

Learning to Filter: Siamese Relation Network for Robust Tracking

Siyuan Cheng^{1,2}, Bineng Zhong^{1*}, Guorong Li³, Xin Liu⁴, Zhenjun Tang¹, Xianxian Li^{1*}, Jing Wang²

¹Guangxi Key Lab of Multi-Source Information Mining & Security,
Guangxi Normal University, Guilin 541004, China

²Department of Computer Science and Technology, Huaqiao University, China

³School of Computer Science and Technology, University of Chinese Academy of Sciences, China,

⁴Seetatech Technology, Beijing, China

siyuancheng@stu.hqu.edu.cn, bnzhong@gxnu.edu.cn, liguorong@ucas.ac.cn

xin.liu@seetatech.com, rrji@xmu.edu.cn, tangzj230@163.com,

lixix@gxnu.edu.cn, roaring@hqu.edu.cn

Abstract

Despite the great success of Siamese-based trackers, their performance under complicated scenarios is still not satisfying, especially when there are distractors. To this end, we propose a novel Siamese relation network, which introduces two efficient modules, i.e. Relation Detector (RD) and Refinement Module (RM). RD performs in a meta-learning way to obtain a learning ability to filter the distractors from the background while RM aims to effectively integrate the proposed RD into the Siamese framework to generate accurate tracking result. Moreover, to further improve the discriminability and robustness of the tracker, we introduce a contrastive training strategy that attempts not only to learn matching the same target but also to learn how to distinguish the different objects. Therefore, our tracker can achieve accurate tracking results when facing background clutters, fast motion, and occlusion. Experimental results on five popular benchmarks, including VOT2018, VOT2019, OTB100, LaSOT, and UAV123, show that the proposed method is effective and can achieve state-of-the-art results. The code will be available at <https://github.com/hqucv/siamrn>

1. Introduction

Visual object tracking, which is the fundamental task in computer vision, has received much attention over the last decades [17, 64, 63, 65]. It aims to capture the position of an arbitrary target accurately by given only its initial state [29]. With the increasing demand for the practical application such as autonomous driving [30], robotics,

*Corresponding author.

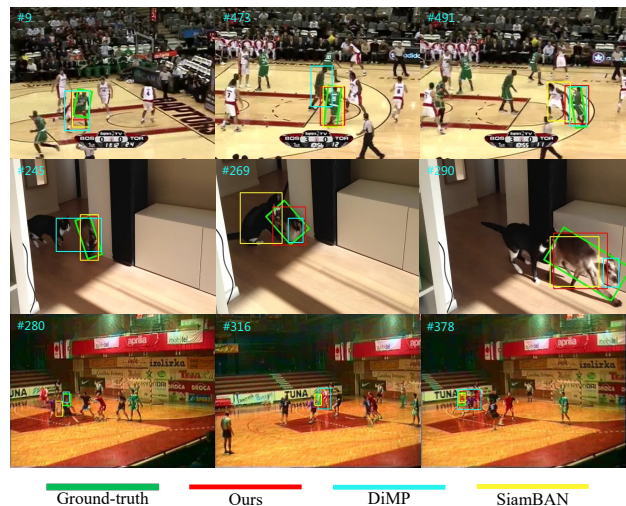


Figure 1. Tracking Result of our Siamese relation network with two state-of-the-art trackers in VOT2018 [22]. Benefiting from our Relation Detector and Refinement Module, when facing similar distractors, appearance change and complex background, our tracker is more robust to these challenges and gets more accurate results. Figure 4 further shows the ability to filtering distractors from the target region.

surveillance [56] and human-computer interaction [33], current trackers require not only accuracy but also speed and robustness to overcome the existing challenges, such as occlusions, fast motions, appearance deformations, illumination change, and background clutters [55], etc.

Owing to the development of CNN and the various network architectures, we can unveil the powerful deep feature for computer vision tasks [24]. In recent years, the Siamese network based trackers [6, 3, 27, 66, 26] have drawn great

attention due to their balanced speed and performance. However, there still exists a gap of robustness between the Siamese trackers and the discriminative trackers that are equipped with an update mechanism [4, 8, 5, 57, 22], which can be manifested in the occasion of tracking failure when encountering similar distractors or background clutters. This is because the training setting of the Siamese trackers is only to match the same target in abundance of image pairs and ignore to distinguish their difference. Therefore, the discriminability for Siamese trackers is deficient when addressing complicated situations. Meanwhile, the target is incompetent to be highlighted due to the Siamese-style cropping strategy, which may introduce the distractive context. Moreover, the classification and regression branches, which are two significant pillars in Siamese-based trackers, are usually optimized independently [6, 26, 27], increasing the probability of mismatch between them during tracking. To be specific, the box corresponding to the position with the highest classification confidence is not the most accurate one for the tracked target.

To address the above issues, we propose a novel Siamese relation network, which introduces two efficient modules named as Relation Detector (RD) and Refinement Module (RM). RD aims to learn the ability to filter the distractors from background via meta-learning. Specifically, we define two categories labeled as target and non-target. Given the training pairs assigned with the label, we perform a two-way one-shot learning, which intends to measure the relationship between them by the nonlinear learnable comparators. Therefore the distractors are filtered out due to low-relevant relationships. The RM is used to integrate the information obtained by RD and the classification branch to refine the tracking results, alleviating the mismatch between classification and regression branches. Furthermore, we equip RD with a contrastive training strategy that attempts not only to learn matching the same target but also pay attention to distinguish the different objects by forming a training triplet and the different combinations of samples, which effectively boosts the discriminant ability of the proposed tracker. Benefiting from the above modules, our tracker possesses remarkable discriminability and robustness when facing background clutters, fast motion, and occlusion, as shown in Figure 1. Experiments (details in Section 4.1) show that our tracker achieves state-of-the-art performance on five popular tracking benchmarks, which confirms the effectiveness of the proposed method.

