

Deeper and Wider Siamese Networks for Real-Time Visual Tracking

Zhipeng Zhang

University of Chinese Academy of Sciences&CASIA

zhipeng.zhang@nlpr.ia.ac.cn

Houwen Peng*

Microsoft Research

houwen.peng@microsoft.com

Abstract

Siamese networks have drawn great attention in visual tracking because of their balanced accuracy and speed. However, the backbone networks used in Siamese trackers are relatively shallow, such as AlexNet [18], which does not fully take advantage of the capability of modern deep neural networks. In this paper, we investigate how to leverage deeper and wider convolutional neural networks to enhance tracking robustness and accuracy. We observe that direct replacement of backbones with existing powerful architectures, such as ResNet [14] and Inception [33], does not bring improvements. The main reasons are that 1) large increases in the receptive field of neurons lead to reduced feature discriminability and localization precision; and 2) the network padding for convolutions induces a positional bias in learning. To address these issues, we propose new residual modules to eliminate the negative impact of padding, and further design new architectures using these modules with controlled receptive field size and network stride. The designed architectures are lightweight and guarantee real-time tracking speed when applied to SiamFC [2] and SiamRPN [20]. Experiments show that solely due to the proposed network architectures, our SiamFC+ and SiamRPN+ obtain up to 9.8%/5.7% (AUC), 23.3%/8.8% (EAO) and 24.4%/25.0% (EAO) relative improvements over the original versions [2, 20] on the OTB-15, VOT-16 and VOT-17 datasets, respectively.

1. Introduction

Visual tracking is one of the fundamental problems in computer vision. It aims to estimate the position of an arbitrary target in a video sequence, given only its location in the initial frame. Tracking at real-time speeds plays a vital role in numerous vision applications, such as surveillance, robotics, and human-computer interaction [10, 21, 30, 43].

Recently, trackers based on Siamese networks [2, 7, 12, 13, 20, 34, 40] have drawn great attention due to their high speed and accuracy. However, the backbone network utilized in these trackers is still the classical AlexNet [18],

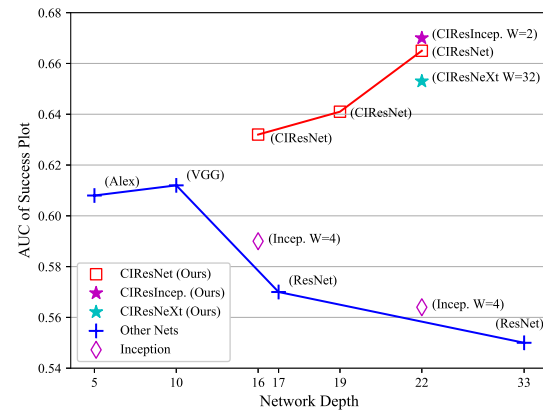


Figure 1: AUC of success plot vs. network depth and width (indicated by W). Here, width refers to the number of branches in a module. The results are obtained using SiamFC [2] with different backbone networks, through evaluation on OTB-13.

rather than modern deep neural networks that have proven more effective for feature embedding. To examine this issue, we replace the shallow backbone with deeper and wider networks, including VGG [29], Inception [33] and ResNet [14]. Unexpectedly, this straightforward replacement does not bring much improvement, and can even cause substantial performance drops when the network depth or width increases, as shown in Fig. 1. This phenomenon runs counter to the evidence that increasing network depth and width is beneficial for elevating model capability [14, 33].

One intuitive reasoning is that these deeper and wider network architectures are primarily designed for image classification tasks, where the precise localization of the object is not paramount. To investigate the concrete reason, we analyze the Siamese network architecture and identify that the receptive field size of neurons, network stride and feature padding are three important factors affecting tracking accuracy. In particular, the receptive field determines the image region used in computing a feature. A larger receptive field provides greater image context, while a small one may not capture the structure of target objects. The network stride affects the degree of localization precision, especially for small-sized objects. Meanwhile, it controls the size of output feature maps, which affects feature discriminability and detection accuracy. Moreover, for a fully-convolutional architecture [2], the feature padding for convolutions induces

*corresponding author

a potential position bias in model training, such that when an object moves near the search range boundary, it is difficult to make an accurate prediction. These three factors together prevent Siamese trackers from benefiting from current deeper and more sophisticated network architectures.

In this paper, we address these issues by designing new residual modules and architectures that allow deeper and wider backbone networks to unleash their power in Siamese trackers. First, we propose a group of cropping-inside residual (CIR) units based on the ‘‘bottleneck’’ residual block [14]. The CIR units crop out padding-affected features inside the block (i.e., features receiving padding signals), and thus prevent convolution filters from learning the position bias. Second, we design two kinds of network architectures, namely deeper and wider networks, by stacking the CIR units. In these networks, the stride and neuron receptive field are formulated to enhance localization precision. Finally, we apply the designed backbone networks to two representative Siamese trackers: SiamFC [2] and SiamRPN [20]. Experiments show that solely due to the proposed network architectures, the Siamese trackers obtain up to 9.8%/5.7%(AUC), 23.3%/8.8%(EAO) and 24.4%/25.0%(EAO) relative improvements over the original versions [2, 20] on the OTB-15, VOT-16 and VOT-17 datasets, respectively. In addition, the designed architectures are lightweight and permit the trackers to run at real-time speed.

The main contributions of this work are twofold.

- We present a systematic study on the factors of backbone networks that affect tracking accuracy, and provides architectural design guidelines for the Siamese tracking framework.
- We design new deeper and wider network architectures for Siamese trackers, based on our proposed no-padding residual units. Experimental results demonstrate that the new architectures provide clear improvements over the baseline trackers. Code and models are available at <https://github.com/researchmm/SiamDW>.

In the remainder of this paper, we first review background on Siamese tracking in Sec. 2. This is followed by an analysis of performance degradation in Sec. 3. Based on the analysis, we propose new residual modules and network architectures in Sec. 4. Experiments and comparisons are reported in Sec. 5. We end the paper with a discussion of related work and the conclusion in Sec. 6 and 7.

