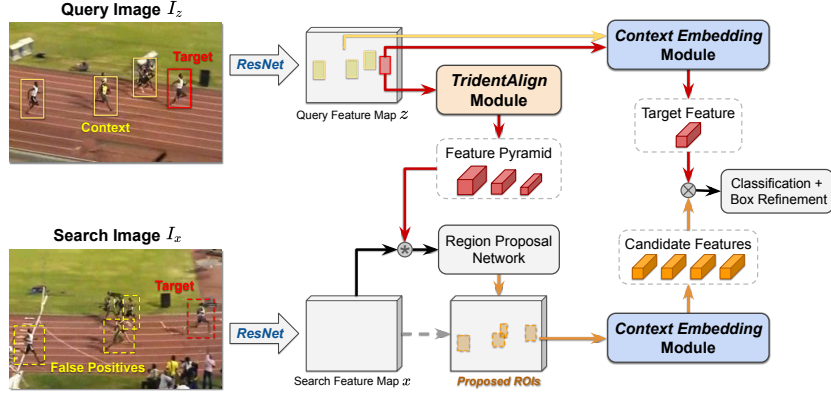


Fig. 1. Overview of the proposed visual tracking framework. Our tracker incorporates the TridentAlign module, which generates a feature pyramid representation of the target that can be utilized for better scale adaptability of the RPN. Moreover, the context embedding module modulates the locally pooled features to incorporate the global context information of a given frame, which encourages the discrimination of the target object from false positives.

However, most existing Siamese trackers are still designed and trained for short-term tracking scenarios [8,9] with strong assumptions of motion smoothness and gradual scale change. In order to impose these assumptions, most trackers assume a small search region around the target and penalize large displacements using hand-crafted window functions. This makes the tracker susceptible to error accumulation and drift due to erroneous localizations, leading to lower performance in numerous long-term tracking scenarios [10,11,12] in which smoothness assumptions are no longer valid. To alleviate these problems, recent Siamese trackers have adopted full-frame search [13,14], so that the target location can be recovered after the out-of-frame disappearance, at the cost of computation time and with sub-real-time speeds. Although a full-frame search can be effective for re-detecting the target, it is susceptible to distractors with appearances similar to that of the target owing to the lack of temporal consistency and global context modeling. Diverse scale variations of the target can also lead to failed re-detection because the target feature representation has a fixed spatial dimension that may fail to represent the wide spatial variation of the target object.

In this paper, we propose a novel real-time visual tracking algorithm to address the aforementioned issues by incorporating the TridentAlign and context embedding modules into our tracking framework, where our region proposal network (RPN) is inspired by the success of recent object detectors. We incorporate the anchor-free FCOS [15] detector head in our RPN to minimize the number of parameters while maximizing the flexibility with regard to large deformations in the aspect ratio. To further enforce the scale adaptiveness of the RPN, we use the target feature representation obtained from the TridentAlign module, which

pools the features from the target region to multiple spatial dimensions to form a feature pyramid, where similar approaches have been shown to be effective in [16,17,18,19]. Moreover, rather than focusing only on the similarity between the locally obtained features, we make full use of the global context information obtained from the given frames. The context-embedded features obtained from the context embedding module are utilized to discriminate the target from the distractor objects in the background region to avoid choosing false-positives. The context embedding module receives hard negative features pooled from a given frame that encompasses potential distractors. It then modulates the local feature representations of the candidate regions, providing additional information on global context.

We compare our method to other Siamese network-based trackers on large-scale visual tracking datasets LaSOT [10], OxUvA [12], TrackingNet [11], and GOT-10k [20] thereby demonstrating performance comparable to that of state-of-the-art trackers. Moreover, our proposed modules require minimal computational overhead, and our tracker is able to run at real-time speed. Our overall tracking framework is shown in Figure 1.

2 Related Work

CNN-based trackers : Conventional online visual tracking algorithms solve the tracking problem via tracking-by-detection, where they attempt to locate the target inside a search region by finding the position where the classifier produces the highest similarity/classification score to that of the target. Given the powerful representation capabilities of CNNs, recent trackers use CNNs for feature representation and classification. One early CNN-based tracker [21] used the feature representation obtained from the denoising autoencoder network. In MDNet [22], VGG [23] features with multi-task training are used, and this tracker is accelerated to real-time speed in [24] by using ROIAlign [25]. Correlation filter-based trackers [26,27] are also widely used on top of pretrained deep features. Notable approaches include the use of the continuous convolutional operator for the fusion of multi-resolution CNN features [28,29], spatially regularized filters [30], the fusion of deep and shallow features [31], and group feature selection [32].

Siamese network-based trackers : Siamese network-based trackers have gained attention owing to their simplicity and high performance. SiamFC [4] proposed a fully convolutional end-to-end approach to visual tracking, enabling increased speed and accuracy, in which correlation filter can be also trained on top of the feature map [33]. SiamRPN [5] added RPN in the final stage for more accurate localization and size estimation, and DaSiamRPN [34] improved the discriminability of the network by adding negative pairs during training to rule out distractors. Both [6] and [7] employed deeper and wider backbone networks based on [2] and [35] with extensive experimental analysis to achieve further performance gains. Other notable approaches include the use of a dynamic network with a transformation learning model [36], pattern detection for local structure-based prediction [37], and cascaded region proposal for progressive refinement

[38]. Recently, noteworthy methods have been proposed to improve the discriminability of the network by performing adaptation at test-time. Particularly, [39] and [40] used gradient information obtained during tracking to update the target representation. Moreover, [41] used a learned update module to obtain the updated accumulated template. An optimization-based model predictor is used in [42] to predict the classifier weights.