Our main contributions can be summarized as follows.

- We introduce a novel Relation Detector (RD) that is trained to obtain the ability to filter the distractors from background via few-shot learning based contrastive training strategy. Benefit from RD, during the tracking procedure, our tracker can distinguish the target in the cluttering background once given the initial state

of the target without further fine-tuning.

- To integrate the information obtained by RD and the classification branch to refine the tracking results, we design a Refinement Module, which can jointly operate the classification and regression to localize the target reducing the mismatch between those two branches.
- Our method achieves state-of-the-art results on five popular benchmarks, including VOT2018, VOT2019, OTB100, LaSOT, and UAV123, which confirms the effectiveness of our proposed method.

2. Related Work

2.1. Siamese network based Trackers

Recently, Siamese network based trackers consider tracking as a metric learning problem and have drawn great attention for their good trade-off between speed and accuracy [15, 47, 66, 27, 26, 6]. Bertinetto *et al.* first introduced SiamFC [3] for visual tracking, which aims to learn the similarity metric between the target template and search region and operate a comparison formed as the cross-correlation to localize the target. With the emergence and success of region proposal network (RPN) [38], Li *et al.* [27] applied it into the Siamese networks framework, referred to as SiamRPN, which solved the scale variation of the target. In order to unveil the powerful deep feature extracted by the deep network, SiamRPN++ [26] and SiamDW [60] tackled the challenge brought by introducing the deep network into the Siamese framework and made the performance highly improved. Chen *et al.* [6] borrowed the idea from FCOS [46] into tracking and designed a simple yet effective anchor-free tracker, which can get rid of the intricate parameters of anchor setting. However, these Siamese-based trackers usually localize the target by using the classification and regression branch and optimize them independently, which can arouse the mismatch between them in tracking procedure [21]. Alternatively, we design a Refinement Module that can jointly operate the optimization of these two branches by collaborating with RD output, which balances these two branches to obtain more confident and accurate tracking results.

2.2. Background Distractors in Visual Tracking

The Siamese network based trackers are susceptible to be affected by background distractors, which severely hampers the robustness of tracker. To find out its causes is the less capable discrimination ability. DaSiamRPN [66] finds the inappropriate distribution of semantic and non-semantic negative samples hamper the discriminability of the tracker, so they replace one of the image in training

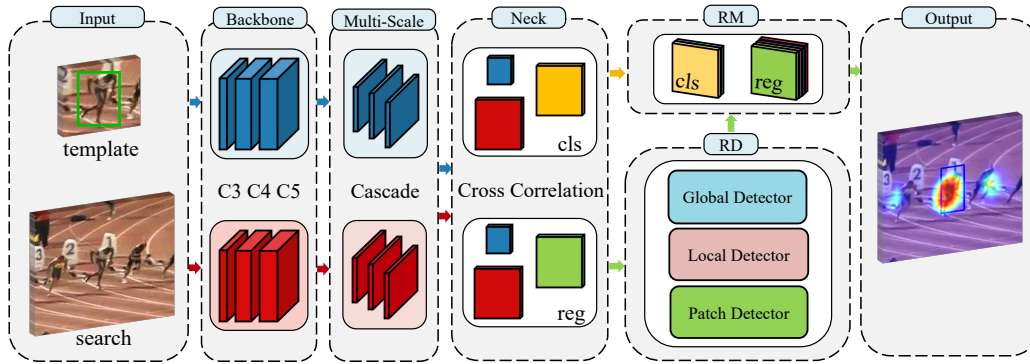


Figure 2. The tracking pipeline of the proposed Siamese Relation Network. The proposed Relation detector (RD) and joint Refinement Module (RM) are presented. During tracking, we feed the features of proposals that are generated from the regression branch by precise ROI pooling [20] into RD to measure the relationships with target-specific feature. Then we convert the output of the RD to a matching score and utilize it in the Refinement Module (RM) to jointly operate the regression and classification branches to predict the target location.

image pairs by negative semantic sample randomly. Recently, some approaches aim to fuse online updating strategy into Siamese framework to pursue more discrimination performance. CFNet [47] interpreted the correlation filter as a differentiable layer to update the model online, and DSiam [15] used a fast transformation module to perform online learning. UpdateNet [59] tried to estimate the target template during tracking to overcome the target appearance variation. However, online update strategy needs heavy computation and the accumulated tracking error will eventually lead to failure. Alternatively, we proposed a Relation Detector (RD), which learns to filter the distractors from background, and further develop a contrastive training strategy that not only to learn matching the same target but also to learn how to distinguish the different objects. Our model is trained offline and can advance the discriminability and robustness of the tracker stably when facing complicated scenarios.

2.3. Few-Shot Learning and Meta Learning

Few-shot learning aims to recognize novel visual categories from very few labeled examples [45]. There is usually only one or very few feasible data during training. Therefore how to overcome this situation to ensure the generalization of the model becomes a challenging task [28]. The mainstream few-shot learning approaches are optimized by recurrent neural network (RNN) [41], further fine-tuned on the target problem [14] or trained to learn an effective metric [36, 44]. A popular trend is to design a general strategy that can guide supervised learning within each task, named as meta-learning [1, 2]. The accumulated meta-knowledge is transferable and the network is able to deal with different tasks. Recently, with the great success of MAML [14], Huang *et al.* [19] and Wang *et al.* [50] used this fine-tune technique to operate fast adaption for

their online tracking task. Inspired by [13, 44], we aim to learn a transferable deep metric for filtering the distractors by measuring the relationships between proposals and the target via the proposed Relation Detector (RD). Unlike [13], which assumes the categories of the objects to be detected are given and focuses on detecting the objects based on given categories of support images. However, there is no concept of category in tracking and it is necessary to determine whether the two objects are the same one.

