

Siamese Instance Search for Tracking

Ran Tao, Efstratios Gavves, Arnold W.M. Smeulders
QUVA Lab, Science Park 904, 1098 XH Amsterdam
{r.tao, egavves, a.w.m.smeulders}@uva.nl

Abstract

In this paper we present a tracker, which is radically different from state-of-the-art trackers: we apply no model updating, no occlusion detection, no combination of trackers, no geometric matching, and still deliver state-of-the-art tracking performance, as demonstrated on the popular online tracking benchmark (OTB) and six very challenging YouTube videos. The presented tracker simply matches the initial patch of the target in the first frame with candidates in a new frame and returns the most similar patch by a learned matching function. The strength of the matching function comes from being extensively trained generically, i.e., without any data of the target, using a Siamese deep neural network, which we design for tracking. Once learned, the matching function is used as is, without any adapting, to track previously unseen targets. It turns out that the learned matching function is so powerful that a simple tracker built upon it, coined Siamese Instance search Tracker, SINT, which only uses the original observation of the target from the first frame, suffices to reach state-of-the-art performance. Further, we show the proposed tracker even allows for target re-identification after the target was absent for a complete video shot.

1. Introduction

At the core of many tracking algorithms is the function by which the image of the target is matched to the incoming frames. The matching function for tracking ideally provides good matching even if the target in the video is occluded, changes its scale, rotates in and out-of-plane or, undergoes uneven illumination, camera motion and other disturbing factors [41, 52]. One way to proceed is to model each of these distortions explicitly in the matching by introducing affine transformations [29], probabilistic matching [6], eigen images [38], illumination invariants [32], occlusion detection [33]. While one explicit matching mechanism may be well-fitted to solve one distortion, it is likely to disturb another.

In this work, rather than explicitly modeling the match-

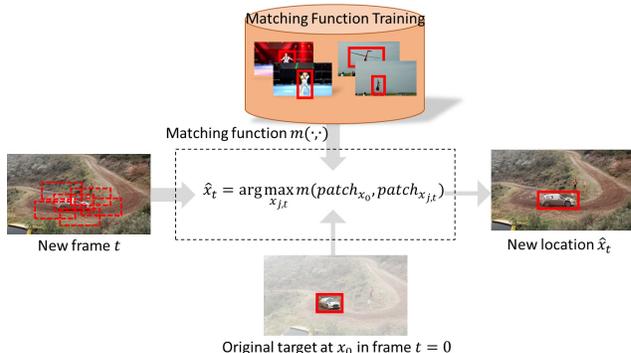


Figure 1: The tracker simply finds the patch that matches best to the original patch of the target in the first frame, using a learned matching function. The matching function is learned once on a rich video dataset. Once it has been learned, it is applied as is, without any adapting, to new videos of previously unseen target objects. We *do not* apply offline target learning and the target is not included in the training video dataset.

ing for particular distortions, we propose to *learn* the matching mechanism. More specifically, we suggest that we learn from external videos that contain various disturbing factors the invariances without, however, explicitly modeling these invariances. If the set of external videos is sufficiently large, the goal is to learn a generically applicable matching function *a priori*. We take extra care that there is absolutely no overlap between the videos we use for training and any of the tracking videos for evaluation. Namely, we *do not* aim to do any offline learning of the tracking targets, since in that case we would essentially learn an object detector. Instead, in the learning we focus on the generic set of object appearance variations in videos. In this way, we optimize the matching function between an arbitrary target and patches from subsequent frames. Once the matching function has been learnt on the external data we do not adapt it anymore and, we apply it as is to new tracking videos of previously unseen target objects.

We focus on learning the matching function suited for

application in trackers. Hence, our aim is not to build a fully fledged tracker which might need explicit occlusion detection [34], model updating [56, 16, 14], tracker combination [56], forget mechanisms [32, 14] and other. We rather focus on the matching function alone, similar to the simplicity of the normalized cross-correlation (NCC) tracker [7, 2]. In this paper, we simply match the initial target in the first frame with the candidates in a new frame and return the most similar one by the learnt matching function, without updating the target, tracker combination, occlusion detection and alike. Figure 1 illustrates the tracking algorithm.

This approach to tracking bears some similarity to instance search [43, 44, 36, 46], where the target specified in the query image is searched for in a pile of images. Introducing matching learning [44] allows for accurate instance search of generic objects even when the relevant images in the search set show drastically different views of the target object from the query image. Here we intend to learn a generic matching function to cope with all sorts of appearance variations from tracking examples. After learning, the matching function is capable of comparing patches recorded under very different conditions for new objects, or, even for new object types that the function has not seen before.

We summarize the contributions of the work as follows. First, we propose to learn a generic matching function for tracking, from external video data, to robustly handle the common appearance variations an object can undergo in video sequences. The learnt function can be applied as is, without any adapting, to new tracking videos of previously unseen target objects. Second, on the basis of the learnt generic matching function, we present a tracker, which reaches state-of-the-art tracking performance. The presented tracker is radically different from state-of-the-art trackers. We apply no model updating, no occlusion detection, no combination of trackers, no geometric matching and alike. In each frame, the tracker simply finds the candidate patch that matches best to the initial patch of the target in the first frame by the learned matching function. Third, to learn the matching function, we use a two-stream Siamese network [3], which we design specifically for tracking. Further, in the absence of any drifting that one would expect by on-the-fly model updating, the proposed tracker allows for successful target object re-identification after the target was absent for a long period of time, *e.g.*, a complete shot.