2. Background on Siamese Tracking

Before analyzing the reasons for the performance degradation shown in Fig. 1, we briefly review the fully-convolutional Siamese tracker SiamFC [2], which serves as the basic framework discussed in this work. The standard

Siamese architecture takes an image pair as input, comprising an exemplar image z and a candidate search image x . The image z represents the object of interest (e.g., an image patch centered on the target object in the first video frame), while x is typically larger and represents the search area in subsequent video frames. Both inputs are processed by a ConvNet φ with parameters θ . This yields two feature maps, which are cross-correlated as

$$f_{\theta}(z, x) = \varphi_{\theta}(z) \star \varphi_{\theta}(x) + b \cdot \mathbb{1}, \quad (1)$$

where $b \cdot \mathbb{1}$ denotes a bias term which takes the value $b \in \mathbb{R}$ at every location. Eq. 1 amounts to performing an exhaustive search of the pattern z over the image x . The goal is to match the maximum value in response map f to the target location.

To achieve this goal, the network is trained offline with random image pairs (z, x) taken from training videos and the corresponding ground-truth label y . The parameters θ of the ConvNet are obtained by minimizing the following logistic loss ℓ over the training set:

$$\arg \min_{\theta} \mathbb{E}_{(z, x, y)} \ell(y, f_{\theta}(z, x)). \quad (2)$$

Previous methods [2, 12, 13, 20, 40, 42] commonly utilize the classic and relatively shallow AlexNet [18] as the backbone network φ in this framework. In our work, we study the problem of how to design and leverage a more advanced ConvNet φ to learn an effective model θ that enhances tracking robustness and accuracy.

3. Analysis of Performance Degradation

In this section, we analyze the underlying reasons for the performance degradation presented in Fig. 1. We conduct ablation experiments on the internal factors of network structures, and identify the ones most responsible for performance drops. We then propose a set of practical guidelines for network architecture design, aimed to alleviate the negative effects.

3.1. Analysis

Quantitative analysis. Performance degradation can directly be attributed to network structure, as it is the only setting that changes in the experiments of Fig. 1. Therefore, we first identify the structural differences among these network architectures¹. As shown in Tab. 2, besides depth and width, there are several other internal network factors that differ among the networks, including stride (STR), padding (PAD), receptive field (RF) of neurons in the last layers, and output feature size (OFS).

To investigate the impact of these factors, we conduct an ablation study. We modify the structures of AlexNet, VGG,

¹Note that the network structures are slightly different from their original versions [14, 29, 33], where the network stride and padding are modified according to SiamFC [2].

Table 1: Analysis of network internal factors on AlexNet, VGG-10, Incep.-22 and ResNet-33. The numbers ①-⑩ represent different versions, in which the convolution kernel size, downsampling layer and padding are modified to show the trends. Details on the modifications are given in the supplementary material due to limited space.

# NUM	①	②	③	④	⑤	⑥	⑦	⑧	⑨	⑩	# NUM	①	②	③	④	⑤	⑥	⑦	⑧	⑨
RF ¹	Max(127)	+24	+16	+8	±0 (87)	±0	-8	-16	+16	+16	RF	+32	+16	+8	±0 (91)	±0	-8	-16	+16	+16
STR	8	8	8	8	8	8	8	8	16	4	STR	8	8	8	8	8	8	8	16	4
OFS	1	3	4	5	6	16	7	8	2	7	OFS	1	3	4	5	16	6	7	2	6
PAD	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗	PAD	✗	✗	✗	✗	✓	✗	✗	✗	✗
Alex	0.56	0.57	0.60	0.60	0.61	0.55	0.59	0.58	0.55	0.59	ResNet	0.56	0.59	0.60	0.62	0.56	0.60	0.60	0.54	0.58
VGG	0.58	0.59	0.61	0.61	0.62	0.56	0.59	0.58	0.54	0.58	Incep. ²	0.58	0.60	0.61	0.63	0.58	0.62	0.61	0.56	0.59

¹ To better show the trends, we denote ±0 as the original RF size of the network. + and - represent increasing and decreasing size over the originals. Max(127) represents the maximum effective RF, which is identical to the size of the exemplar image, i.e. 127 × 127 pixels.

² For the Inception network, its RF size lies in a range. Here we only list the theoretically maximum size, to align with ResNet for comparison.

	Alex	VGG-10	Incep.-16	Res.-17	Incep.-22	Res.-33
RF	87	103	23~183	227	39~519	739
STR	8	8	8	8	16	16
OFS	6	4	16	16	8	8
PAD	✗	✗	✓	✓	✓	✓
W	1	1	4	1	4	1
AUC	0.61	0.61	0.59	0.57	0.56	0.55

Table 2: Internal factors of different networks: receptive field (RF) of neurons in the last layer of network, stride (STR), output feature size (OFS), padding (PAD) and width (W). Since Inception contains multiple branches in one block, its RF lies within a range.

Inception and ResNet, and expose the effects of the internal factors. As shown in Tab. 1, when network stride (STR) increases from 4 or 8 to 16, the performance drops significantly (⑩ vs. ③ vs. ⑨ on AlexNet and VGG, ⑨ vs. ② vs. ⑧ on Incep. and ResNet). This illustrates that Siamese trackers prefer mid-level features (stride 4 or 8), which are more precise in object localization than high-level features (stride ≥ 16). For the maximum size of receptive field (RF), the optima lies in a small range. Specifically, for AlexNet, it ranges from 87-8 (Alex⑦) to 87+16 (Alex③) pixels; while for Incep.-22, it ranges from 91-16 (Incep.⑦) to 91+8 (Incep.③) pixels. VGG-10 and ResNet-17 also exhibit similar phenomena. In these cases, the optimal receptive field size is about 60%~80% of the input exemplar image z size (e.g. 91 vs 127). Intriguingly, this ratio is robust for various networks in our study, and it is insensitive to their structures. It illustrates that the size of RF is crucial for feature embedding in a Siamese framework. The underlying reason is that RF determines the image region used in computing a feature. A large receptive field covers much image context, resulting in the extracted feature being insensitive to the spatial location of target objects. On the contrary, a small one may not capture the structural information of objects, and thus it is less discriminative for matching. Therefore, only RF in a certain size range allows the feature to abstract the characteristics of the object, and its ideal size is closely related to the size of the exemplar image. For the output feature size, it is observed that a small size (OFS ≤ 3) does not benefit tracking accuracy. This is because small feature maps lack enough spatial structure description of target objects, and thus are not robust in image similarity calcula-