3. Method

In this section, we first give an overview of our Siamese relation network, then we introduce the Relation Detector in detail, followed by the description of contrast training strategy. Finally, we present our joint Refinement Module.

3.1. Overview

We adopt SiamBAN [6] as our strong baseline which employs ResNet-50 [16] as the backbone for Siamese network. A cascade structure is designed in our network in order to leverage the multi-level features for more accurate prediction. During inference, the input of the network is a pair of images. One image is the first frame of a sequence cropped to size of 127×127 and used as template image. The other one is the subsequent frame cropped to 255×255 and used as search image. After processing them with the backbone we get multi-scale feature maps, which are used for the operation of cross-correlation. Then we utilize Relation Detector (RD), which has learned transferable knowledge based on deep metric, to measure the relationships between target and proposals generated by regression branch. In the end, integrating the outputs of RD and classification results, Refinement Module generates the prediction of the tracking result. Compared to a two-stage tracker like SPM [51], which first generates candidates in the same

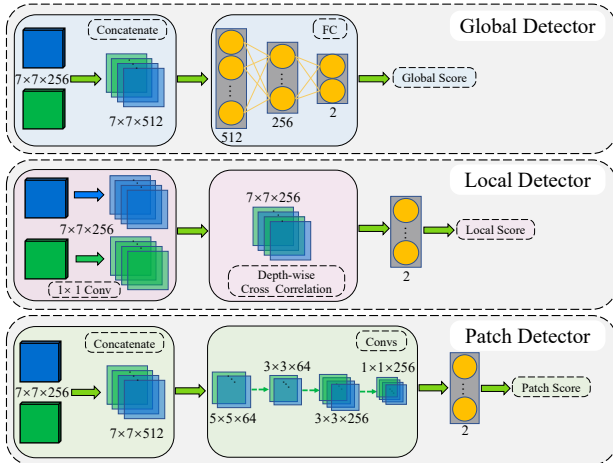


Figure 3. The architecture of Relation Detector (RD). It consists of three different nonlinear comparators, named as Gloabl Detector, Local Detecor and Patch Detector. We measure the relationships between target and ROI (region of interest) by these detectors which is meta learned to filter the distractors. Each detector finally generates a score to measure the relationship of the input pair.

category and then predicts final result within a two-round classification and regression procedure, we only carry out one-round and intend to jointly refine the two branches. The tracking pipeline of our method is shown in Figure 2.

3.2. Relation Detector

The characteristic of the tracking task is that the tracked target can be an arbitrary object, which can be something we have never seen before. A naive way is to calculate the difference between proposals and targets by fixed hand-craft linear comparator [42, 48] (Cosine distance, Euclidean distance, Mahalanobis distance). However, it can be easily invalid when facing indistinguishable distractors, so we propose a novel Relation Detector (RD) equipped with the adaptive nonlinear comparator that has a substantial discriminative ability to filter the distractors by measuring the similarity relationship with the tracked target. Since training such a detector requires samples of the target in the same sequence, which are usually deficient, the conventional training strategy can not work well while the way of few-shot learning can overcome this challenge. Therefore we conduct the few-shot learning to train the network so that our Relation Detector can transfer the meta-knowledge gained during massive diverse few-shot task learning.

The detailed structure of the RD is shown in Figure 3. It consists of three different non-linear learnable comparators, i.e. Global Detector, Local Detector, and Patch Detector. Given the support feature f_s and query proposal features f_q with the size of $7 \times 7 \times C$, the Global Detector aims to compare their global characteristics through deep embed-

ding, while the Local Detector tries to learn a more detailed comparison in pixel-wise and channel-wise. For the Patch Detector, we design it for learning the relationships between different patches.

However, directly applying the ROI features to the RD may cause the problem of feature misalignments. Thus we introduce a self-attention module [53], which can align the pairs of ROI features before putting them into RD. We conduct the experiment to analyze their performance and effectiveness, which will be further discussed in section 4.3.

3.3. Contrastive Training Strategy

Unlike the conventional learning framework, the few-shot learning task has the characteristic of lacking labeled samples in each category [45]. It aims to construct a classifier to assign a label \hat{y} to each sample \hat{x} in the query set, through some known labeled samples which are considered as support set. When the support set contains N different categories, and each category has K labeled samples, we define it as N -way K -shot. In our training, we define two categories, i.e. target and non-target, and conduct our experiments as two-way one-shot learning.

Generation of Contrastive Training Samples. Only matching objects of the same instance is insufficient because the ability to distinguish the different also matters. Thus we exploit the potential relationships of the training samples and construct the training triplet (s_c, q_c, s_n) , where s_c and s_n are positive support and negative support images, and q_c are query images. s_c and q_c are extracted from the same video while s_n is from different video.

During every single shot, we not only match the objects belonging to the target category but also distinguish distractors in non-target class and the model learns to measure the relationships among different sample combinations generated by the input triplet. In detail, we define the ground-truth of positive supports as s_p , and use p_p to represent the positive proposal that generated by s_c and q_c . Similarly, the ground-truth of negative supports are denoted as n_n , and we use p_n for negative proposals that are generated by s_c and q_c . Then we combine them to different pairs, as (s_p, p_p) , (s_p, p_n) , $(n_n, p_p/p_n)$, and keep the ratio as 1:2:1. We adopt MSE loss as the loss function and calculate the matching loss on these formed pairs.