2. Related Work

Matching functions in tracking One of the most basic concept of tracking is the direct matching between the intensity values of the target patch and the patches taken from the incoming image. The oldest tracking algorithm does just that by normalized cross-correlation [7, 2]. Its simplicity is also its strength, still being in use as part of the TLD-tracker [23]. Subsequent trackers have reconsid-

ered the matching function by focusing on the various distortions to the target image faced in tracking. The Lucas and Kanade tracker [29] adds an affine transformation to the matching function. MST [6] relies on probabilistic matching. FRT [1] uses the earth mover’s distance matching. And IVT [38] matches by the metric of eigen images obtained during tracking. LIT [30] is successful with L1-metric matching on graphs of fragments. SPT [50] uses super-pixels for matching, HBT [12] uses HOG-features in a probabilistic approach, and FBT [32] uses color invariants for robustness against illumination variations. Different from all methods above, which pursue explicit modeling of the matching function, this paper aims to *learn* the matching function from example videos annotated with the correct boxes.

Recent tracking methods In recognition of the hardness of the task, composite trackers have been introduced. TLD [23] integrates the NCC matching for recovery with a differential tracker and a complex updating model. Struck [14] is based on structural SVM with the displacement as the continuous output, with a cautious update mechanism. More recently, MEEM [56] successfully learns and updates a discriminative tracker, keeping a set of historical snapshots as experts who derive the per frame prediction based on an entropy regularized optimization. Alien [34] is a successful long-term tracker relying on oversampling of local features and RANSAC-based geometric matching. In the very recent MUSTer [18] one component stores short-term memories of the target for short-term tracking using integrated correlation filters, where the long-term memory is based on RANSAC matching again. Finally, the AND-OR tracker [51] proposes a discriminative learning of hierarchical, compositional and-or graphs that account for the appearance and structural variations of the object. In this paper, we focus on simple tracking inference scheme, namely finding the patch that matches best to the initial target in the first frame. The complexity, instead, is incorporated externally, where the matching function is trained to be robust against appearance variations. Hence, rather than learning on-the-fly, we learn what can be encountered *in general* without requiring target-specific learning. Once learned, the matching function can be built in the successful, aforementioned composite trackers to enhance their performance.

Deep learning in tracking [49] uses a stacked denoising autoencoder to learn tracking features. The features are performing poorly, however. [25] learns a target classifier online, which is fundamentally hampered by a lack of data. [17] focuses on learning target-specific saliency map using pre-trained ImageNet network. [47] pre-trains a convolutional neural network for measuring generic objectness on ImageNet 2014 detection set and adapts the network online to predict the target-specific objectness. Compared to previous works, this work focuses on a different part

of a tracker. We employ deep neural networks to learn a generic matching function from rich external data to compare patches for use in tracking.

Instance Search Instance search from one example, also known as particular object retrieval, is related to object tracking, especially when localized [22, 43]. The most popular paradigm is based on matching local image descriptors between the query and the candidate image [40, 36, 20, 37, 45, 43] and is especially accurate for buildings [36]. Recently, [44] proposed to learn a robust representation for instance search of less textured, more generic objects, showing good accuracy despite the significant appearance changes between the query and the candidate images. We derive some inspiration from [44]. We propose to learn a robust matching function for matching arbitrary, generic objects that may undergo all sorts of appearance variations. We focus, however, on tracking. Instead of focusing on a specific category *e.g.*, shoes, and learning from images with a white background [44], we learn in this work a universal matching model suited for tracking that applies to all categories and all realistic imaging conditions.

Siamese architecture [3] proposes the two-stream Siamese architecture for signature verification. Later, the two-stream network architecture has been applied to face verification [5, 42], ground-to-aerial image matching [27], local patch descriptor learning [13, 53] and stereo matching [54]. In this work, we design a Siamese network-architecture to learn robust and generic representation for object tracking, aiming to be invariant to all sorts of appearance variations in practical tracking scenarios.

Fast localization Tracking also bears resemblance to the object localization problem. Usually, it requires efficient processing of multiple regions in one frame. [26] proposes efficient region computation by reordering the encoding, pooling and classification steps for the ‘shallow’ representations such as Fisher vector [35]. Recent work by Girshick [10] proposes an efficient way of processing multiple regions in one single pass through the deep neural network for fast object detection. Inspired by [10], we incorporate the region-of-interest pooling layer into our network for fast processing of multiple regions in one frame for tracking.

3. Siamese Instance Search Tracker

In the following we describe the proposed method for tracking, which is coined Siamese INstance search Tracker, SINT for abbreviation. We first present the matching function, which is the core of the tracker. Then we describe the simple online tracking inference.