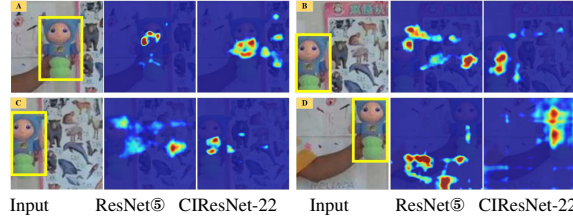


Figure 2: Visualization of position bias learnt in the model w/ and w/o padding (ResNet⑤ in Tab. 2 vs. ours). (A) presents the target at the image center, while (B-D) show it moving to boundaries due to imprecise tracked position in the previous frame.

tion. Moreover, as shown in Tab. 1 (⑤ vs. ⑥ on AlexNet and VGG, ④ vs. ⑤ on Incep. and ResNet), we observe that network padding has a highly negative impact on the final performance. To examine this further, we conduct a qualitative experiment.

Qualitative analysis. The Siamese networks feed pairs of exemplar and search images as training data, and learn a feature embedding for matching. If the networks contain padding operations, the embedding features of an exemplar image are extracted from the original exemplar image plus additional (zero)-padding regions. Differently, for the features of a search image, some of them are extracted only from image content itself, while some are extracted from image content plus additional (zero)-padding regions (e.g. the features near the border). As a result, there is inconsistency between the embeddings of target object appearing at different positions in search images, and therefore the matching similarity comparison degrades. Fig. 2 presents a visualization example of such inconsistency-induced effect in the testing phase. It shows that when the target object moves to image borders, its peak does not precisely indicate the location of the target. This is a common case caused by tracker drifts, when the predicted target location is not precise enough in the previous frame.

3.2. Guidelines

According to the above analysis, we summarize four basic guidelines to alleviate the negative impacts of structural factors in a network architecture.

- *Siamese trackers prefer a relatively small network stride.* Network stride affects the overlap ratio of receptive fields

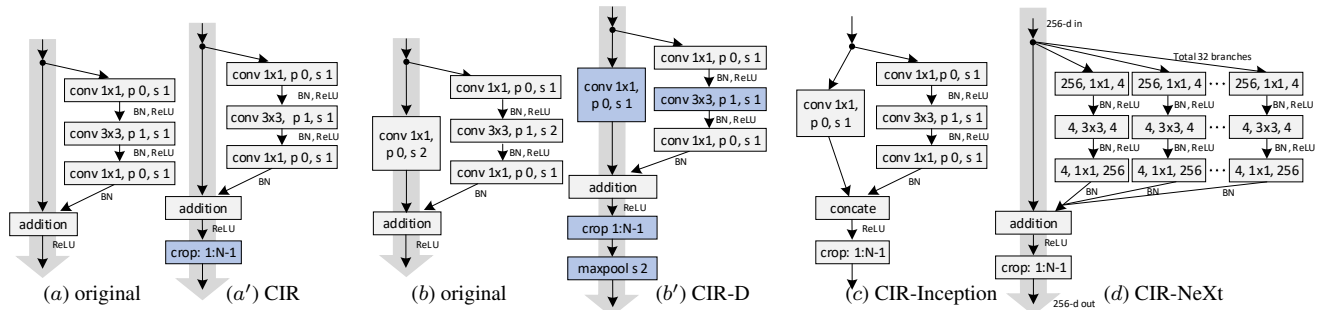


Figure 3: The proposed cropping-inside residual units. (a) and (b) are the original residual unit and downsampling unit, while (a') and (b') are our proposed ones. (c) and (d) are the proposed wide residual units. The grey arrows indicate the shortcut paths for easy information propagation, while the blue boxes highlight the differences from the original units. The letters 'p' and 's' indicate the padding size and stride, respectively. The settings of 'p' and 's' in (d) are the same as in (c).

for two neighboring output features. It thereby determines the basic degree of localization precision. Therefore, when network depth increases, the stride should not increase accordingly. With regards to accuracy and efficiency, an empirically effective choice is to set the stride to 4 or 8.

- *The receptive field of output features should be set based on its ratio to the size of the exemplar image.* An empirically effective ratio is 60%~80% for an exemplar image. A proper ratio allows the network to extract a set of features, each of which captures the information of different spatial parts of a target object. This leads the extracted features to be robust in calculating region similarity. Particularly, the maximum RF should not be larger than the exemplar image, otherwise the performance will drop significantly.
- *Network stride, receptive field and output feature size should be considered as a whole when designing a network architecture.* These three factors are not independent of one another. If one changes, the others will change accordingly. Considering them together can help the designed network to extract more discriminative features in a Siamese framework.
- *For a fully convolutional Siamese matching network, it is critical to handle the problem of perceptual inconsistency between the two network streams.* There are two feasible solutions. One is to remove the padding operations in networks, while the other is to enlarge the size of both input exemplar and search images, and then crop out padding-affected features.

4. Deeper and Wider Siamese Networks

In this section, we design new modules, i.e. cropping-inside residual (CIR) units, to eliminate the underlying position bias. Then, we build up deeper and wider backbone networks by stacking the new residual modules. The network stride and receptive field size are well controlled according to the guidelines. We further apply the designed networks to two representative Siamese trackers, i.e. SiamFC [2] and SiamRPN [20].

4.1. Cropping-Inside Residual (CIR) Units

The residual unit [14] is a key module in network architecture design due to its easy optimization and powerful representation. It consists of 3 stacked convolution layers and a shortcut connection that bypasses them, as shown in Fig. 3(a). The three layers are 1×1 , 3×3 and 1×1 convolutions, where the 1×1 layers are responsible for reducing and then restoring dimensions, leaving the 3×3 layer as a bottleneck with smaller input and output dimensions. This bottleneck convolution includes zero-padding of size 1, to ensure a compatible output size before the addition.