Hard Negative Mining. At the early stage of training, applying easy samples can make the model convergent stably. To further enhance the discriminative power of the model, we introduce the hard negative mining technique in the middle and later stage of training. In detail, we perform the hard negative mining in two different ways (online and offline). For the online ways, we first compute the intersection of union (IoU) between proposals and ground-truth and then remove part of the proposals with IoU value greater than 0.2. Then from the left proposals, we choose the pro-

posals with the highest confidence score as the hard negative samples. For the offline ways, we generate an index table for approximate nearest neighbor queries during the pre-processing of large scale datasets including VID [39], GOT [18] and LaSOT [11]. Then given an image of the tracked object, we use the index table to retrieve nearest neighbors of the tracked object in the embedding space [49] and can find N different hard negative samples.

3.4. Refinement Module

Classification and regression are two pillars in CNN-based trackers, which are usually optimized independently leading to the issue of mismatch between them. To be specific, the box corresponding to the position with the highest classification confidence is not the most accurate, not even the tracked target. We design a refinement module, which effectively integrates the proposed RD into the Siamese framework. At first we utilize the output of the RD and convert it to a matching score sized of $25 \times 25 \times 1$. Next, we conduct an element-wise multiplication between the matching score and cross-correlation feature map of the classification branch, which can filter the distractors in the background by suppressing the false positive position. Then we pass the refined correlation feature through a convolution layer to generate a refined classification confidence score. Figure 4 provides some examples of the obtained confidence map. With the refinement module, we combine the information of the regression branch with the classification branch, and jointly operate them to predict the target location so that the problem of mismatch can be alleviated.

3.5. Ground-truth and Loss function

In this section, we illustrate the loss function used training our model. Let (i, j) denote the point in the feature map, and $r_{i,j}$ is the relation score that to be regressed. (g_{x_c}, g_{y_c}) , (g_{x_0}, g_{y_0}) , (g_{x_1}, g_{y_1}) , (g_w, g_h) represent the center position, left-top position, right-bottom position, and size of the target box respectively.

Classification and Regression Branches. Followed by Chen *et al.* [6], we use the ellipse figure region to design the label. There are two ellipses namely E_1 and E_2 . The center and axes length of E_1 are set to (g_{x_c}, g_{y_c}) and $(\frac{g_w}{2}, \frac{g_h}{2})$ while that of E_2 are set to (g_{x_c}, g_{y_c}) and $(\frac{g_w}{4}, \frac{g_h}{4})$. Thus the two ellipse are,

$$\frac{(p_i - g_{x_c})^2}{(\frac{g_w}{2})^2} + \frac{(p_j - g_{y_c})^2}{(\frac{g_h}{2})^2} = 1, \quad (1)$$

$$\frac{(p_i - g_{x_c})^2}{(\frac{g_w}{4})^2} + \frac{(p_j - g_{y_c})^2}{(\frac{g_h}{4})^2} = 1, \quad (2)$$

For the classification branch, the location (p_i, p_j) falling within the ellipse E_2 is defined as positive, and the one which falls outside the ellipse E_1 is considered as negative,

if (p_i, p_j) falls between the ellipse E_1 and E_2 , we ignore it. Cross-entropy loss is used for calculate the loss function (referred as \mathcal{L}_{cls}) of the classification branch. For the regression branch, the positive location (p_i, p_j) is used to regress target box. The ground truth denoted as $d_{i,j}^l$ for the regression branch is defined as,

$$d_{(i,j)}^l = p_i - g_{x_0}, d_{(i,j)}^t = p_j - g_{y_0}, \quad (3)$$

$$d_{(i,j)}^r = g_{x_1} - p_i, d_{(i,j)}^b = g_{y_1} - p_j. \quad (4)$$

We adopt IoU(Intersection over Union) loss for regression, which is defined as,

$$\mathcal{L}_{reg} = 1 - IoU. \quad (5)$$

Relation Detector. The label $y_{i,j}$ for training Relation Detector depends on training pairs mentioned in section 3.3. We adopt MSE loss for our matching score regression, which is formed as,

$$\mathcal{L}_{matching} = (r_{i,j} - y_{i,j})^2 \quad (6)$$

Our model is set to train end-to-end and the total loss function is defined as,

$$\mathcal{L} = \lambda_1 * \mathcal{L}_{cls} + \lambda_2 * \mathcal{L}_{reg} + \lambda_3 * \mathcal{L}_{matching}, \quad (7)$$

we empirically set the $\lambda_1, \lambda_2, \lambda_3$ as 1 without hyperparameter searching.

3.6. Training and Inference

Training. We train our Siamese relation network on large-scale datasets, including ImageNet VID [39], YouTube-BoundingBoxes [37], COCO [32], ImageNet DET [40], GOT10k [18] and LaSOT [11]. The training input is an image triplet, including a template patch and a search patch extracted from the same sequence with size of 127×127 pixels and 255×255 pixels respectively [3], and a negative search patch which is extracted from another sequence sized as 255×255 . We first select two patches from the same sequence in the triplet, and collect at most 16 positive samples and 48 negative samples on it to train the classification branch and the regression branch [27, 26, 6]. Then, the negative search patch in the triplet is used to generate the training input of our Relation Detector, as mentioned in 3.3. In addition, we start using online hard negative mining at epoch 5 and epoch 15 for offline. Our entire network can be trained end-to-end and doesn't need any further fine-tuning.