Context-aware trackers : Conventional visual tracking algorithms treat all candidate regions in the search image independently. They perform tracking by choosing the candidate region with the highest similarity score to the target, where other candidate regions have no influence on this decision process. Online adaptation-based approaches partially address this issue by updating the tracker with previously obtained self-labeled training examples at test-time [22,40,42]; however, the possibility of error accumulation and drift persists. Several tracking approaches take a broader context area into consideration [43,44]. Particularly, [43] employs a context-aware network in the initial adaptation process to choose an expert autoencoder that is used throughout the online tracking process. Furthermore, [44] proposed a generic context-aware correlation filter learning scheme that can add contextual information for better background suppression. Nevertheless, the context modeling in these approaches is still limited to a fixed spatial region around the target, which lacks the global context information of all the possible distractors in a given scene.

Long-term trackers : With growing interest on long-term visual tracking tasks and more benchmarks becoming available to the public [10,45,12], tracking algorithms focused on solving long-term tracking scenarios have emerged. Long-term visual tracking problems are typically defined as problems involving the tracking of a target object in a relatively long (*e.g.*, > 1 min) video sequence, during which the target may disappear for a brief period [12]. The main challenge of long-term tracking tasks is recovering the target after its disappearance or a tracking failure. To solve this problem, [34,46] attempted to expand the search area when the confidence value falls below a predefined threshold. Moreover, [13,14] performed a full-frame search without any locality assumption. Our approach shares some similarities with [13], but it differs significantly in these aspects: (1) We use the proposed TridentAlign module to construct a feature pyramid representation of the target for improved scale adaptability and part-based predictions in the region proposal stage, whereas [13] pools the target feature into a 1×1 kernel, losing spatial information; (2) Our tracker introduces the context embedding module, which employs hard negative features obtained from potential distractors as contextual information, and the local feature representations can be modulated accordingly during tracking. In contrast, [13] only uses local features without any consideration of the context, which makes it susceptible to distractors; (3) Our tracker runs at a real-time speed of 57 fps with a ResNet-18 backbone and 42 fps with a ResNet-50 backbone. It thus achieves higher performance on LaSOT [10] with a significantly lighter backbone network (ResNet-18) and is approximately 9 times faster than the model proposed in [13], which uses a significantly deeper backbone network (ResNet-50).

3 Proposed Method

Inspired by two-stage object detection networks [3,16,17], our framework largely includes two stages: the region proposal stage and classification stage. Given a pair of input RGB images, that is, the query image (initial frame) I_z and search image (current frame) I_x along with the shared backbone network $\varphi(\cdot)$, the respective feature maps $z = \varphi(I_z)$ and $x = \varphi(I_x)$ are obtained and then passed to the RPN. In the region proposal stage, the RPN generates proposals for target region-of-interest (RoIs) in the search image, given the target information obtained from the query image. Subsequently, in the classification stage, the classifier performs binary classification on the ROI proposals obtained in the previous stage, with a positive result indicating the target region and a negative indicating the background region. Figure 1 shows the overall flow of the proposed method.

In the following subsections, we provide more detailed explanations of the proposed TridentAlign module used in the region proposal stage and context embedding module used in the classification stage. Then, we describe the online tracking procedure for our approach.

3.1 Region Proposal with Scale Adaptive TridentAlign

The initial target bounding box coordinates and input feature maps are given as $z, x \in \mathbb{R}^{h \times w \times c}$, where the channel dimension of the input feature maps are reduced from the original outputs of the ResNet backbone by employing 1×1 conv layers. Using these feature maps, the TridentAlign module performs multiple ROIAlign [25] operations on z with varying spatial dimensions to obtain target feature representations $z_i \in \mathbb{R}^{s_i \times s_i \times c}$, where s_i is the spatial dimension of the pooled features. These features form a feature pyramid denoted as $Z = \{z_1, z_2, \dots, z_K\}$, where K is the total number of features in the feature pyramid. Then, the depth-wise cross-correlation between search feature map x and each target feature z_i in the feature pyramid Z is calculated as

$$\hat{x}_i = x \circledast z_i, \quad (1)$$

where \circledast denotes the depth-wise cross-correlation operator with zero padding. As a result, each $\hat{x}_i \in \mathbb{R}^{h \times w \times c}$ is obtained for the corresponding z_i and is concatenated along the channel dimension to form the multi-scale correlation map $[\hat{x}_1, \hat{x}_2, \dots, \hat{x}_K] = \hat{x} \in \mathbb{R}^{h \times w \times Kc}$. The correlation map is then refined as in $f_{att}(\hat{x}) \in \mathbb{R}^{h \times w \times c}$ using a self-attention block analogous to that employed in [47]. In this self-attention block, the adaptive channel and spatial attention weights are applied to focus on a specific position and target scale, followed by a 1×1 conv layer, thereby reducing the channel dimension back to c .