CIR Unit. As discussed in Sec. 3, network padding may introduce position bias in the Siamese framework. Hence, it is necessary to remove the padding in residual units when utilizing them to build a Siamese network. To this end, we augment the residual unit with a cropping operation, which is incorporated after the feature addition, as shown in Fig. 3(a'). The cropping operation removes features whose calculation is affected by the zero-padding signals. Since the padding size is one in the bottleneck layer, only the outermost features on the border of the feature maps are cropped out. This simple operation neatly removes padding-affected features in residual unit.

Downsampling CIR (CIR-D) Unit. The downsampling residual unit is another key building block for network design. It is utilized to reduce the spatial size of feature maps while doubling the number of feature channels. Similar to the residual unit, the downsampling unit also contains padding operations, as shown in Fig. 3(b). Thus, we also modify its structure to remove the negative effects caused by the padding. As shown in Fig. 3(b'), we change the convolutional stride from 2 to 1 within both the bottleneck layer and shortcut connection. Cropping is again inserted after the addition operation to remove the padding-affected features. Finally, max-pooling is employed to perform spatial downsampling of the feature map. The key idea of these modifications is to ensure that only the features influenced by padding are removed, while keeping the intrinsic block structure unchanged. If we were only to insert cropping after the addition operation, as done in the proposed CIR unit, without changing the position of downsampling, the

Table 3: Architectures of designed backbone networks for Siamese trackers. The CIR-D units are used in the first block of the ‘conv3’ stage, except for CIResNet-43, in which CIR-D is located in the fourth block.

Stage	CIResNet-16	CIResNet-19	CIResNet-22	CIResInception-22	CIResNeXt-22	CIResNet-43
conv1	$7 \times 7, 64, \text{stride } 2$					
	$2 \times 2 \text{ max pool, stride } 2$					
conv2	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 1$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 2$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64, C = 32 \\ 1 \times 1, 256 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 14$
conv3	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \\ [1 \times 1, 128] \times 4 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128, C = 32 \\ 1 \times 1, 512 \end{bmatrix} \times 4$	
	cross correlation Eq. 1					
# RF	77	85	93	13~93	93	105
# OFS	7	6	5	5	5	6
# Params	1.304 M	1.374 M	1.445 M	1.695 M	1.417 M	1.010 M
# FLOPs	2.43 G	2.55 G	2.65 G	2.71 G	2.52 G	6.07 G

features after cropping would not receive any signal from the outermost pixels in the input image. As the network depth increases, this would effectively cause even more image content to be removed, resulting in noisy/incomplete extracted features.

CIR-Inception and CIR-NeXt Units. We further equip the CIR unit with a multi-branch structure, enabling it to be used in building wide networks. Similar to Inception [33] and ResNeXt [39], we widen the CIR unit with multiple feature transformations, generating the CIR-Inception and CIR-NeXt modules as shown in Fig. 3(c-d). Specifically, in the CIR-Inception structure, we insert a 1×1 convolution into the shortcut connection, and merge the features of the two branches by concatenation, rather than by addition. In CIR-ResNeXt, we split the bottleneck layer into 32 transformation branches, and aggregate them by addition. Moreover, for the downsampling units of CIR-Inception and CIR-NeXt, the modifications are the same as those in CIR-D (Fig. 3(b’)), where the convolution stride is reduced and max-pooling is added. These two multi-branch structures enable the units to learn richer feature representations.

4.2. Network Architectures

By stacking the above CIR units, we build up deeper and wider networks. The constructions follow our design guidelines. First, we determine the network stride. A stride of 8 is used to build a 3-stage network, while a stride of 4 is employed in constructing a 2-stage one. Then, we stack CIR units. We control the number of units and the position of downsampling units in each stage. The goal is to ensure that the receptive field size of neurons in the final layer lies within the derived range, i.e. 60%-80% of the exemplar image. Additionally, when network depth increases, the receptive field may exceed this range. Therefore, we halve the stride to 4 to control the receptive field.

Deeper Networks. We construct deeper networks using CIR and CIR-D units. The structures are similar to ResNet [14], but with different network stride, receptive field size, and building blocks. In Tab. 3, we present four

deep cropping-inside residual networks, i.e. CIResNet-16, 19, 22 and 43. Since these networks have similar structure, we present details for just two of them: CIResNet-22 and CIResNet-43.

CIResNet-22 has 3 stages (stride=8) and consists of 22 weighted convolution layers. Except for the first 7×7 convolution, the others are all CIR units. A cropping operation (with size of 2) follows the 7×7 convolution to remove padding-affected features. The feature downsampling in the first two stages are performed by a convolution and a max-pooling of stride 2, following the original ResNet [14]. In the third stage, downsampling is performed by the proposed CIR-D unit, which is located at the first block in this stage (4 blocks in total). When the feature map size is downsampled, the number of filters is doubled to increase feature discriminability. The spatial size of the output feature map is 5×5 , with each feature receiving signals from a region of size 93×93 pixels on the input image plane, i.e. the corresponding size of receptive field.

We further increase network depth to 43 layers in building CIResNet-43. Because of its large depth, CIResNet-43 is designed with only 2 stages, to keep its receptive field size within the suggested range. In the second stage of CIResNet-43, there are 14 blocks, where the fourth one has a CIR-D unit for feature downsampling. It is worth noticing that CIResNet-43 almost reaches the maximum depth of backbone networks that can achieve real-time speed in the SiamFC [2] framework. It has 6.07G FLOPs (multiply-adds) and runs at an average of ~ 35 fps in SiamFC framework on a GeForce GTX 1080 GPU.

Wider Networks. We construct two types of wide network architectures using CIR-Inception and CIR-NeXt units, respectively. Here, we only present a 22-layer structure as an example, since other wider networks are similar to this case. As presented in Tab. 3, the wide networks, i.e. CIResInception-22 and CIResNeXt-22, have similar structure to CIResNet-22 in terms of network stride, building block number and output feature size. But the network

widths are increased by 2 and 32 times respectively, through the multi-branch building blocks. Also, the receptive field size becomes diverse (i.e. 13~93) in CIResInception-22 due to multi-branch concatenation, but the maximum size still remain within the suggested range.