Inference. During inference, the patch of the target in the first frame is used as the template and fed it into the backbone to extract the template feature f_z . We cache it during tracking to avoid duplicating computation in subsequent tracking. Besides, we also generate and cache the




Figure 4. Visualization of confidence maps. In left part, the 1st column represents the search images with ground truth during tracking. The 2nd column means the confidence maps without our Relation Detector (RD). The 3rd shows the confidence maps that are equipped with Relation Detector (RD), right part is the same. Our Relation Detector (RD) uses the gained meta knowledge to filter the distractors from the target region and advances the discriminability of our tracker.

ROI feature f_z^{roi} of the template via precise ROI pooling [20]. For subsequent frames, we crop the search patch based on the tracking result in the previous frame and extract its feature denoted as f_x . Then we perform the prediction in the search region to get the regression map $P_{w \times h \times 4}^{reg-all}$ and generate proposals. Next, features of proposals are cropped and concatenated with the cached target ROI feature f_z^{roi} . The obtained features are fed in RD to measure the relationships between proposals and the target. We convert this relationship to a matching score $s_{w \times h \times 1}^{matching}$ and do the element-wise multiplication with the correlation map f_{cls}^{corr} in classification branch. In this way, we fuse the regression result into the classification branch instead of calculating them independently. Then we generate the classification map $P_{w \times h \times 2}^{cls-all}$ through the refined correlation map f_{cls}^{*corr} . Finally, we can get predicted boxes by $P_{w \times h \times 4}^{reg-all}$ and $P_{w \times h \times 2}^{cls-all}$.

4. Experiments

4.1. Implementation Details

We adopt ResNet-50 [16] which is pre-trained on ImageNet [25] as the backbone, and freeze the weights of the first two layers for stable training. We optimize the model with stochastic gradient descent (SGD) and set the batch size to 28. Our model is trained for 20 epochs, and at the first 5 epochs, we use a warm-up strategy to set the learning rate from 0.001 to 0.005. For the last 15 epochs, the learning rate decays exponentially from 0.005 to 0.00005. The weights of the backbone are only released in the last 10 epochs and tuned with the learning rate at a one tenth level of current. Weight decay and momentum are set as 0.0001 and 0.9. Our approach is implemented in Python using Pytorch with 4 NVIDIA TITAN V GPUs.

4.2. Comparison with the state-of-the-art

We evaluate the proposed algorithm against state-of-the-art methods on five tracking benchmarks.

VOT2018 [22] benchmark consists of 60 sequences with different challenging factors. The overall performance of the tracker is evaluated using the EAO (Expected Average Overlap), which takes both accuracy (average overlap during successful tracking periods) and robustness (failure rate) into consideration. The detailed comparisons with the top-performing trackers are reported in Table 1. Among previous approaches, Ocean [61] (offline version) achieves the best EAO value and SiamRPN++ [26] gets the highest accuracy. These two methods are both based on the Siamese network. Compared with DiMP [4], our model achieves a performance gain of 2.9%. When compared to Ocean, our model achieves the same EAO value that is the highest value among the proposed state-of-the-art trackers. In comparison with the baseline tracker SiamBAN [6], we have a substantial improvement of 4.7% in robustness. These results prove that our Siamese Relation Network has learned the strong ability to filter the distractors as expected. Furthermore, we conduct a comparison with the state-of-the-art trackers in term of EAO on different visual attributes and the result is shown in Figure 6. On attributes of camera motion and illumination change our tracker ranks first, and ranks second and third on attributes of occlusion and size change. This shows our tracker are capable to tackle above challenges.

VOT2019 [23] public dataset is one of the most recent datasets for evaluating single object trackers. Compared with VOT2018, the video sequences in VOT2019 have a 20% difference and more challenging videos which contain fast motion and similar distractors. Table 2 shows that our tracker achieves the optimal performance in EAO, which is 1.2% higher than that of Ocean which ranked at second place. In addition, in comparison with the baseline SiamBAN, we achieve a 9.0% improvement of robustness, which proves the strong discriminability of our tracker.

OTB100 [54] is one of the most widely used benchmarks for visual object tracking and consists of 100 well-annotated video sequences. Different from VOT settings, it has a regulation known as one-pass evaluation (OPE) and evaluates the tracker on two basic indications: a precision score and an area under curve (AUC) of success plot. Compared with the top-performing trackers as shown in Figure 8, our model has a good performance in challenging videos and achieves the leading position among state-of-the-art trackers.

LaSOT [11] is a large-scale dataset with 1400 sequences in total. The video sequences in LaSOT have an average sequence length of more than 2500 frames, which is much longer than previous datasets. We conduct one-pass evaluation with success and precision scores to evaluate our tracker. The result is shown in Figure 10. Compared with nine methods [9, 4, 6, 20, 26, 12, 35, 43, 7], our approach achieves the fourth place in terms of AUC and precision without any long-term strategies, which is 1.3% and 1.0%

	ECO	UPDT	SiamRPN	LADCF	ATOM	SiamRPN++	DiMP	SiamBAN	Ocean	Ours
	[8]	[5]	[27]	[57]	[20]	[26]	[4]	[6]	[61]	
EAO \uparrow	0.281	0.379	0.384	0.389	0.401	0.417	0.441	0.452	0.470	0.470
Accuracy \uparrow	0.484	0.576	0.586	0.503	0.590	0.604	0.597	0.597	0.598	0.595
Robustness \downarrow	0.276	0.184	0.276	0.159	0.204	0.234	0.152	0.178	0.169	0.131

Table 1. Detail comparisons on VOT2018 with the state-of-the-art in terms of EAO, accuracy and robustness. The best three results are shown in red, blue and green colors, respectively. DiMP is the ResNet-50 version (DiMP-50), and Ocean is the offline version, the same below.