3.1. Matching Function

To learn a matching function robust to all sorts of distortions as described earlier, we need a model that operates on pairs of data, (x_j, x_k) . A network architecture that has been

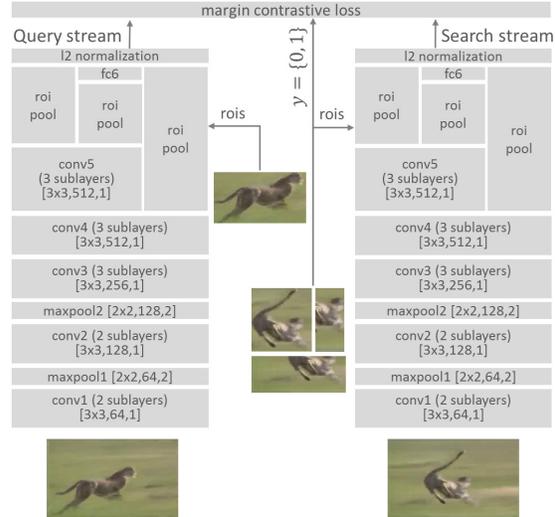


Figure 2: The proposed Siamese invariance network to learn the generic matching function for tracking. ‘conv’, ‘maxpool’, ‘roi pool’ and ‘fc’ stand for convolution, max pooling, region-of-interest pooling and fully connected layers respectively. Numbers in square brackets are kernel size, number of outputs and stride. The fully connected layer has 4096 units. All conv layers are followed by rectified linear units (ReLU) [31].

successfully shown to work well on pairs of data is the two-stream Siamese architecture [3, 5]. A Siamese architecture builds on top of convolutional networks. Next, we analyze the different components of the proposed two-stream network which we coin *Siamese Invariance Network*.

Network architecture We use a Siamese architecture composed of two branches. The Siamese network processes the two inputs separately, through individual networks that usually take the form of a convolutional neural network. For individual branches, we design and compare two different network architectures, a small one similar to AlexNet [24] and a very deep one inspired by VGGNet [39] (Figure 2)¹. In the following we highlight the distinctive designs of the networks as compared to the successful AlexNet and VGGNet.

Being largely a localization task the tracking problem is naturally susceptible to rough discretizations. Aiming for precise localization, we design our network with very few maxing pooling layers, fewer than the networks in [24, 39]. Indeed, as max pooling maintains only the strongest of the activations from a local neighborhood to use as input for the subsequent layers, the spatial resolution of the activations is aggressively reduced, at the very least by 50% only

¹Due to the space limit, only the very deep network is shown here. We put the illustration of the other AlexNet-like network in the supplementary material.

in the simple case of 2×2 local neighborhoods. An advantage of max pooling is it introduces invariance to local deformations. However, this is more important for object categorization, where the objects vary a lot in appearance. In tracking even if the target object changes its appearance over time, it still remains the same object in all frames. Moreover, it is important to be able to follow the small appearance changes, such as local deformations, of the object over time. Regarding the two architectures we propose, for the AlexNet-like small net we do not include any max pooling layer, while for the VGG-like large net, we only have two max pooling at the very early stage (see Figure 2), as the lower level layers learn filters of very small receptive fields and their max pooling layers are important to maintain robustness to local noise.

In tracking one typically needs to evaluate hundreds of candidate regions for the next frame. Although one can simply pass through the candidate regions independently, this would lead to a severe computation overhead, especially since there is a significant overlap between the candidate regions. Therefore, we employ a region pooling layer [10] for the fast processing of multiple overlapping regions. Each branch of the Siamese architecture takes as input one image and a set of bounding box regions. The network first processes the entire image for a few layers, then the region pooling layer converts the feature map from a particular region into a fixed-length representation. Having a fixed length representation, one can now proceed to the subsequent layers.

The layers in a deep network capture progressively more abstract representations [55]. Typically, the filters of the lower layers get activated the most on lower level visual patterns, such as edges and angles, whereas higher layers get activated the most on more complex patterns, such as faces and wheels. Also, the deeper one layer is, the more invariant it is to appearance changes but also less discriminative, especially for instance-level distinction. In tracking we do not know the type of target object we want to track, whether it is highly textured with rich low level patterns or not. We do not know either the complexity of the background, whether there are confusing objects in which case higher discrimination would probably be more helpful. For this reason we propose to use the outputs from multiple layers as the intermediate representation that is then fed to the loss function. Similar observations have also been made in [28, 15] for different tasks, semantic segmentation and fine-grained localization specifically. All activations are pooled using the region pooling layers.

Given that modern convolutional neural networks use rectified linear units that do not bound the output values, the nonlinear activations can vary a lot in the range of values they produce. As such and without considerations, the network output and the loss function will be heavily influ-

enced by the scale of the generated features and not their representation quality. To avoid this we propose to add an ℓ_2 normalization layer before the loss layer. The normalization layer is applied on each of the layer activations that are fed to the loss layer and has the property of maintaining the direction of the feature, while forcing features from different scales to lie on the same unit sphere.

Compared to standard convolutional neural networks, AlexNet and VGGNet [24, 39], our architecture has several differences, highlighted above. However, we also explicitly design our networks to be compatible to AlexNet and VGGNet. In this way, we are able to initialize the weights of our networks using the ImageNet-pretrained AlexNet and VGGNet to avoid training from scratch, something that would likely lead to overfitting. Last, note that we keep the parameters of the two convolutional network branches tied together, as there would be an increased danger of overfitting otherwise.

Network input Our training data consist of videos of objects, whose bounding box location is provided to us. To emulate the instance search paradigm and to avoid confusion, we coin the first stream of our network as query stream, whereas the second stream of our network as search stream. For the query stream we randomly pick one frame from the video and use the annotated patch of the target. Since we want to be robust to as many types of variations that we might face when tracking novel objects as possible, for the search stream we randomly pick another video frame that does not need to be adjacent to the frame of the query stream. From the frame of the search stream we sample boxes and the ones that overlap more than ρ_+ with the ground truth are deemed positives, while the ones that overlap less than ρ_- with the ground truth are deemed negatives. From these we form positive and negative pairs of data that we use for the training.