4.3. Applications

We apply the designed deeper and wider networks to two representative Siamese trackers: the classical SiamFC [2] and the most recently proposed SiamRPN [20]. In both of these two trackers, we replace the original shallow backbones, i.e. the 5-layer AlexNet [18], with our designed networks, which is the only modification to the original frameworks.

5. Experiments

This section presents the results of our deeper and wider Siamese networks on multiple benchmark datasets, with comparisons to the state-of-the-art tracking algorithms. Ablation studies are also provided to analyze the effects of the components in the proposed networks.

5.1. Experimental Details

Training. The parameters of our networks are initialized with the weights pre-trained on ImageNet [28]. During training, we freeze the weights of the first 7×7 convolution layer, and gradually fine-tune other layers from back to front. We unfreeze the weights of the layers in each block (i.e. the proposed cropping-inside residual units) after every five training epochs. There are 50 epochs in total, the same as in [2, 20]. The learning rates are decreased logarithmically from $10^{-3}/10^{-2}$ to $10^{-7}/10^{-5}$ for SiamFC and SiamRPN, respectively. The weight decay is set to 10^{-4} , and the momentum is set to 0.9 (for both SiamFC and SiamRPN). We use synchronized SGD [19] on 4 GPUs, with each GPU hosting 32 images, hence the mini-batch size is 128 images per iteration.

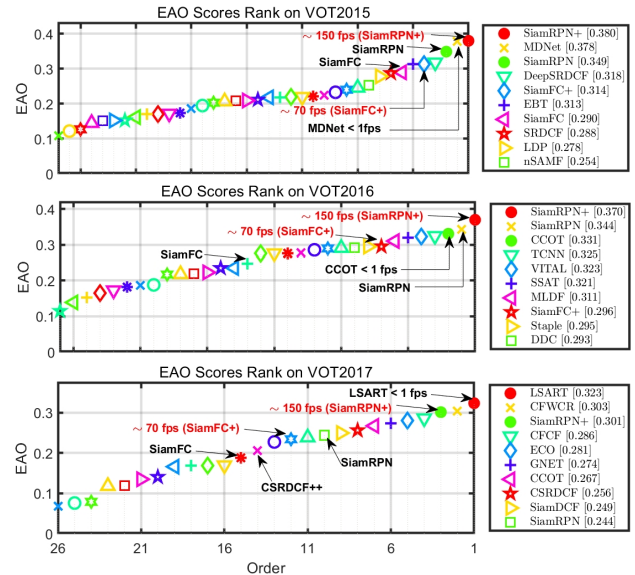
The training image pairs for SiamFC are collected from the ImageNet VID dataset [28], while for SiamRPN, it is generated from ImageNet VID [28] and Youtube-BB [27], which is the same as those in the original frameworks [2, 20]. The size of an exemplar image is 127×127 pixels, while the size of a search image is 255×255 pixels.

Testing. Tracking follows the same protocols as in SiamFC [2] and SiamRPN [20]. The embedding $\varphi_\theta(z)$ of the target object is computed once at the first frame, and then is continuously matched to subsequent search images $\varphi_\theta(x)$. To handle scale variations, SiamFC searches for the object over three scales $1.0482^{\{-1,0,1\}}$ and updates the scale by linear interpolation with a factor of 0.3629 to provide damping. SiamRPN searches over only one scale since it employs proposal refinement to handle scale change. The penalty for large change of proposal size and aspect ratio is

Backbone	OTB(AUC)		VOT-17(EAO)		FPS	
	SiamFC	SiamRPN	SiamFC	SiamRPN	SiamFC	SiamRPN
AlexNet	0.608[2]	0.637[20]	0.188[17]	0.244[20]	101[2]	190[20]
CIResNet-16	0.632	0.651	0.202	0.260	75	160
CIResNet-19	0.640	0.660	0.225	0.279	73	155
CIResNet-22	0.662	0.662	0.234	0.301	70	150
CIResIncep.-22	0.666	0.673	0.215	0.296	67	145
CIResNeXt-22	0.654	0.660	0.230	0.285	72	155
CIResNet-43	0.638	0.652	0.207	0.265	35	75

Table 4: Performance of our network architectures in SiamFC and SiamRPN. To compare with the original results reported in [2, 17, 20], SiamFCs are evaluated on OTB-2013 and VOT-17, while SiamRPNs are evaluated on OTB-2015 and VOT-17. The speed (FPS) is measured on a GeForce GTX 1080 GPU.

Figure 4: Expected average overlap (EAO) plot for VOT-15, 16 and 17. The listed methods, such as EBT[41], LDP[24], nSAMF[22], TCNN[25], MLDF[36], CFWCR[36] and CFCF[11] are compared in VOT challenges [9, 16, 17].



set to 0.439.

Our networks and trackers are implemented using Python 3.6 and PyTorch 0.3.1. The experiments are conducted on a PC with a GeForce GTX 1080 GPU and a Xeon E5 2.4GHz CPU.

5.2. Comparison with Baselines

We first compare our deeper and wider networks to the baseline AlexNet in the SiamFC and SiamRPN frameworks. As presented in Tab. 4, on OTB-13, OTB-15 and VOT-17 datasets, our proposed networks outperform the baseline AlexNet. In particular, SiamFC equipped with a CIResNet.-22 backbone obtains relative improvements of 9.8% (AUC) and 24.4% (EAO) over the original AlexNet on OTB-2015 and VOT-17, respectively. Meanwhile, SiamRPN armed with CIResNet-22 achieves 4.4% and 23.3% relative gains.

Table 5: Performance comparisons on five tracking benchmarks. Red, Green and Blue fonts indicate the top-3 trackers, respectively.