	MemDTC	SPM	SiamRPN++	SiamMask	ATOM	DCFST	DiMP	SiamBAN	Ocean	Ours
	[58]	[51]	[26]	[52]	[20]	[62]	[4]	[6]	[61]	
EAO \uparrow	0.228	0.275	0.285	0.287	0.292	0.317	0.321	0.327	0.329	0.341
Accuracy \uparrow	0.485	0.577	0.599	0.594	0.603	0.585	0.582	0.602	0.590	0.593
Robustness \downarrow	0.587	0.507	0.482	0.461	0.411	0.376	0.371	0.396	0.376	0.306

Table 2. Detail comparisons on VOT2019 real-time experiments.




Figure 5. Expected averaged overlap performance on VOT2018.

higher than baseline.

UAV123 [34] is an aerial video benchmark that contains 123 sequences captured from low-altitude UAVs. Different from other benchmarks, the perspective of UAV123 is from top to bottom, and the size of the target is relatively small. Compared with 9 state-of-art real-time trackers [4, 20, 26, 6, 66, 8, 10, 31], our tracker ranks second in AUC and first in precision, see in Figure 9.

4.3. Ablation Study

Discussion on Multi-level Prediction and Relation Detection. The multi-level feature contains different information on the target, in order to fully exploit the potential of them, we investigate the impact of layer-wise aggregation. The Relation Detector performs the relation detection on a different level as well to capture the multiple relationships between target and proposals. From the results shown in Table 3 we can find that when only utilizing a single-layer feature, *conv4* performs best. For different combinations of two-level features, we can see that the combination of *conv4* and *conv5* is the best. After aggregating three-level features, our model gets the highest score on AUC.

Discussion on Different Relation Head. We test the performance of the three kinds of relation detectors known as Global Detector, Local Detector, and Patch Detector,




Figure 6. Comparison of EAO on VOT2018 for the following visual attributes: camera motion, illumination change, occlusion, size change, and motion change. Frames that do not correspond to any of the five attributes are marked as unassigned. The values in parentheses indicate the EAO range of each attribute and overall of the trackers.




Figure 7. Expected averaged overlap performance on VOT2019.

which has been mentioned in 3.2. Trying to find out the relation we need most, we conduct the experiments follow-




Figure 8. Success and precision plots on OTB100.




Figure 9. Success and precision plots on UAV123.




Figure 10. Success and precision plots on LaSOT.

ing the 2-way 1-shot setting on OTB100. The results of the ablation study are shown in Table 3. For adopting a single head, the Local Detector gets the best score on AUC. We infer that it can compare the ROI feature most in detail. Combining two types of relation heads achieves better performance than using the single one. The best performance is achieved when utilizing all types of relations heads. It seems that these three kinds of heads are complementary with each other for distinguishing the target from background clutters.

Discussion on Contrastive Training Strategy and Hard Negative Mining. During the few-shot training phase, in order to improve the discriminability of our model, we perform a contrastive training strategy described in section 3.3. when conduction experiments on OTB100, we find that if using the naive strategy (one-way one-shot learning), the model achieves 0.688 in AUC, which drops 1.3% compared to the contrastive way. It seems that our contrastive training strategy not only learns the similarity between tar-

C3	C4	C5	Global	Local	Patch	AUC \uparrow
✓			✓	✓	✓	0.676
	✓		✓	✓	✓	0.686
		✓	✓	✓	✓	0.669
✓	✓		✓	✓	✓	0.691
✓		✓	✓	✓	✓	0.684
	✓	✓	✓	✓	✓	0.693
✓	✓	✓	✓			0.678
✓	✓	✓		✓		0.686
✓	✓	✓			✓	0.672
✓	✓	✓	✓	✓		0.696
✓	✓	✓	✓		✓	0.683
✓	✓	✓	✓	✓	✓	0.693
✓	✓	✓	✓	✓	✓	0.701

Table 3. Quantitative comparison results of our trackers with different level feature combinations and different relation heads combinations on OTB100. C3, C4, C5 represent *conv3*, *conv4*, *conv5*. Global, Local, Patch represent three types of relation heads proposed in 3.2.

get and proposals but also can distinguish the difference between them. Moreover, our proposed hard negative mining technique enhances 0.3% in AUC. It maybe because that adding hard negative samples in time will improve the robustness of the mode learned from massive simple samples.

5. Conclusions

In the paper, we propose a simple yet effective model called Relation Detector (RD) equipped with contrastive training strategy, which is meta trained to gain the ability to learn to filter the distractors from the target region by measuring the relationships between them. Moreover, we design a refinement module to jointly operate the classification and regression process to localize the target, which can alleviate the mismatch of these two branches to get more confident and accurate results. Extensive experiments are conducted on five popular benchmarks and our method obtains the state-of-the-art results with real-time running speed.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 61972167, 61802135), the Project of Guangxi Science and Technology (No. 2020AC19194), the Guangxi “Bagui Scholar” Teams for Innovation and Research Project, the Guangxi Collaborative Innovation Center of Multi-source Information Integration and Intelligent Processing, the Guangxi Talent Highland Project of Big Data Intelligence and Application, and the Open Project Program of the National Laboratory of Pattern Recognition (NLPR) (No. 202000012).