Loss In the end, the two branches in the Siamese Invariance Network are connected with a single loss layer. For tracking we want the network to generate feature representations, that are close by enough for positive pairs, whereas they are far away at least by a minimum for negative pairs. Bearing these requirements in mind and inspired by [5], we employ the margin contrastive loss

$$\mathcal{L}(x_j, x_k, y_{jk}) = \frac{1}{2}y_{jk}D^2 + \frac{1}{2}(1 - y_{jk})\max(0, \epsilon - D^2), \quad (1)$$

where $D = \|f(x_j) - f(x_k)\|_2$ is the Euclidean distance of two ℓ_2 -normalized latent representations, $y_{jk} \in \{0, 1\}$ indicates whether x_j and x_k are the same object or not, and ϵ is the minimum distance margin that pairs depicting different objects should satisfy.

Data As tracking is an inherently online task, where no training data related to the target object are available, it is important to emphasize that *the network is learnt on external videos that do not appear in the tracking evaluation sets*. The data should be varying enough, covering a good amount of semantics and not focus on particular objects, otherwise the tuned network parameters will overfit to particular object categories. Furthermore, as we do not explicitly learn types of invariances, namely we do not learn “illumination invariance” separately from “scale invariance”, therefore in the external data we do not need any specific variation labels. The only requirement is the box annotations within the video following a particular object.

3.2. Tracking Inference

Once we have completed the learning of the matching function, we are ready to deploy it as is to tracking, without any further adapting. We propose a simple tracking strategy. As the only reliable data we have for the target object is its location at the first frame, at each frame we compare the sampled candidate boxes with the target object at the first frame. We pass all the candidate boxes from the search stream of our network and pick the candidate box that matches best to the original target,

$$\hat{x}_t = \arg \max_{x_{j,t}} m(x_{t=0}, x_{j,t}), \quad (2)$$

where $x_{j,t}$ are all the candidate boxes at frame t , m is the learned matching function, $m(x, y) = f(x)^T f(y)$.

Candidate sampling We employ the radius sampling strategy [14]. More specifically, around the predicted location of the previous frame we sample locations evenly on circles of different radii. Different from [14], to handle scale variations we generate at each sampled location multiple candidate boxes at different scales.

Box refinement Provided that the box prediction is accurate enough, [8, 11] showed that a refinement step of the boxes can improve localization accuracy significantly. To this end we adopt their strategy and refine at each frame the predicted bounding box further.

As in [11] we train four Ridge regressors for the (x, y) coordinates of the box center and the width and height (w, h) of the box based on the first frame. The regressors are not updated during tracking in order to avoid the risk of contaminating the regressors with noisy data. For each frame, the regressors take the representation of the picked candidate box as input and produce a refined box.

4. Experiments

4.1. Implementation Details

Candidate Sampling We use the radius sampling strategy [14] to generate candidate boxes. We use 10 radial

and 10 angular divisions. The search radius is set to be the longer axis of the initial box in the first frame. At each sample location, we generate three scaled versions of the initial box with the scales being $\{\frac{\sqrt{2}}{2}, 1, \sqrt{2}\}$.

Network training We use the ALOV dataset [41] for training and validation. We choose ALOV for training as it covers many types of variations one could expect in tracking. *We exclude the 12 videos in ALOV that are also in tracking benchmark (OTB) [52], as we evaluate the proposed tracker on OTB*. After removing the 12 videos, the training set and the tracking evaluation set have no common objects. From every two frames in a video, we generate multiple pairs. One element in a pair is the groundtruth bounding box in one frame and the other element is a box sampled in the other frame. The pair is considered to be positive if the sampled box has a intersection-over-union overlap larger than 0.7 with the corresponding groundtruth box and considered to be negative if the overlap is smaller than 0.5. The training pairs and validation pairs are generated from different videos, and therefore from different objects. For training, in total we have sampled from ALOV dataset 60,000 pairs of frames and each pair of frames has 128 pairs of boxes. For validation, we have gathered 2,000 pairs of frames and the same as for training each pair of frames contains 128 pairs of boxes.

Instead of training the two-stream Siamese network from scratch, we load the pre-trained network parameters and fine tune the Siamese network. Specifically, we use the networks pre-trained for ImageNet classification, available in the Caffe library [21]. The initial fine tuning learning rate is 0.001 and the weight decay parameter is 0.001. The learning rate is decreased by a factor of 10 after every 2 epochs. We stop tuning when the validation loss does not decrease any more.

4.2. Dataset and evaluation metrics

Dataset To evaluate the tracking performance, we use the online tracking benchmark (OTB) [52]. OTB is a collection of 50 videos. 51 tracking sequences are defined with bounding box annotations. The dataset covers various challenging aspects in object tracking, such as fast motion, deformation, background clutter and occlusion.