With the refined correlation map, we use the detection head module employed in [15], where each branch outputs binary classification labels and bounding box regression values. For a single location (i, j) inside the output map, classification labels $p_{i,j}$ with bounding box regression values $t_{i,j}$ are predicted from the respective branches. At training stage, if a location (i, j) is inside the ground-truth

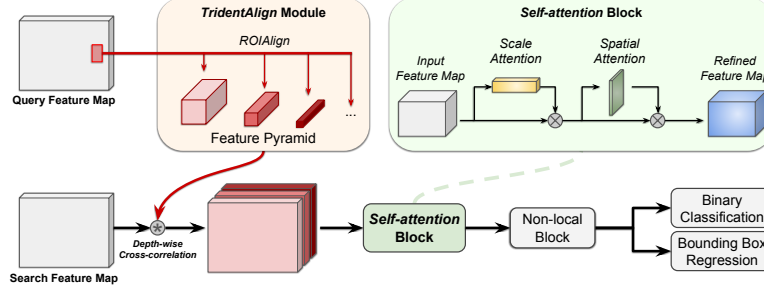


Fig. 2. Overview of the proposed region proposal network. The feature pyramid representation of the target is constructed using our TridentAlign module, wherein each feature undergoes a depth-wise cross-correlation operation with the search feature map. The correlated feature maps are concatenated along the channel dimension; here, a self-attention block is used to focus more on a certain spatial area with a certain target scale. Followed by a non-local block [48] and binary classification/bounding box regression branches, ROI can be obtained.

(GT) target bounding box, it is considered a positive sample and we assign the GT label $c_{i,j}^* = 1$ and GT regression target $t_{i,j}^* = (l^*, t^*, r^*, b^*)$, where l^*, t^*, r^* , and b^* are the distances from (i, j) to the four sides (left, top, right, bottom) of the bounding box, respectively. For negative samples, we assign $c_{i,j}^* = 0$. To train the overall RPN, we use the same loss as the one used in [15]:

$$L_{rpn}(\{p_{i,j}\}, \{t_{i,j}\}) = \frac{1}{N_{pos}} \sum_{i,j} L_{cls}(p_{i,j}, c_{i,j}^*) + \frac{\lambda}{N_{pos}} \sum_{i,j} \mathbf{1}_{\{c_{i,j}^* > 0\}} L_{reg}(t_{i,j}, t_{i,j}^*), \quad (2)$$

where N_{pos} is the number of positive samples, L_{cls} is the focal loss [17], and L_{reg} is the linear IoU loss. The loss is summed over all locations of the output map, and L_{reg} is only calculated for positive samples. Subsequently, a non-maximum suppression (NMS) operation is performed to obtain the top N candidate ROIs. The overall architecture of the RPN is illustrated in Figure 2.

3.2 Classification with Context-Embedded Features

Given the candidate ROIs obtained from the preceding region proposal stage, ROIAlign operations are performed on the search feature map x to obtain a set of candidate features $X = \{x_1, x_2, \dots, x_N\}$, where each $x_i \in \mathbb{R}^{s \times s \times c}$ with N candidate regions. Using all of the candidate features in X to generate the global context information, we aim to modulate each feature x_i to obtain the context-embedded feature $\tilde{x}_i \in \mathbb{R}^{s \times s \times c}$. First, element-wise maximum and averaging operations are performed over all features in X to obtain max-pooled and average-pooled features, which are concatenated along the channel dimension as $x_{ctx} \in \mathbb{R}^{s \times s \times 2c}$. Then, inside the context embedding module, context generator

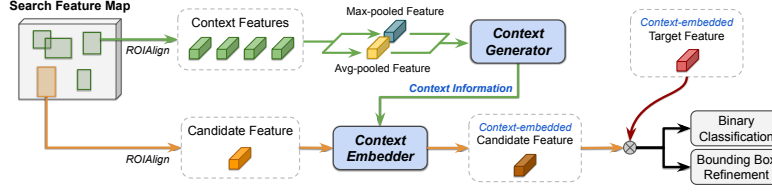


Fig. 3. Overview of our context embedding framework. Given the candidate ROI and context regions, the respective feature representations are obtained by performing ROIAlign operations on each region. Using the context features, max-pooled and average-pooled features are obtained via element-wise maximum and averaging operations. The context generator receives these features to generate the global context information, which the context embedder embeds into the candidate features. Finally, the context-embedded candidate feature can be compared with the context-embedded target features for binary classification and bounding box refinement.

$g_1(\cdot)$ receives x_{cxt} to generate the global context information, and context embedder $g_2(\cdot)$ receives both the candidate feature x_i and context information from g_1 to generate the context-embedded feature \tilde{x}_i . The overall context embedding scheme is illustrated in Figure 3. For our context generator and embedder design, we test 4 variants: (1) simple concatenation, (2) simple addition, (3) CBAM [47], and (4) FILM [49] based modules. The details of each variant are listed in Table 1.

For the simple concatenation-based module (Table 1(1)), x_{cxt} is directly used as context information and the context embedder g_2 receives $[x_i, x_{cxt}]$ as input, where $[\cdot, \cdot]$ denotes concatenation along the channel dimension. In the simple addition-based module (Table 1(2)), context information is generated in a form of additive modulation δ , which is added to the original candidate feature x_i . For the CBAM-based module (Table 1(3)), context information is generated and applied to x_i as channel attention m_c and spatial attention m_s . Finally, the FILM-based module (Table 1(4)) modulates x_i by applying an affine transformation with coefficients γ and β .

Table 1. Variants of the context embedding module. We test 4 possible implementations of the context embedding module. For 3-layer CNNs, we use 1×1 kernels with output channels set to c , followed by ReLU activation.