Tracker	Year	OTB-2013		OTB-2015		VOT15			VOT16			VOT17		
		AUC	Prec.	AUC	Prec.	A	R	EAO	A	R	EAO	A	R	EAO
SRDCF [5]	2015	0.63	0.84	0.60	0.80	0.56	1.24	0.29	0.54	0.42	0.25	0.49	0.97	0.12
SINT [34]	2016	0.64	0.85	-	-	-	-	-	-	-	-	-	-	-
Staple [1]	2016	0.60	0.80	0.58	0.78	0.57	1.39	0.30	0.54	0.38	0.30	0.52	0.69	0.17
SiamFC [2]	2016	0.61	0.81	0.58	0.77	0.53	0.88	0.29	0.53	0.46	0.24	0.50	0.59	0.19
ECO-HC [4]	2017	0.65	0.87	0.64	0.86	-	-	-	0.54	0.3	0.32	0.49	0.44	0.24
PTAV [8]	2017	0.66	0.89	0.64	0.85	-	-	-	-	-	-	-	-	-
DSiam [12]	2017	0.64	0.81	-	-	-	-	-	-	-	-	-	-	-
CFNet [35]	2017	0.61	0.80	0.59	0.78	-	-	-	-	-	-	-	-	-
StructSiam [40]	2018	0.64	0.88	0.62	0.85	-	-	-	-	-	0.26	-	-	-
TriSiam [7]	2018	0.62	0.82	0.59	0.78	-	-	-	-	-	-	-	-	0.20
SiamRPN [20]	2018	-	-	0.64	0.85	0.58	1.13	0.35	0.56	0.26	0.34	0.49	0.46	0.24
SiamFC+	Ours	0.67	0.88	0.64	0.85	0.57	-	0.31	0.54	0.38	0.30	0.50	0.49	0.23
SiamRPN+	Ours	0.67	0.92	0.67	0.90	0.59	-	0.38	0.58	0.24	0.37	0.52	0.41	0.30

This verifies that our designed architectures resolve the performance degradation problem shown in Fig. 1. Also, it shows the effectiveness of our proposed CIR units for Siamese networks.

It is worth noting that when the depth of CIResNets increases from 16 to 22 layers, the performance improves accordingly. But when increasing to 43 layers, CIResNet does not obtain further gains. There are two main reasons. 1) The network stride is changed to 4, such that the overlap between the receptive fields of two adjacent features is large. Consequently, it is not as precise as networks with stride of 8 in object localization. 2) The number of output feature channels is halved, compared to the other networks in Tab. 3 (i.e. 256 vs. 512 channels). The overall parameter size is also smaller. These two reasons together limit the performance of CIResNet-43. Furthermore, wider networks also bring gains for Siamese trackers. Though CIResNet-22 contain more transformation branches, its model size is smaller (see Tab. 3). Therefore, its performance is inferior to CIResIncep.-22 and CIResNet-22.

5.3. Comparison to State-of-the-Art Trackers

We further compare our enhanced Siamese trackers to state-of-the-art tracking algorithms. We select some of the currently best performing trackers in general, as well as other recent Siamese trackers for comparison. Our enhanced trackers employ the best performing CIResNet-22 as the backbone, and are denoted as SiamFC+ and SiamRPN+. The comparison is conducted on five datasets: OTB-2013, OTB-2015, VOT15, VOT16 and VOT17.

OTB Benchmarks. The evaluation on OTB-2013 and OTB-2015 follows the standard protocols proposed in [37, 38]. Two metrics, i.e. precision and area under curve (AUC) of success plots, are used to rank the trackers. The results are reported within Tab. 5. It shows that our SiamFC+ and SiamRPN+ surpass other Siamese trackers, such as the recent proposed StructSiam [40] and TriSiam [7]. This demonstrates the effectiveness of our designed architecture. Moreover, compared with other state-of-the-art algorithms, such as ECO-HC [4] and CFNet [35], our trackers are still superior in terms of precision and speed.

Table 6: Ablation for residual unit vs. CIR unit on SiamFC.

	CIResNet-20	CIResNet-22	CIResIncep.-22
Res. Unit	0.204	0.213	0.227
CIR Unit	0.271	0.301	0.282

Table 7: Ablation over different downsampling settings used in SiamFC. See main text for explanations.

	CIResNet-20	CIResNet-22	CIResIncep.-22
Setting 1	0.264	0.292	0.266
Setting 2	0.259	0.287	0.275
CIR-D	0.271	0.301	0.282

VOT Benchmarks. The evaluation on VOT benchmarks is performed by the official toolkit [17], in which accuracy (A), robustness (R) and expected average overlap (EAO) serve as the metrics.

VOT-15. We compare our SiamFC+ and SiamRPN+ with the state-of-the-art trackers on the *vor-2015* challenge. The results are reported in Fig. 4 (top). Our SiamRPN+ achieves the best result, slightly better than MDNet [26]. Also, SiamRPN+ runs much faster than MDNet, which operates at ~ 1 fps. Compared to the baselines, i.e. SiamFC and SiamRPN, our deeper network enhanced trackers obtain 8.8% and 8.9% relative improvements, respectively.

VOT-16. The video sequences in VOT-16 are the same as those in VOT-15, but the ground-truth bounding boxes are precisely re-annotated. We compare our trackers to the top-10 trackers in the challenge. As shown in Fig. 4 (middle), SiamRPN+ ranks first in terms of EAO. It surpasses CCOT [6], the champion of the *vor-2016* challenge, by 3.9 points, as well as the most recent VITAL [31] by 4.7 points. Moreover, SiamFC+ also improves over the baseline by a large margin, i.e. 6.0 points on EAO.

VOT-17. Fig. 4 (bottom) shows the comparison with the trackers in the *vor-2017* challenge. Our SiamRPN+ achieves an EAO of 3.01, slightly inferior to the best performing LSART tracker [32]. However, SiamRPN+ runs at 150 fps, which is 150 \times faster than LSART. Compared to real-time trackers, SiamRPN+ ranks first in terms of accuracy and robustness. Surprisingly, even the plain SiamFC+ surpasses the real-time tracker champion of *vor-2017* CSRDCF++ [23] by 2.2 points. This further verifies the effectiveness of our deeper network architecture designed for Siamese trackers.