Evaluation metrics We follow the evaluation protocol of [52], where two metrics are used: success plot and precision plot. Both metrics measure the percentage of successfully tracked frames. For the success plot, a frame is declared to be successfully tracked if the estimated bounding box and the groundtruth box have an intersection-over-union overlap larger than a certain threshold. For precision plot, tracking on a frame is considered successful if the distance between the centers of the predicted box and the groundtruth box is under some threshold. A plot is given by varying the threshold values. Tracking algorithms are

	AUC	Prec@20
(a) pretrained-alexnet-fc6	42.8	66.3
(b) firstframe-Siamese-finetuned-alexnet-fc6	44.0	67.9
(c) Siamese-finetuned-alexnet-fc6	47.4	72.0
(d) pretrained-alexnet-fc6-nomaxpooling	50.0	70.8
(e) Siamese-finetuned-alexnet-fc6-nomaxpooling	53.9	74.8
(f) Siamese-finetuned-alexnet-conv45fc6-nomaxpooling	55.0	76.2
(g) Siamese-finetuned-vgg16-conv45fc6-nomaxpooling	59.2	83.6

Table 1: Evaluation of different architectural and design choices of the Siamese invariance network for tracking on the OTB dataset [52]. We use the recommended evaluation methods, namely the area under the curve (AUC) for the success plot and the precision at 20 (Prec@20) for the precision plot.

ranked based on the area under curve (AUC) score for the success plot and precision at threshold 20 (Prec@20) for the precision plot. We use the available toolkit provided by the benchmark to generate plots and numbers. In the following, we also use success rate where needed, *i.e.*, the percentage of successfully tracked frames.

4.3. Design evaluation

We first validate our design choices of the network. In this sets of experiments, box refinement is not considered.

Network tuned generically on external video data vs. network pre-tuned on ImageNet vs. network fine tuned target-specifically on first frame In this experiment, we show the effectiveness of the Siamese network tuned on external data. To that end, we compare the Siamese fine tuned AlexNet-style network using ALOV (denoted as “Siamese-finetuned-alexnet-fc6”) with the ImageNet pre-tuned AlexNet (“pretrained-alexnet-fc6”) and the Siamese fine tuned network using the training pairs gathered in the first frame (“firstframe-Siamese-finetuned-alexnet-fc6”). In this comparison, all three use a single layer *fc6* for feature representation. As shown in the rows (a)-(c) of Table 1, the Siamese fine tuned network using ALOV (c) significantly improves over the pre-tuned net (a), while fine tuning on the first frame (b) gives a marginal improvement. We conclude that Siamese networks fine tuned using large amount of external data are to be preferred.

To max pool or not to max pool? We now examine our design choice of having no maxing pooling layers in the network (“pretrained-alexnet-fc6-nomaxpooling” vs. “pretrained-alexnet-fc6” and “Siamese-finetuned-alexnet-fc6-nomaxpooling” vs. “Siamese-finetuned-alexnet-fc6”). As shown in Table 1, (d) vs. (a) and (e) vs. (c), including max pooling layers deteriorates accuracy, as expected due to the reduction of the resolution of the feature maps which causes poor localization. When inspecting the results

	sr@0.3	sr@0.5	sr@0.7
pretrained-alexnet-fc6	68.3	46.2	19.6
pretrained-alexnet-fc6-nomaxpooling	75.3	58.1	32.6
Siamese-finetuned-alexnet-fc6	74.6	56.2	25.4
Siamese-finetuned-alexnet-fc6-nomaxpooling	79.3	67.6	38.8

Table 2: Success rates (sr) of the tracker at three intersection-over-union overlap ratios for different network architectures. From the table it is clear that a network architecture without max pooling delivers a more precise localization and hence a better matching function.

when no max pooling layers are included, the success rate improvement is higher at higher intersection-over-union overlap ratios, see Table 2. We conclude that max pooling layers are not necessary for our Siamese invariance network with small AlexNet-style architecture.

Multi-layer features vs. single-layer features Next, we evaluate whether it is more advantageous to use features from a single layer or from multiple layers. We compare “Siamese-finetuned-alexnet-conv45fc6-nomaxpooling”, which uses the outputs of layers *conv4*, *conv5* and *fc6* as features, with “Siamese-finetuned-alexnet-fc6-nomaxpooling”, which uses the output of *fc6* as feature. Table 1 shows that using multi-layer features is helpful ((e) vs. (f)). We conclude that using features from multiple layers is advantageous.

Large net vs. small net Lastly, we compare a VGGNet-style architecture with an AlexNet-style architecture (“Siamese-finetuned-vgg16-conv45fc6-nomaxpooling” vs. “Siamese-finetuned-alexnet-fc6-nomaxpooling”). Both use as features the outputs of three layers. As shown in the last two rows (f) and (g) of Table 1, using a deeper network improves the performance significantly.

4.4. State-of-the-art comparison

Overall comparison In addition to the 29 trackers included in the benchmark [52], *e.g.*, TLD [23], Struck [14] and SCM [57], we also include the most recent trackers for comparison. The included recent trackers are TGPR [9], MEEM [56], SO-DLT [47], KCFDP [19] and MUSTer [18].