	Generator $g_1(x_{cxt})$		Embedder $g_2(g_1(x_{cxt}), x_i)$
	Type	Output	Operation
(1) Simple concat.	Identity	$x_{cxt} \in \mathbb{R}^{s \times s \times 2c}$	3-layer CNN, $g_2([x_i, x_{cxt}])$
(2) Simple add.	3-layer CNN	$\delta \in \mathbb{R}^{s \times s \times c}$	$x_i + \delta$
(3) CBAM-based	3-layer CNN	$m_c \in \mathbb{R}^{1 \times 1 \times c}, m_s \in \mathbb{R}^{s \times s \times 1}$	$(x_i \otimes m_c) \otimes m_s$
(4) FILM-based	3-layer CNN	$\gamma, \beta \in \mathbb{R}^{s \times s \times c}$	$\gamma \otimes x_i + \beta$

Algorithm 1: TACT

Input : Tracking sequence of length L , $\{I^1, I^2, \dots, I^L\}$
Initial target bounding box coordinates
Output: Target bounding box coordinates for each frame

Initialization at $t = 1$
Compute query feature map $z = \varphi(I^1)$ for initial frame I^1
Build target feature pyramid Z from z using TridentAlign
Using same z as search feature map; obtain candidate features using RPN
Obtain context-embedded target feature \tilde{z}_0

For later frames in tracking sequence
for $t = 2$ **to** L **do**
 Compute search feature map $x = \varphi(I^t)$ for frame I^t
 Using Z and x , obtain ROIs with candidate features X using RPN
 For every $x_i \in X$, calculate context-embedded feature \tilde{x}_i to form \tilde{X}
 Compute $\tilde{x}_i \otimes \tilde{z}_0$ for every $\tilde{x}_i \in \tilde{X}$
 For every ROI, obtain softmax classification scores and box refinement values
 Choose refined ROI with highest classification score as output
end

Finally, each context-embedded candidate feature \tilde{x}_i is compared with the context-embedded target feature $\tilde{z}_0 \in \mathbb{R}^{s \times s \times c}$ by element-wise multiplication as in $\tilde{x}_i \otimes \tilde{z}_0$. Binary classification and bounding box refinement operations are subsequently performed. For every \tilde{x}_i , a classification label c_i and refined bounding box coordinates t_i are obtained. At training stage, the GT classification label $c_i^* = 1$ is assigned to candidate boxes with $\text{IoU}(t_i, t_i^*) > \tau_p$, where t_i^* is the GT box coordinates, and $c_i^* = 0$ is assigned to candidate boxes with $\text{IoU}(t_i, t_i^*) < \tau_n$. In our experiments, we use $(\tau_p, \tau_n) = (0.5, 0.4)$. To train our context embedding module and classifier, we minimize the loss function given as

$$L_{det}(\{c_i\}, \{t_i\}) = \frac{1}{N_{pos}} \sum_i L_{cls}(c_i, c_i^*) + \frac{\lambda}{N_{pos}} \sum_i \mathbf{1}_{\{c_i^* > 0\}} L_{reg}(t_i, t_i^*), \quad (3)$$

where the loss functions L_{cls} and L_{reg} are the same as those in Eq. (2).

3.3 TridentAlign and Context Embedding Tracker

Herein, we propose **TridentAlign** and **Context embedding Tracker (TACT)**. The overall tracking procedure is organized and shown as Algorithm 1. The tracking process is purposely made simple to achieve real-time speed. Furthermore, our tracking algorithm performs a full-frame search for every frame without any motion smoothness assumption based on the previous positions of the target; therefore, it is possible to run our tracker on a batch of multiple frames of offline videos. Increasing the batch size from 1 to 8 results in a large boost in tracking speed: we obtain $57 \rightarrow 101$ fps with the ResNet-18 backbone and $42 \rightarrow 65$ fps with the ResNet-50 backbone.

4 Experiments

In this section, we specify the implementation details and the experimental setup, then we compare the performance of our proposed tracking algorithm with that of other approaches on four large-scale tracking benchmarks, namely LaSOT [10], OxUvA [12], TrackingNet [11], and GOT-10k [20]. We also perform ablation experiments for the individual components to further analyze the effectiveness of our proposed method.

4.1 Implementation Details

Parameters: We resized the input images to 666×400 , where original aspect ratio is preserved by adding zero-padding to the right or bottom side of the images. We used ResNet-18 and ResNet-50 [2] as the backbone feature extractor network, followed by 1×1 conv layers, where the channel dimension of the output features was set to $c = 256$. We reduced the stride of the last residual block to 1 to obtain feature maps with a size of 42×25 . Regarding the RPN stage, the TridentAlign module generates a feature pyramid of size $K = 3$ with spatial dimensions $s_i \in \{3, 5, 9\}$. A total of $N = 64$ ROI proposal boxes are obtained via NMS with an overlap threshold value of 0.9. In the subsequent classification stage, the spatial dimension of the pooled candidate features obtained by ROIAlign was set to $s = 5$.

Training data: To train the model, we used training splits of the ImageNetVID [50], YouTubeBB [51], GOT-10k [20], and LaSOT [10] datasets, and pairs of query and search images are uniformly sampled from the video sequences in these datasets. When sampling an image pair, a video sequence was chosen randomly where the probability of choosing a certain dataset is proportional to its total size. For a sampled image pair, we performed random data augmentation such as horizontal flips and the addition of gaussian noise, blurring, and color jitter. The bounding box coordinates were also randomly augmented by $\pm 1\%$ of their original width/height.