References

- [1] Proceedings of the 34th international conference on machine learning, ICML 2017, sydney, nsw, australia, 6-11 august 2017. volume 70 of *Proceedings of Machine Learning Research*. PMLR, 2017. [3](#)
- [2] Proceedings of the 35th international conference on machine learning, ICML 2018, stockholmsmässan, stockholm, sweden, july 10-15, 2018. volume 80 of *Proceedings of Machine Learning Research*. PMLR, 2018. [3](#)
- [3] Luca Bertinetto, Jack Valmadre, João F. Henriques, Andrea Vedaldi, and Philip H. S. Torr. Fully-convolutional siamese networks for object tracking. In *ECCV Workshops (2)*, volume 9914 of *Lecture Notes in Computer Science*, pages 850–865, 2016. [1](#), [2](#), [5](#)
- [4] Goutam Bhat, Martin Danelljan, and Luc Van Gool. Learning discriminative model prediction for tracking. In *ICCV*, pages 6181–6190. IEEE, 2019. [2](#), [6](#), [7](#)
- [5] Goutam Bhat, Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. Unveiling the power of deep tracking. In *ECCV (2)*, volume 11206 of *Lecture Notes in Computer Science*, pages 493–509. Springer, 2018. [2](#), [7](#)
- [6] Zedu Chen, Bineng Zhong, Guorong Li, Shengping Zhang, and Rongrong Ji. Siamese box adaptive network for visual tracking. In *CVPR*, pages 6667–6676. IEEE, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [7] Kenan Dai, Yunhua Zhang, Dong Wang, Jianhua Li, Huchuan Lu, and Xiaoyun Yang. High-performance long-term tracking with meta-updater. In *CVPR*, pages 6297–6306. IEEE, 2020. [6](#)
- [8] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. ECO: efficient convolution operators for tracking. In *CVPR*, pages 6931–6939. IEEE Computer Society, 2017. [2](#), [7](#)
- [9] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *CVPR*, pages 7181–7190. IEEE, 2020. [6](#)
- [10] Martin Danelljan, Gustav Häger, Fahad Shahbaz Khan, and Michael Felsberg. Learning spatially regularized correlation filters for visual tracking. In *ICCV*, pages 4310–4318. IEEE Computer Society, 2015. [7](#)
- [11] Heng Fan, Liting Lin, and Fan Yang. Lasot: A high-quality benchmark for large-scale single object tracking. In *CVPR*, pages 5374–5383. Computer Vision Foundation / IEEE, 2019. [5](#), [6](#)
- [12] Heng Fan and Haibin Ling. Siamese cascaded region proposal networks for real-time visual tracking. In *CVPR*, pages 7952–7961. Computer Vision Foundation / IEEE, 2019. [6](#)
- [13] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. In *CVPR*, pages 4012–4021. IEEE, 2020. [3](#)
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR, 2017. [3](#)
- [15] Qing Guo, Wei Feng, Ce Zhou, Rui Huang, Liang Wan, and Song Wang. Learning dynamic siamese network for visual object tracking. In *ICCV*, pages 1781–1789. IEEE Computer Society, 2017. [2](#), [3](#)
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778. IEEE Computer Society, 2016. [3](#), [6](#)
- [17] João F. Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(3):583–596, 2015. [1](#)
- [18] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *CoRR*, abs/1810.11981, 2018. [5](#)
- [19] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Bridging the gap between detection and tracking: A unified approach. In *ICCV*, pages 3998–4008. IEEE, 2019. [3](#)
- [20] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yunqing Jiang. Acquisition of localization confidence for accurate object detection. In *ECCV (14)*, volume 11218 of *Lecture Notes in Computer Science*, pages 816–832. Springer, 2018. [3](#), [6](#), [7](#)
- [21] Wei Ke, Tianliang Zhang, Zeyi Huang, Qixiang Ye, Jianzhuang Liu, and Dong Huang. Multiple anchor learning for visual object detection. In *CVPR*, pages 10203–10212. IEEE, 2020. [2](#)
- [22] Matej Kristan, Ales Leonardis, and Jiri Matas. The sixth visual object tracking VOT2018 challenge results. In *ECCV Workshops (1)*, volume 11129 of *Lecture Notes in Computer Science*, pages 3–53. Springer, 2018. [1](#), [2](#), [6](#)
- [23] Leonardis A et al. Kristan M, Matas J. The seventh visual object tracking VOT2019 challenge results. In *ICCV Workshops*, pages 2206–2241. IEEE, 2019. [6](#)
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1106–1114, 2012. [1](#)
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1106–1114, 2012. [6](#)
- [26] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *CVPR*, pages 4282–4291. Computer Vision Foundation / IEEE, 2019. [1](#), [2](#), [5](#), [6](#), [7](#)
- [27] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *CVPR*, pages 8971–8980. IEEE Computer Society, 2018. [1](#), [2](#), [5](#), [7](#)
- [28] Fei-Fei Li, Robert Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(4):594–611, 2006. [3](#)
- [29] Xi Li, Weiming Hu, Chunhua Shen, Zhongfei Zhang, Anthony R. Dick, and Anton van den Hengel. A survey of appearance models in visual object tracking. *ACM Trans. Intell. Syst. Technol.*, 4(4):58:1–58:48, 2013. [1](#)
- [30] Xi Li, Weiming Hu, Chunhua Shen, Zhongfei Zhang, Anthony R. Dick, and Anton van den Hengel. A survey of appearance models in visual object tracking. *ACM Trans. Intell. Syst. Technol.*, 4(4):58:1–58:48, 2013. [1](#)