As described earlier, the proposed SINT focuses on the tracking matching function, while having a simple online inference. As a preliminary demonstration that SINT can be further improved by employing more advanced online components, we also evaluate a variant of SINT, coined SINT+, which uses an adaptive candidate sampling strategy suggested by [48] and optical flow [4]. In SINT+, the sampling range is adaptive to the image resolution, set to be $30/512 * w$ in this experiment, where w is the image width. Optical flow is used in SINT+ to filter out motion inconsistent candidates. Specifically, given the pixels covered by

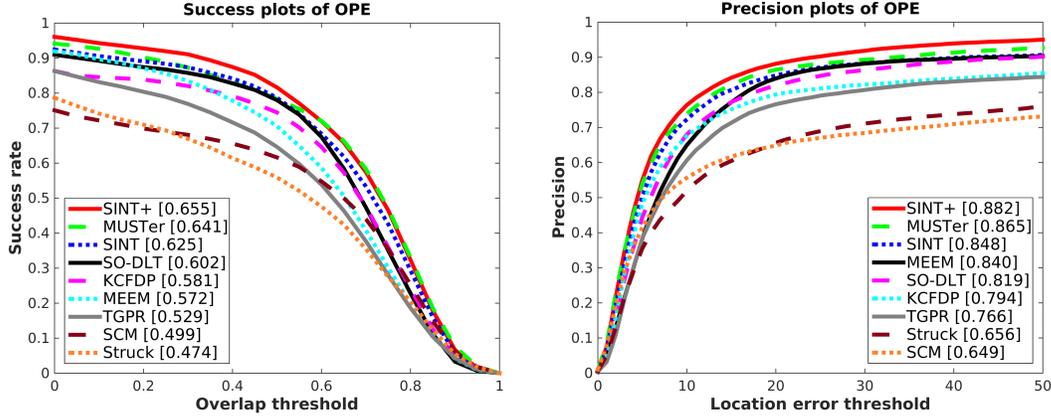


Figure 3: State-of-the-art comparison on OTB [52]. In spite of the fact that the online part of the proposed SINT is just selecting the patch that matches best to the target in the first frame, SINT is on par with state-of-the-art tracker. SINT+, using a better candidate sampling than SINT and optical flow as an additional component, achieves the best performance.

	OPE	TRE	SRE
MEEM	57.2 / 84.0	58.5 / 83.2	51.8 / 76.9
MUSTer	62.1 / 83.6	60.9 / 81.1	56.2 / 78.9
SINT	62.5 / 84.8	64.3 / 84.9	57.9 / 80.6

Table 3: Robustness evaluation on OTB, measured in AUC/Prec@20. OPE is one-pass evaluation. TRE and SRE are temporal and spatial robustness evaluation. The results of MEEM are taken from [56] and the results of MUSTer are obtained using the publicly available code.

the predicted box in the previous frame and the estimated optical flow, we know where those pixels are in the current frame and we remove the candidate boxes that contain less than 25% of those pixels, as these candidates are deemed inconsistent to the motion.

Figure 3 shows the overall performance. For clarity, only the top performing trackers are shown. Despite relying on a simple NCC-like tracking inference, SINT reaches state-of-the-art performance, being tantalizingly close to MUSTer [18] and more accurate than others by a considerable margin. SINT+, with an adaptive sampling and a simple use of optical flow, further improves SINT, outperforming clearly all state-of-the-art other trackers.

Temporal and spatial robustness To verify the robustness of the proposed tracker, we conduct the temporal robustness evaluation (TRE) and spatial robustness evaluation (SRE) defined by the benchmark. The results are summarized in Table 3. Compared to MEEM and MUSTer, SINT is temporally and spatially the same as robust, if not better.

Per distortion type comparison Further, the 50 videos in the benchmark are annotated with 11 distortion types (*e.g.*,

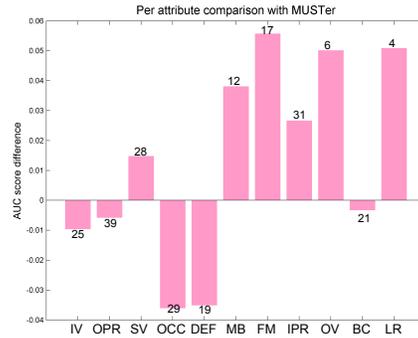


Figure 4: Per attribute comparison on AUC score of the proposed SINT with MUSTer [18]. The bars are the performance difference between SINT and MUSTer. Positive means SINT is better. The integer number at each bar is the number of tracking sequences belonging to that group.

illumination variation, occlusion *etc.*). To gain more insights, we evaluate the performance of SINT for individual attributes and compare with MUSTer [18]. SINT performs better in 6 and 7 out of the 11 groups for the AUC and the Prec@20 metrics respectively. Due to the space limit, Figure 4 only shows the plot for AUC. It is observed that MUSTer is better mainly in “occlusion” and “deformation”, whereas SINT is better in “motion blur”, “fast motion”, “in-plane rotation”, “out of view” and “low resolution”.

Failure modes of SINT When similar objects appear in view, the tracker may jump from the target to another as it only looks for the maximum similarity with the original patch of the target in the first frame (Figure 5: left). And, when there is large occlusion, the matching function might suffer (Figure 5: right).



Figure 5: Failure cases of SINT: similar confusing object (left) and large occlusion (right). Examples are from OTB sequences ‘Bolt’ and ‘Lemming’ respectively. In the left example, the tracker fires on another Jamaican runner in the same uniform as the target. In the right example, the target is heavily occluded by the lighter.

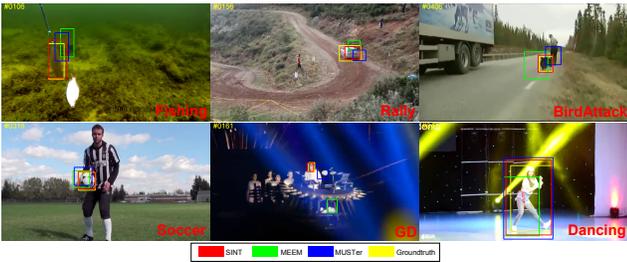


Figure 6: Example frames from the 6 test sequences.