Training details: We optimized the sum loss functions $L = L_{rpn} + L_{det}$, where the losses are given in Eq. (2) and Eq. (3) with $\lambda = 1$. We used the Adam [52] optimizer with a batch size of four pairs to train our network. The learning rate was set to 10^{-4} , and the weight decay coefficient was set to 10^{-5} . For initialization, we used pretrained weights from the ResNet architectures, and during training, we freeze the weights of the residual blocks, except for the last block. We first trained the network for 2×10^6 iterations without the context embedding module and decayed the learning rate by a factor of 0.5 halfway. Then, we added the context embedding module and trained the network for another 10^6 iterations with a learning rate of 10^{-5} . We allocated an initial burn-in phase of 10^4 iterations, during which only the RPN was trained. In this way, we prevent negative candidate ROIs from overwhelming the classification stage, which can stabilize the training process. Our model was implemented in Python using the PyTorch [53] library. For run-time measurements, we run and time our model on a single Nvidia RTX 2080Ti GPU.

Table 2. Comparison on the **LaSOT** test set.

	TACT-18	TACT-50	GlobalTrack [13]	ATOM [54]	DiMP-50 [42]	SiamRPN++ [6]	DASiam [34]	SPLT [55]	MDNet [22]	Ocean [56]	SiamFC [4]	CFNet [33]
AUC	0.556	0.575	0.521	0.518	0.569	0.496	0.448	0.426	0.397	0.560	0.336	0.275
Precision	0.583	0.607	0.529	0.506	-	0.491	0.427	0.396	0.373	0.566	0.339	0.259
Norm. Precision	0.638	0.660	0.599	0.576	0.650	0.569	-	0.494	0.460	-	0.420	0.312
FPS	57	42	6	30	43	35	110	25.7	0.9	25	58	43

Table 3. Comparison on the **OxUvA** test set.

(%)	TACT-50	GlobalTrack [13]	SPLT [55]	MBMD [57]	DASiam _{LT} [34]	EBT [58]	SiamFC _{+R} [4]	SINT [59]	LCT [46]	TLD [60]	MDNet [22]	ECO-HC [29]
MaxGM	70.9	60.3	62.2	54.4	41.5	28.3	45.4	32.6	39.6	43.1	34.3	31.4
TPR	80.9	57.4	49.8	60.9	68.9	32.1	42.7	42.6	29.2	20.8	47.2	39.5
TNR	62.2	63.3	77.6	48.5	0.0	0.0	48.1	0.0	53.7	89.5	0.0	0.0

4.2 Quantitative Evaluation

Evaluation datasets and metrics: We evaluated our tracker (hereafter denoted as TACT) on test sets including four large-scale visual tracking benchmark datasets: LaSOT [10], OxUvA [12], TrackingNet [11], and GOT-10k [20]. The parameters were fixed for all benchmarks and experiments. LaSOT and OxUvA are long-term tracking benchmarks whose average sequence length is longer than 1 min, whereas TrackingNet and GOT-10k have shorter sequences but include a larger number of sequences with more various classes of objects.

The **LaSOT** [10] dataset is a large-scale and long-term tracking dataset consisting of 1,400 long-term sequences with an average sequence length of 2,512 frames (83 secs), which are annotated with the bounding box coordinates of the target object. We evaluated our tracker on the test set that includes 280 sequences under a one-pass evaluation setting, where the performance metrics are the area-under-curve (AUC) of the success plot, location precision, and normalized precision. The **OxUvA** [12] dataset is used to evaluate long-term tracking performance where its development and test sets have 200 and 166 sequences, respectively, with an average length of 4,260 frames (142 secs). In addition to evaluating the accuracy of the predicted boxes, the tracker must also report whether the target is present/absent in a given frame. The performance metric is the maximum geometric mean (MaxGM) over the true positive rate (TPR) and the true negative rate (TNR). **TrackingNet** [11] is a large-scale dataset of more than 30,000 videos collected from YouTube, of which 511 are included in the test set. Similar to the other benchmarks, it uses precision, normalized precision, and the AUC of the success plot as performance metrics. **GOT-10k** [20] is a tracking dataset focused on the one-shot experiment setting in which the training and test sets have disjoint object classes. It contains 10,000 videos of which 420 are used as the test set. Trackers are evaluated by calculating the success rate (SR, with threshold values 0.5 and 0.75) and average overlap (AO) value.

Table 4. Comparison on the **TrackingNet** test set.

(%)	TACT-18	TACT-50	GlobalTrack [13]	ATOM [54]	DiMP-50 [42]	SiamRPN++ [6]	DA-Siam [34]	UPDT [31]	MDNet [22]	SiamFC [4]	CFNet [33]	ECO [29]
Precision	70.1	70.8	65.6	64.8	68.7	69.4	59.1	55.7	56.5	53.3	53.3	49.2
Norm. Precision	78.4	78.8	75.4	77.1	80.1	80.0	73.3	70.2	70.5	66.6	65.4	61.8
Success	73.4	74.0	70.4	70.3	74.0	73.3	63.8	61.1	60.6	57.1	57.8	55.4

Table 5. Comparison on the **GOT-10k** test set.