5.4. Ablation Study

In Tab. 6–8, we evaluate the effects of different factors in our designed networks on the VOT-16 dataset.

With vs. without CIR unit. The cropping-inside residual unit is a key component in our network architectures. To evaluate its impact, we replace it with the original residual unit [14] in the networks. As shown in Tab. 6, this replacement causes remarkable performance drops, e.g. a degradation of 8.8 points from 0.301 to 0.213 on CIResNet-22. It clearly verifies the importance of padding removal in the CIR unit, which essentially eliminates position bias in learning. The predicted heatmaps of CIResNet-22 in Fig. 2 also prove this point.

With vs. without CIR-D unit. We compare three different downsampling settings in networks: 1) directly using the original downsampling residual unit, i.e. Fig. 3(b), 2) inserting a cropping operation after addition in the downsampling residual unit, and 3) the proposed CIR-D unit, i.e. Fig. 3(b'). Tab. 7 presents the results. It shows that the first two settings are comparable, but inferior to the third one. This indicates our CIR-D unit is effective. In particular, the cropping introduced in the second setting does not bring improvements, since it removes parts of internal features (i.e. not the padding-affected features), resulting in information loss from the original input.

Impact of receptive field, feature size and stride. We tune the sizes of these factors and show their impacts on final performance. Specifically, we vary the convolutional kernel size in the last cropping-inside residual block to change the size of receptive field and output feature. Taking CIResNet-22 as an example, we vary the kernel size from 1 to 6, which causes the feature size to change from 7 to 2. To change the network stride, we replace one CIR unit with a CIR-D unit in the networks. Tab. 8 shows the results. We can observe that when RF becomes large, the performance drops significantly (i.e. from ④ to ① in Tab. 8). The underlying reason is that a large RF covers much image context, resulting in the extracted feature being insensitive to the spatial location of the target object. For output features, it is observed that a small size (OFS ≤ 3) does not benefit accuracy. Also, a large network stride, i.e. 16, is not as precise as a medium sized one, i.e. 8. These results echo our analysis and guidelines presented at the beginning of this paper.

6. Discussions

Network Architectures. The problem addressed in this paper can be seen as a subtask of network architecture design, which mainly develops in two ways: making networks deeper [29, 14] or wider [33, 39]. To make networks deeper, ResNets [14] introduce an identity mapping that makes training ultra deep networks possible. To make networks wider, GoogLeNet [33] and its descendants employ

Table 8: Analysis of network internal factors.

# NUM	①	②	③	④	⑤	⑥	⑦	⑧
RF	+24	+16	+8	± 0 (93)	-8	-16	+16	+16
OFS	2	3	4	5	6	7	6	2
STR	8	8	8	8	8	8	4	16
CIResNet-16	0.22	0.23	0.24	0.26	0.25	0.24	0.23	0.20
CIResNet-19	0.23	0.26	0.26	0.28	0.27	0.26	0.24	0.21
CIResNet-22	0.25	0.27	0.28	0.30	0.29	0.27	0.26	0.23
CIResIncep.-22	0.24	0.26	0.27	0.28	0.27	0.26	0.25	0.22

an Inception module to introduce multiple feature transformations and thus enhance model representation capacity.

Our work takes advantage of these deep and wide network architectures, and modifies them to adapt effectively to Siamese networks. Two key principles for Siamese backbone design are presented. One is to remove padding operations inside network architectures, the other is to control the receptive field size and network stride. Both of them are shown to have significant impact on tracking performance. Moreover, this is the first work that systematically studies how to design robust backbone networks in visual tracking.

Siamese Trackers. Siamese trackers follow a tracking by similarity matching strategy. The pioneering works are SINT [34] and SiamFC [2], which each employ a Siamese network to offline train a similarity metric between the object target and candidate image patches. A large amount of follow-up work [7, 12, 13, 20, 40] have been proposed, and they fall into two camps. One improves matching precision with high-level semantic information or a localization model [13, 15, 40, 20]. The other enhances the offline Siamese model with online updates [12, 35, 42].

There is a recent work also studying how to leverage deep networks for visual tracking [3]. But it approaches the problem from the direction of data augmentation and feature fusion. Differently, our work studies how to design network architectures and successfully equips Siamese trackers with deeper and wider backbones.

7. Conclusion

In this paper, we design deep and wide network architectures for Siamese trackers. This is motivated by the observation that direct replacement of backbones with existing powerful networks does not bring improvements. We carefully study the key causes and determine that receptive field size, network padding and stride are crucial factors. Experiments on five benchmarks demonstrate the effectiveness of the proposed architectures, leading to competitive performance on five datasets.

Acknowledgement. This work was done in Microsoft Research Asia. Thanks to Steve Lin and Zhirong Wu for helpful discussions. Zhipeng Zhang is partly supported by NSFC U1803119, 2016QY01W0106 and JQ18018.