4.5. Additional sequences and re-identification

We now further illustrate the strength of the proposed SINT on 6 newly collected sequences from YouTube. We downloaded the sequences so that they are extra challenging in terms of tracking distortions as defined by [52]. Figure 6 shows example frames from these sequences. The sequences have considerable degrees of scale change (“Fishing”, “Rally”, “BirdAttack” and “GD”), fast motion (“BirdAttack”, “Soccer” and “Dancing”), out-of-plane rotation (“Rally” and “Dancing”), non-rigid deformation (“Fishing”, “BirdAttack” and “Dancing”), low contrast (“Fishing”), illumination variation (“GD” and “Dancing”) and poorly textured objects (“Fishing” and “BirdAttack”).

We evaluate the proposed tracker, SINT, with MEEM [56] and MUSTer [18] on these sequences. The performance is summarized in Table 4, where we adopt the AUC score metric from the benchmark [52]. Results show that SINT is again a competitive tracker, outperforming MUSTer [18] and MEEM [56].

We, furthermore, observe that provided a window sampling over the whole image using [58], SINT is accurate in target re-identification, after the target was missing for a significant amount of time from the video. We illustrate this in Figure 7, where we track *Yoda*. As shown in Figure 7, the tracker has good capability of discovering the target when it re-enters the scene after being absent for a complete shot.

	MEEM [56]	MUSTer [18]	SINT
<i>Fishing</i>	4.3	11.2	53.7
<i>Rally</i>	20.4	27.5	53.4
<i>BirdAttack</i>	40.7	50.2	66.7
<i>Soccer</i>	36.9	48.0	72.5
<i>GD</i>	13.8	34.9	35.8
<i>Dancing</i>	60.3	54.7	66.8
mean	29.4	37.8	58.1

Table 4: Comparison on AUC score of the proposed SINT with MEEM [56] and MUSTer [18].



Figure 7: The capability of the tracker to re-discover the target, illustrated on a 1500-frame, 12-shot *Star Wars* video. One object (*Yoda*) is appearing in 6 of the shots, while being absent in the intermediate ones. Red dots indicate *Yoda* is present while black dots indicate *Yoda* is absent. Y-axis is the matching score with the target at the first frame. The results show good capability of the tracker to discover the target when it re-enters the scene.

5. Conclusion

This work presents Siamese Instance search Tracker, SINT. It tracks the target, simply by matching the initial target in the first frame with candidates in a new frame and returns the most similar one by a learned matching function. The strength of the tracker comes from the powerful matching function, which is the focus of the work. The matching function is learned on ALOV [41], based on the proposed two-stream very deep neural network. We take extra care that there is absolutely no overlap between the training videos and any of the videos for evaluation. Namely, we *do not* aim to do any pre-learning of the tracking targets. Once learned, the matching function is used as is, without any adapting, to track arbitrary, previously unseen targets. It turns out the matching function is very effective in coping with common appearance variations an object can have in videos. The simple tracker built upon the matching function, reaches state-of-the-art performance on OTB [52], without updating the target, tracker combination, occlusion detection and alike. Further, SINT allows for target re-identification after the target was absent for a complete shot, demonstrated on a 1500-frame, 12-shot *Star Wars* video.