- [31] Yang Li and Jianke Zhu. A scale adaptive kernel correlation filter tracker with feature integration. In *ECCV Workshops (2)*, volume 8926 of *Lecture Notes in Computer Science*, pages 254–265. Springer, 2014. 7
- [32] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV (5)*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer, 2014. 5
- [33] Liwei Liu, Junliang Xing, Haizhou Ai, and Xiang Ruan. Hand posture recognition using finger geometric feature. In *ICPR*, pages 565–568. IEEE Computer Society, 2012. 1
- [34] Matthias Mueller, Neil Smith, and Bernard Ghanem. A benchmark and simulator for UAV tracking. In *ECCV (1)*, volume 9905 of *Lecture Notes in Computer Science*, pages 445–461. Springer, 2016. 7
- [35] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, pages 4293–4302. IEEE Computer Society, 2016. 6
- [36] Boris N. Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. TADAM: task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, pages 719–729, 2018. 3
- [37] 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 *CVPR*, pages 7464–7473. IEEE Computer Society, 2017. 5
- [38] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015. 2
- [39] Olga Russakovsky, Jia Deng, and Hao Su. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015. 5
- [40] Olga Russakovsky, Jia Deng, and Hao Su. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015. 5
- [41] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy P. Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1842–1850. JMLR.org, 2016. 3
- [42] Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learning. In *NIPS*, pages 4077–4087, 2017. 4
- [43] Yibing Song, Chao Ma, and Xiaohe Wu. VITAL: visual tracking via adversarial learning. In *CVPR*, pages 8990–8999. IEEE Computer Society, 2018. 6
- [44] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208. IEEE Computer Society, 2018. 3
- [45] Sebastian Thrun. Is learning the n-th thing any easier than learning the first? In *NIPS*, pages 640–646. MIT Press, 1995. 3, 4
- [46] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: fully convolutional one-stage object detection. In *ICCV*, pages 9626–9635. IEEE, 2019. 2
- [47] Jack Valmadre, Luca Bertinetto, João F. Henriques, Andrea Vedaldi, and Philip H. S. Torr. End-to-end representation learning for correlation filter based tracking. In *CVPR*, pages 5000–5008. IEEE Computer Society, 2017. 2, 3
- [48] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *NIPS*, pages 3630–3638, 2016. 4
- [49] Paul Voigtlaender, Jonathon Luiten, and Philip H. S. Torr. Siam R-CNN: visual tracking by re-detection. In *CVPR*, pages 6577–6587. IEEE, 2020. 5
- [50] Guangting Wang, Chong Luo, Xiaoyan Sun, Zhiwei Xiong, and Wenjun Zeng. Tracking by instance detection: A meta-learning approach. In *CVPR*, pages 6287–6296. IEEE, 2020. 3
- [51] Guangting Wang, Chong Luo, and Zhiwei Xiong. Spm-tracker: Series-parallel matching for real-time visual object tracking. In *CVPR*, pages 3643–3652. Computer Vision Foundation / IEEE, 2019. 3, 7
- [52] Qiang Wang, Li Zhang, and Luca Bertinetto. Fast online object tracking and segmentation: A unifying approach. In *CVPR*, pages 1328–1338. Computer Vision Foundation / IEEE, 2019. 7
- [53] Xiaolong Wang, Ross B. Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803. IEEE Computer Society, 2018. 4
- [54] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Online object tracking: A benchmark. In *CVPR*, pages 2411–2418. IEEE Computer Society, 2013. 6
- [55] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Object tracking benchmark. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(9):1834–1848, 2015. 1
- [56] Ruiyue Xu, Yepeng Guan, and Yizhen Huang. Multiple human detection and tracking based on head detection for real-time video surveillance. *Multim. Tools Appl.*, 74(3):729–742, 2015. 1
- [57] Tianyang Xu, Zhen-Hua Feng, Xiao-Jun Wu, and Josef Kittler. Learning adaptive discriminative correlation filters via temporal consistency preserving spatial feature selection for robust visual object tracking. *IEEE Trans. Image Process.*, 28(11):5596–5609, 2019. 2, 7
- [58] Tianyu Yang and Antoni B. Chan. Learning dynamic memory networks for object tracking. In *ECCV (9)*, volume 11213 of *Lecture Notes in Computer Science*, pages 153–169. Springer, 2018. 7
- [59] Lichao Zhang, Abel Gonzalez-Garcia, Joost van de Weijer, Martin Danelljan, and Fahad Shahbaz Khan. Learning the model update for siamese trackers. In *ICCV*, pages 4009–4018. IEEE, 2019. 3
- [60] Zhipeng Zhang and Houwen Peng. Deeper and wider siamese networks for real-time visual tracking. In *CVPR*, pages 4591–4600. Computer Vision Foundation / IEEE, 2019. 2
- [61] Zhipeng Zhang, Houwen Peng, and Jianlong Fu. Ocean: Object-aware anchor-free tracking. In *ECCV (21)*, volume 12366 of *Lecture Notes in Computer Science*, pages 771–787. Springer, 2020. 6, 7

- [62] Linyu Zheng, Ming Tang, Yingying Chen, Jinqiao Wang, and Hanqing Lu. Learning feature embeddings for discriminant model based tracking, 2020. [7](#)
- [63] Bineng Zhong, Bing Bai, Jun Li, Yulun Zhang, and Yun Fu. Hierarchical tracking by reinforcement learning-based searching and coarse-to-fine verifying. *IEEE Trans. Image Process.*, 28(5):2331–2341, 2019. [1](#)
- [64] Bineng Zhong, Hongxun Yao, Sheng Chen, Rongrong Ji, Tat-Jun Chin, and Hanzi Wang. Visual tracking via weakly supervised learning from multiple imperfect oracles. *Pattern Recognit.*, 47(3):1395–1410, 2014. [1](#)
- [65] Qinqin Zhou, Bineng Zhong, Yulun Zhang, Jun Li, and Yun Fu. Deep alignment network based multi-person tracking with occlusion and motion reasoning. *IEEE Trans. Multim.*, 21(5):1183–1194, 2019. [1](#)
- [66] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware siamese networks for visual object tracking. In *ECCV (9)*, volume 11213 of *Lecture Notes in Computer Science*, pages 103–119. Springer, 2018. [1, 2, 7](#)