(%)	TACT-18	TACT-50	ATOM [54]	DiMP-50 [42]	SiamMask [61]	Ocean [56]	CFNet [33]	SiamFC [4]	GOTURN [62]	CCOT [28]	ECO [29]	CF2 [63]	MDNet [22]
SR_{0.50}	64.8	66.5	63.4	71.7	58.7	72.1	40.4	35.3	37.5	32.8	30.9	29.7	30.3
SR_{0.75}	44.7	47.7	40.2	49.2	36.6	-	14.4	9.8	12.4	10.7	11.1	8.8	9.9
AO	55.9	57.8	55.6	61.1	51.4	61.1	37.4	34.8	34.7	32.5	31.6	31.5	29.9

Comparison to other trackers: We evaluated our proposed tracker on the LaSOT test set and provide the results in Table 2, where we denote our tracker with the ResNet-18 backbone as **TACT-18** and that with the ResNet-50 backbone as **TACT-50**. Both variants of TACT outperform other recent ResNet backbone-based trackers, which are GlobalTrack [13], ATOM [54], SiamRPN++ [6], and SPLT [55]. Moreover, our tracker runs faster than these algorithms, at real-time speed. To further test the long-term tracking capabilities of TACT, we evaluated our tracker on the OxUvA test set and show the results in Table 3. To predict the presence/absence of the target, we simply used a threshold value of 0.95. Output confidence scores below the given threshold were considered to indicate absence of the target. The results show that our tracker outperforms other long-term tracking algorithms in terms of the MaxGM and TPR metrics by a substantial margin, even when compared to GlobalTrack [13] and SPLT [55], which are trackers specifically designed for long-term tracking applications.

We also evaluated TACT on the relatively short-term and large-scale tracking benchmarks, which are TrackingNet and GOT-10k. The evaluation results for TrackingNet are shown in Table 4, where both variants of TACT exhibit competitive results in terms of the precision and success rate metrics, outperforming most trackers. Furthermore, Table 5 shows that the proposed method also obtains consistent results with regard to the comparison performed on GOT-10k: both variants of TACT were able to demonstrate high performance on all metrics. Even without any temporal smoothness assumptions or manual parameter tuning, TACT was able to achieve superior performance on the short-term tracking datasets compared to that of conventional trackers that are focused on short-term tracking applications.

To analyze the effectiveness of our proposed TridentAlign and context embedding modules, we show the success plots for eight different challenge attributes of the LaSOT dataset in Figure 4. TACT achieved substantial performance improvements on the attributes of aspect ratio change, deformation, rotation, and scale variation. Compared to other full-frame search-based tracker GlobalTrack, TACT performed better by a large margin owing to its TridentAlign module, which facilitates robustness to scale variations and large deformations of the

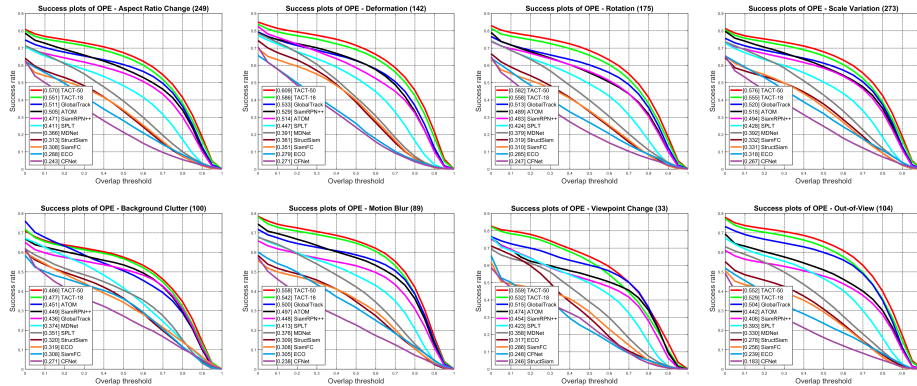


Fig. 4. Success plots for eight challenge attributes of the LaSOT dataset

target. Moreover, GlobalTrack only considers local features without any global context information; therefore, it is more prone to being affected by distractors similar to the target as shown in the background clutter plots, with inferior performance to that of ATOM and SiamRPN++. Using our context embedding module, robustness to background clutter can be reinforced via the global context modeling and embedding scheme. TACT also shows improvements with regard to the motion blur, viewpoint change, and out-of-view attributes owing to its full-frame search based design, which allows it to successfully recover from prolonged out-of-frame target disappearances and brief drifts.

In Figure 5, we present qualitative results produced by GlobalTrack, ATOM, SiamRPN++, SPLT, MDNet, and TACT-50 for selected videos from the LaSOT dataset. The results show that TACT successfully tracks the target despite challenging conditions such as large deformation in *kite-6*, occlusion in *fox-5*, and background clutter in *bicycle-9* and *skateboard-19*, whereas the other trackers fail. For additional attribute plots and qualitative results for other videos from the LaSOT dataset, please refer to the attached supplementary document and video.

4.3 Ablation Study

To provide more in-depth analysis of and insight into the proposed TACT, we performed additional ablation experiments for the individual components. For the following experiments, we used the test set of the LaSOT dataset and the success plot AUC metric to compare the performance of different variants of TACT.

Component-wise ablation: To validate the contribution of each individual component to the performance gain, we compared different variants of TACT by adding or removing the proposed modules. Table 6 shows the results of the abla-

Table 6. Ablation analysis of individual components. Adding each component contributes to consistent performance gains over the baseline model. As a performance measure, the AUC of the success plot on the LaSOT test set is shown.