References

- [1] A. Adam, E. Rivlin, and I. Shimshoni. Robust fragments-based tracking using the integral histogram. In *CVPR*, 2006. 2
- [2] K. Briechle and U. D. Hanebeck. Template matching using fast normalized cross correlation. In *Proceedings of SPIE: Optical Pattern Recognition XII*, 2001. 2
- [3] J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. LeCun, C. Moore, E. Säckinger, and R. Shah. Signature verification using a siamese time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04):669–688, 1993. 2, 3
- [4] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *TPAMI*, 33(3):500–513, 2011. 6
- [5] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005. 3, 4
- [6] D. Comaniciu, V. Ramesh, and P. Meer. Real-time tracking of non-rigid objects using mean shift. In *CVPR*, 2000. 1, 2
- [7] R. O. Duda, P. E. Hart, et al. *Pattern classification and scene analysis*, volume 3. Wiley New York, 1973. 2
- [8] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *TPAMI*, 32(9):1627–1645, 2010. 5
- [9] J. Gao, H. Ling, W. Hu, and J. Xing. Transfer learning based visual tracking with gaussian processes regression. In *ECCV*, 2014. 6
- [10] R. Girshick. Fast r-cnn. *arXiv preprint arXiv:1504.08083*, 2015. 3, 4
- [11] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 5
- [12] M. Godec, P. M. Roth, and H. Bischof. Hough-based tracking of non-rigid objects. In *ICCV*, 2011. 2
- [13] X. Han, T. Leung, Y. Jia, R. Sukthankar, and A. C. Berg. Matchnet: Unifying feature and metric learning for patch-based matching. In *CVPR*, 2015. 3
- [14] S. Hare, A. Saffari, and P. H. Torr. Struck: Structured output tracking with kernels. In *ICCV*, 2011. 2, 5, 6
- [15] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *CVPR*, 2015. 4
- [16] S. Holzer, S. Ilic, D. J. Tan, M. Pollefeys, and N. Navab. Efficient learning of linear predictors for template tracking. *IJCV*, 2015. 2
- [17] S. Hong, T. You, S. Kwak, and B. Han. Online tracking by learning discriminative saliency map with convolutional neural network. In *ICML*, 2015. 2
- [18] Z. Hong, Z. Chen, C. Wang, X. Mei, D. Prokhorov, and D. Tao. Multi-store tracker (muster): A cognitive psychology inspired approach to object tracking. In *CVPR*, 2015. 2, 6, 7, 8
- [19] D. Huang, L. Luo, M. Wen, Z. Chen, and C. Zhang. Enable scale and aspect ratio adaptability in visual tracking with detection proposals. In *BMVC*, 2015. 6
- [20] H. Jégou, M. Douze, and C. Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *ECCV*, 2008. 3
- [21] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 5
- [22] Y. Jiang, J. Meng, and J. Yuan. Randomized visual phrases for object search. In *CVPR*, 2012. 3
- [23] Z. Kalal, K. Mikolajczyk, and J. Matas. Tracking-learning-detection. *TPAMI*, 34(7):1409–1422, 2012. 2, 6
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 3, 4
- [25] H. Li, Y. Li, F. Porikli, et al. Deeptrack: Learning discriminative feature representations by convolutional neural networks for visual tracking. In *BMVC*, 2014. 2
- [26] Z. Li, E. Gavves, K. E. van de Sande, C. G. Snoek, A. W. Smeulders, et al. Codemaps-segment, classify and search objects locally. In *ICCV*, 2013. 3
- [27] T.-Y. Lin, Y. Cui, S. Belongie, J. Hays, and C. Tech. Learning deep representations for ground-to-aerial geolocalization. In *CVPR*, 2015. 3
- [28] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 4
- [29] B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *IJCAI*, 1981. 1, 2
- [30] X. Mei and H. Ling. Robust visual tracking using l1 minimization. In *ICCV*, 2009. 2
- [31] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. 3
- [32] H. T. Nguyen and A. W. Smeulders. Robust tracking using foreground-background texture discrimination. *Int. J. Comput. Vision*, 2006. 1, 2
- [33] J. Pan and B. Hu. Robust occlusion handling in object tracking. In *CVPR*, 2007. 1
- [34] F. Pernici and A. D. Bimbo. Object tracking by oversampling local features. *TPAMI*, 2014. 2
- [35] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *ECCV*, 2010. 3
- [36] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *CVPR*, 2007. 2, 3
- [37] D. Qin, C. Wengert, and L. van Gool. Query adaptive similarity for large scale object retrieval. In *CVPR*, 2013. 3
- [38] D. A. Ross, J. Lim, R.-S. Lin, and M.-H. Yang. Incremental learning for robust visual tracking. *IJCV*, 2008. 1, 2
- [39] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 3, 4
- [40] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *ICCV*, 2003. 3
- [41] A. W. Smeulders, D. M. Chu, R. Cucchiara, S. Calderara, A. Dehghan, and M. Shah. Visual tracking: an experimental survey. *TPAMI*, 36(7):1442–1468, 2014. 1, 5, 8

- [42] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *CVPR*, 2014. 3
- [43] R. Tao, E. Gavves, C. G. Snoek, and A. W. M. Smeulders. Locality in generic instance search from one example. In *CVPR*, 2014. 2, 3
- [44] R. Tao, A. W. M. Smeulders, and S.-F. Chang. Attributes and categories for generic instance search from one example. In *CVPR*, 2015. 2, 3
- [45] G. Toliás, Y. Avrithis, and H. Jégou. To aggregate or not to aggregate: Selective match kernels for image search. In *ICCV*, 2013. 3
- [46] G. Toliás, Y. Avrithis, and H. Jégou. Image search with selective match kernels: aggregation across single and multiple images. *International Journal of Computer Vision*, 2015. 2
- [47] N. Wang, S. Li, A. Gupta, and D.-Y. Yeung. Transferring rich feature hierarchies for robust visual tracking. *arXiv preprint arXiv:1501.04587*, 2015. 2, 6
- [48] N. Wang, J. Shi, D.-Y. Yeung, and J. Jia. Understanding and diagnosing visual tracking systems. In *ICCV*, 2015. 6
- [49] N. Wang and D.-Y. Yeung. Learning a deep compact image representation for visual tracking. In *NIPS*, 2013. 2
- [50] S. Wang, H. Lu, F. Yang, and M.-H. Yang. Superpixel tracking. In *ICCV*, 2011. 2
- [51] T. Wu, Y. Lu, and S. Zhu. Online object tracking, learning and parsing with and-or graphs. *TPAMI*, 2015. 2
- [52] Y. Wu, J. Lim, and M.-H. Yang. Online object tracking: A benchmark. In *CVPR*, 2013. 1, 5, 6, 7, 8
- [53] S. Zagoruyko and N. Komodakis. Learning to compare image patches via convolutional neural networks. In *CVPR*, 2015. 3
- [54] J. Žbontar and Y. LeCun. Computing the stereo matching cost with a convolutional neural network. In *CVPR*, 2015. 3
- [55] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014. 4
- [56] J. Zhang, S. Ma, and S. Sclaroff. Meem: Robust tracking via multiple experts using entropy minimization. In *ECCV*, 2014. 2, 6, 7, 8
- [57] W. Zhong, H. Lu, and M.-H. Yang. Robust object tracking via sparsity-based collaborative model. In *CVPR*, 2012. 6
- [58] C. L. Zitnick and P. Dollár. Edge boxes: Locating object proposals from edges. In *ECCV*, 2014. 8