References

- [1] Luca Bertinetto, Jack Valmadre, Stuart Golodetz, Ondrej Miksik, and Philip HS Torr. Staple: Complementary learners for real-time tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1401–1409, 2016. 7
- [2] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *European conference on computer vision*, pages 850–865. Springer, 2016. 1, 2, 4, 5, 6, 7, 8
- [3] Goutam Bhat, Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. Unveiling the power of deep tracking. *arXiv preprint arXiv:1804.06833*, 2018. 8
- [4] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, Michael Felsberg, et al. Eco: Efficient convolution operators for tracking. In *CVPR*, volume 1, page 3, 2017. 7
- [5] Martin Danelljan, Gustav Hager, Fahad Shahbaz Khan, and Michael Felsberg. Learning spatially regularized correlation filters for visual tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4310–4318, 2015. 7
- [6] Martin Danelljan, Andreas Robinson, Fahad Shahbaz Khan, and Michael Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *European Conference on Computer Vision*, pages 472–488. Springer, 2016. 7
- [7] Xingping Dong and Jianbing Shen. Triplet loss in siamese network for object tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 459–474, 2018. 1, 7, 8
- [8] Heng Fan and Haibin Ling. Parallel tracking and verifying: A framework for real-time and high accuracy visual tracking. In *Proc. IEEE Int. Conf. Computer Vision, Venice, Italy*, 2017. 7
- [9] Michael Felsberg, Amanda Berg, Gustav Hager, Jorgen Ahlberg, Matej Kristan, Jiri Matas, Ales Leonardis, Luka Cehovin, Gustavo Fernandez, Tomas Vojir, et al. The thermal infrared visual object tracking vot-tir2015 challenge results. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 76–88, 2015. 6
- [10] Hamed Kiani Galoogahi, Ashton Fagg, and Simon Lucey. Learning background-aware correlation filters for visual tracking. In *ICCV*, pages 1144–1152, 2017. 1
- [11] Erhan Gundogdu and A Aydin Alatan. Good features to correlate for visual tracking. *IEEE Transactions on Image Processing*, 27(5):2526–2540, 2018. 6
- [12] Qing Guo, Wei Feng, Ce Zhou, Rui Huang, Liang Wan, and Song Wang. Learning dynamic siamese network for visual object tracking. In *The IEEE International Conference on Computer Vision (ICCV) (Oct 2017)*, 2017. 1, 2, 7, 8
- [13] Anfeng He, Chong Luo, Xinmei Tian, and Wenjun Zeng. A twofold siamese network for real-time object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4834–4843, 2018. 1, 2, 8
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 2, 4, 5, 8
- [15] David Held, Sebastian Thrun, and Silvio Savarese. Learning to track at 100 fps with deep regression networks. In *European Conference on Computer Vision*, pages 749–765. Springer, 2016. 8
- [16] Matej Kristan, Ale? Leonardis, and Ji?i Matas. et. The visual object tracking vot2016 challenge results. In *Proceedings, European Conference on Computer Vision (ECCV) workshops*, pages 777–823, 8Oct. 2016. 6
- [17] M. Kristan, A. Leonardis, J. Matas, and M. Felsberg. et. The visual object tracking vot2017 challenge results. In *2017 IEEE International Conference on Computer Vision Workshop (ICCVW)*, volume 00, pages 1949–1972, Oct. 2017. 6, 7
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1, 2, 6
- [19] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989. 6
- [20] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8971–8980, 2018. 1, 2, 4, 6, 7, 8
- [21] Xi Li, Weiming Hu, Chunhua Shen, Zhongfei Zhang, Anthony Dick, and Anton Van Den Hengel. A survey of appearance models in visual object tracking. *ACM transactions on Intelligent Systems and Technology (TIST)*, 4(4):58, 2013. 1
- [22] Yang Li and Jianke Zhu. A scale adaptive kernel correlation filter tracker with feature integration. In *European conference on computer vision*, pages 254–265. Springer, 2014. 6
- [23] Alan Lukezic, Tomas Vojir, Luka Cehovin Zajc, Jiri Matas, and Matej Kristan. Discriminative correlation filter with channel and spatial reliability. In *CVPR*, volume 6, page 8, 2017. 7
- [24] Alan Lukežič, Luka Čehovin Zajc, and Matej Kristan. Deformable parts correlation filters for robust visual tracking. *IEEE transactions on cybernetics*, 48(6):1849–1861, 2018. 6
- [25] Hyeonseob Nam, Mooyeol Baek, and Bohyung Han. Modeling and propagating cnns in a tree structure for visual tracking. *arXiv preprint arXiv:1608.07242*, 2016. 6
- [26] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4293–4302, 2016. 7
- [27] Esteban Real, Jonathon Shlens, Stefano Mazzocchi, Xin Pan, and Vincent Vanhoucke. Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In *Computer Vision and Pattern Recogni-*

- tion (CVPR), 2017 IEEE Conference on, pages 7464–7473. IEEE, 2017. 6
- [28] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 6
- [29] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1, 2, 8
- [30] Arnold WM Smeulders, Dung M Chu, Rita Cucchiara, Simone Calderara, Afshin Dehghan, and Mubarak Shah. Visual tracking: An experimental survey. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (1):1, 2013. 1
- [31] Yibing Song, Chao Ma, Xiaohe Wu, Lijun Gong, Linchao Bao, Wangmeng Zuo, Chunhua Shen, Rynson Lau, and Ming-Hsuan Yang. Vital: Visual tracking via adversarial learning. *arXiv preprint arXiv:1804.04273*, 2018. 7
- [32] Chong Sun, Huchuan Lu, and Ming-Hsuan Yang. Learning spatial-aware regressions for visual tracking. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8962–8970, 2018. 7
- [33] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 1, 2, 5, 8
- [34] Ran Tao, Efstratios Gavves, and Arnold WM Smeulders. Siamese instance search for tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1420–1429, 2016. 1, 7, 8
- [35] Jack Valmadre, Luca Bertinetto, João Henriques, Andrea Vedaldi, and Philip HS Torr. End-to-end representation learning for correlation filter based tracking. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5000–5008. IEEE, 2017. 7, 8
- [36] Lijun Wang, Wanli Ouyang, Xiaogang Wang, and Huchuan Lu. Visual tracking with fully convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 3119–3127, 2015. 6
- [37] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Online object tracking: A benchmark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2411–2418, 2013. 7
- [38] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Object tracking benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1834–1848, 2015. 7
- [39] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kai Ming He. Aggregated residual transformations for deep neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5987–5995. IEEE, 2017. 5, 8
- [40] Yunhua Zhang, Lijun Wang, Jinqing Qi, Dong Wang, Mengyang Feng, and Huchuan Lu. Structured siamese network for real-time visual tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 351–366, 2018. 1, 2, 7, 8
- [41] Gao Zhu, Fatih Porikli, and Hongdong Li. Beyond local search: Tracking objects everywhere with instance-specific proposals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 943–951, 2016. 6
- [42] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware siamese networks for visual object tracking. In *European Conference on Computer Vision*, pages 103–119. Springer, 2018. 2, 8
- [43] Wangmeng Zuo, Xiaohe Wu, Liang Lin, Lei Zhang, and Ming-Hsuan Yang. Learning support correlation filters for visual tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 1