TridentAlign	Context Embedding	AUC (TACT-18)	AUC (TACT-50)
\times	\times	0.535	0.552
\checkmark	\times	0.545	0.564
\checkmark	\checkmark	0.556	0.575

Table 7. Ablation analysis of the context embedding module. Among possible variants of the context embedding module, the FILM-based module shows the best performance. As a performance measure, the AUC of the success plot for the LaSOT test set is shown.

	No context	Simple concat.	Simple add.	CBAM-based	FILM-based
AUC (TACT-18)	0.545	0.532	0.552	0.551	0.556

tion analysis on individual components for both TACT-18 and TACT-50. Starting from the variants without the TridentAlign or context embedding modules, adding each component consistently improves the performance of both TACT-18 and TACT-50, which validates the effectiveness of our proposed approach. Our proposed TridentAlign and context embedding modules contribute to +2.1% and +2.3% to the performance gains of TACT-18 and TACT-50, respectively. All models were trained under the same settings using the same training datasets until convergence.

Variants of the context embedding module: We test four possible designs for the context embedding module, as introduced in Table 1. For the experiments, we started from a baseline model of TACT-18, which was trained without the context embedding module. Then, we added the context embedding module on top of the baseline model and trained the final model for additional iterations (as specified in Section 4.1). Table 7 shows the results for different variants of the context embedding module. Among all variants, the FILM-based module shows the highest performance gain followed by the simple addition-based module and the CBAM-based module. In contrast, the simple concatenation-based module degrades the performance. These results are somewhat consistent with other affine transform-based feature-wise transformation methods utilized in previous literature [49], where its multiplicative and additive operations provide adequate conditioning information for a given feature without hampering the discriminability of the original feature space.



Fig. 5. Qualitative Results on LaSOT. Tracking results for the sequences *kite-6*, *zebra-17*, *fox-5*, *bicycle-9*, *giraffe-10*, and *skateboard-19*. The color of the bounding box denotes a specific tracker. Yellow numbers on the top-left corner indicate frame indexes.

5 Conclusion

In this paper, we proposed a novel visual tracking method that aims to deal with large scale variations and deformations of the target while improving its discriminability by utilizing the global context information of the surroundings. Built upon a two-stage object detection framework, our proposed TACT incorporates the TridentAlign and context embedding modules to overcome the limitations of conventional tracking algorithms. The TridentAlign module constructs a target feature pyramid that encourages the adaptability of the tracker to large scale variations and deformations by fully utilizing the spatial information of the target. The context embedding module generates and embeds the global context information of a given frame into a local feature representation for improved discriminability of the target against distractor objects. The proposed modules are designed efficiently such that the overall framework can run at a real-time speed. Experimental results on four large-scale visual tracking benchmarks validate the strong performance of the TACT on long-term and short-term tracking tasks, achieving improved performance on challenging scenarios.

Acknowledgments

This work was supported by IITP grant funded by the Ministry of Science and ICT of Korea (No.2017-0-01780).

References

1. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: NIPS. (2012)
2. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. (2016)
3. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS. (2015)
4. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. arXiv preprint arXiv:1606.09549 (2016)
5. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: CVPR. (2018)
6. Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., Yan, J.: Siamrpn++: Evolution of siamese visual tracking with very deep networks. In: CVPR. (2019)
7. Zhang, Z., Peng, H.: Deeper and wider siamese networks for real-time visual tracking. In: CVPR. (2019)
8. Wu, Y., Lim, J., Yang, M.H.: Object tracking benchmark. IEEE TPAMI **37** (2015) 1834–1848
9. Čehovin, L., Leonardis, A., Kristan, M.: Visual object tracking performance measures revisited. IEEE TIP **25** (2016) 1261–1274
10. Fan, H., Lin, L., Yang, F., Chu, P., Deng, G., Yu, S., Bai, H., Xu, Y., Liao, C., Ling, H.: Lasot: A high-quality benchmark for large-scale single object tracking. In: CVPR. (2019)
11. Muller, M., Bibi, A., Giancola, S., Alsubaihi, S., Ghanem, B.: Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In: ECCV. (2018)
12. Valmadre, J., Bertinetto, L., Henriques, J.F., Tao, R., Vedaldi, A., Smeulders, A.W., Torr, P.H., Gavves, E.: Long-term tracking in the wild: A benchmark. In: ECCV. (2018)
13. Huang, L., Zhao, X., Huang, K.: Globaltrack: A simple and strong baseline for long-term tracking. In: AAAI. (2019)
14. Voigtlaender, P., Luiten, J., Torr, P.H., Leibe, B.: Siam r-cnn: Visual tracking by re-detection. In: CVPR. (2020)
15. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: ICCV. (2019)
16. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. IEEE TPAMI **37** (2015) 1904–1916
17. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. (2017)
18. Li, Y., Chen, Y., Wang, N., Zhang, Z.: Scale-aware trident networks for object detection. In: ICCV. (2019)
19. Hariharan, B., Arbeláez, P., Girshick, R., Malik, J.: Hypercolumns for object segmentation and fine-grained localization. In: CVPR. (2015)
20. Huang, L., Zhao, X., Huang, K.: Got-10k: A large high-diversity benchmark for generic object tracking in the wild. IEEE TPAMI (2019) 11
21. Wang, N., Yeung, D.Y.: Learning a deep compact image representation for visual tracking. In: NIPS. (2013)
22. Nam, H., Han, B.: Learning multi-domain convolutional neural networks for visual tracking. In: CVPR. (2015)

23. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint*, arXiv:1409.1556 (2014)
24. Jung, I., Son, J., Baek, M., Han, B.: Real-time mdnet. In: ECCV. (2018)
25. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV. (2017)
26. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. *IEEE TPAMI* **37** (2015) 583–596
27. Bolme, D.S., Beveridge, J.R., Draper, B.A., Lui, Y.M.: Visual object tracking using adaptive correlation filters. In: CVPR. (2010)
28. Danelljan, M., Robinson, A., Khan, F.S., Felsberg, M.: Beyond correlation filters: Learning continuous convolution operators for visual tracking. In: ECCV. (2016)
29. Danelljan, M., Bhat, G., Khan, F., Felsberg, M.: Eco: Efficient convolution operators for tracking. In: CVPR. (2017)
30. Danelljan, M., Hager, G., Shahbaz Khan, F., Felsberg, M.: Convolutional features for correlation filter based visual tracking. In: ICCV Workshop. (2015)
31. Bhat, G., Johnander, J., Danelljan, M., Shahbaz Khan, F., Felsberg, M.: Unveiling the power of deep tracking. In: ECCV. (2018)
32. Xu, T., Feng, Z.H., Wu, X.J., Kittler, J.: Joint group feature selection and discriminative filter learning for robust visual object tracking. In: ICCV. (2019)
33. Valmadre, J., Bertinetto, L., Henriques, J., Vedaldi, A., Torr, P.H.S.: End-to-end representation learning for correlation filter based tracking. In: CVPR. (2017)
34. Zhu, Z., Wang, Q., Li, B., Wu, W., Yan, J., Hu, W.: Distractor-aware siamese networks for visual object tracking. In: ECCV. (2018)
35. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR. (2015)
36. Guo, Q., Feng, W., Zhou, C., Huang, R., Wan, L., Wang, S.: Learning dynamic siamese network for visual object tracking. In: ICCV. (2017)
37. Zhang, Y., Wang, L., Qi, J., Wang, D., Feng, M., Lu, H.: Structured siamese network for real-time visual tracking. In: ECCV. (2018)
38. Fan, H., Ling, H.: Siamese cascaded region proposal networks for real-time visual tracking. In: CVPR. (2019)
39. Li, P., Chen, B., Ouyang, W., Wang, D., Yang, X., Lu, H.: Gradnet: Gradient-guided network for visual object tracking. In: ICCV. (2019)
40. Choi, J., Kwon, J., Lee, K.M.: Deep meta learning for real-time target-aware visual tracking. In: ICCV. (2019)
41. Zhang, L., Gonzalez-Garcia, A., Weijer, J.v.d., Danelljan, M., Khan, F.S.: Learning the model update for siamese trackers. In: ICCV. (2019)
42. Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: ICCV. (2019)
43. Choi, J., Jin Chang, H., Fischer, T., Yun, S., Lee, K., Jeong, J., Demiris, Y., Young Choi, J.: Context-aware deep feature compression for high-speed visual tracking. In: CVPR. (2018)
44. Mueller, M., Smith, N., Ghanem, B.: Context-aware correlation filter tracking. In: CVPR. (2017)
45. Moudgil, A., Gandhi, V.: Long-term visual object tracking benchmark. In: ACCV. (2018)
46. Ma, C., Yang, X., Zhang, C., Yang, M.H.: Long-term correlation tracking. In: CVPR. (2015)
47. Woo, S., Park, J., Lee, J.Y., So Kweon, I.: Cbam: Convolutional block attention module. In: ECCV. (2018)
48. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: CVPR. (2018)

49. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: AAAI. (2018)
50. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV **115** (2015) 211–252
51. Real, E., Shlens, J., Mazzocchi, S., Pan, X., Vanhoucke, V.: Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In: CVPR. (2017)
52. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. ICLR (2015)
53. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: NeurIPS. (2019)
54. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: Atom: Accurate tracking by overlap maximization. In: CVPR. (2019)
55. Yan, B., Zhao, H., Wang, D., Lu, H., Yang, X.: 'skimming-perusal'tracking: A framework for real-time and robust long-term tracking. In: ICCV. (2019)
56. Zhipeng Zhang, Houwen Peng, J.F.B.L.W.H.: Ocean: Object-aware anchor-free tracking. In: ECCV. (2020)
57. Zhang, Y., Wang, D., Wang, L., Qi, J., Lu, H.: Learning regression and verification networks for long-term visual tracking. arXiv preprint arXiv:1809.04320 (2018)
58. Zhu, G., Porikli, F., Li, H.: Beyond local search: Tracking objects everywhere with instance-specific proposals. In: CVPR. (2016)
59. Tao, R., Gavves, E., Smeulders, A.W.: Siamese instance search for tracking. In: CVPR. (2016)
60. Kalal, Z., Mikolajczyk, K., Matas, J.: Tracking-learning-detection. IEEE TPAMI **34** (2011) 1409–1422
61. Wang, Q., Zhang, L., Bertinetto, L., Hu, W., Torr, P.H.: Fast online object tracking and segmentation: A unifying approach. In: CVPR. (2019)
62. Held, D., Thrun, S., Savarese, S.: Learning to track at 100 fps with deep regression networks. In: ECCV. (2016)
63. Ma, C., Huang, J.B., Yang, X., Yang, M.H.: Hierarchical convolutional features for visual tracking. In: ICCV. (2015)