

## Learning the Model Update for Siamese Trackers

Lichao Zhang<sup>1</sup>, Abel Gonzalez-Garcia<sup>1</sup>, Joost van de Weijer<sup>1</sup>, Martin Danelljan<sup>2</sup>, Fahad Shahbaz Khan<sup>3,4</sup>

<sup>1</sup> Computer Vision Center, Universitat Autònoma de Barcelona, Spain

<sup>2</sup> Computer Vision Laboratory, ETH Zürich, Switzerland

<sup>3</sup> Inception Institute of Artificial Intelligence, UAE

<sup>4</sup> Computer Vision Laboratory, Linköping University, Sweden

{lichao, agonzalez, joost}@cvc.uab.es, martin.danelljan@vision.ee.ethz.ch, fahad.khan@liu.se

### Abstract

*Siamese approaches address the visual tracking problem by extracting an appearance template from the current frame, which is used to localize the target in the next frame. In general, this template is linearly combined with the accumulated template from the previous frame, resulting in an exponential decay of information over time. While such an approach to updating has led to improved results, its simplicity limits the potential gain likely to be obtained by learning to update. Therefore, we propose to replace the handcrafted update function with a method which learns to update. We use a convolutional neural network, called UpdateNet, which given the initial template, the accumulated template and the template of the current frame aims to estimate the optimal template for the next frame. The UpdateNet is compact and can easily be integrated into existing Siamese trackers. We demonstrate the generality of the proposed approach by applying it to two Siamese trackers, SiamFC and DaSiamRPN. Extensive experiments on VOT2016, VOT2018, LaSOT, and TrackingNet datasets demonstrate that our UpdateNet effectively predicts the new target template, outperforming the standard linear update. On the large-scale TrackingNet dataset, our UpdateNet improves the results of DaSiamRPN with an absolute gain of 3.9% in terms of success score. Code and models are available at <https://github.com/zhanglichao/updatenet>.*

### 1. Introduction

Generic visual object tracking is the task of predicting the location of a target object in every frame of a video, given its initial location. Tracking is one of the fundamental problems in computer vision, spanning a wide range of applications including video understanding [34], surveillance [12], and robotics [29]. It is a highly challenging task



Figure 1. **Qualitative comparison between model updates.** We learn to update the model template using *UpdateNet*. When combined with Siamese trackers such as SiamFC [1], our learned updating strategy can be effectively adapted to current circumstances, unlike the simple linear update commonly used.

due to frequent appearance changes, various types of occlusions, the presence of distractor objects, and environmental aspects such as motion blur or illumination changes.

Currently, there are two prevalent tracking paradigms: Siamese tracking methods [1, 27, 41, 50] and tracking-by-detection methods [3, 7, 11, 19, 32, 47, 48]. In this work, we consider Siamese trackers, since they provide competitive accuracy while achieving impressive computational efficiency. The basic principle of these trackers is to match an object appearance template with a corresponding feature representation of the search region in the test frame. The features for the object template and the search region are acquired through a deep neural network trained offline on a large dataset. Such a training strategy has shown to provide excellent visual descriptors for the tracking task [1, 50].

In the original Siamese tracker [1], the object template is initialized in the first frame and then kept fixed during the remainder of the video. However, appearance changes are often large and failing to update the template could lead to early failure of the tracker. In such scenarios, it is

important to adapt the model to the current target appearance. To accommodate this problem, more recent Siamese trackers [27, 41, 50] have implemented a simple linear update strategy using a running average with a fixed learning rate [36]. This strategy assumes a constant rate of appearance change across all frames in the video, as well as across different videos. In practice, the update requirements for the object template greatly vary for different tracking situations, which depend on a complex combination of external factors such as motion, blur, or background clutter. Therefore, a simple linear update is often inadequate to cope with changing updating needs and to generalize to all potentially encountered situations. Moreover, this update is also constant in all spatial dimensions, which does not allow for localized partial updates. This is especially damaging in situations such as partial occlusions, where only a certain part of the template needs to be updated. Finally, excessive reliance on the initial template may suffer from catastrophic drift and the inability to recover from tracking failures.

In this paper, we propose to *learn* the target template update itself. Our learned update strategy utilizes target and image information, and is thus adaptive to the present circumstances of each particular situation. In our approach, the updated template is computed as a function of (i) the initial ground-truth template, (ii) the accumulated template from all previous frames, and (iii) the feature template at the predicted object location in the current frame. Hence, the new accumulated template contains an effective historical summary of the object’s current appearance, as it is continually updated using the most recent information while being robust due to the strong signal given by the initial object appearance. More specifically, the aforementioned template update function is implemented as a convolutional neural network, *UpdateNet*. This is a compact model that can be combined with any Siamese tracker to enhance its online updating capabilities while maintaining its efficiency properties. Furthermore, it is sufficiently complex to learn the nuances of effective template updating and be adaptive enough to handle a large collection of tracking situations.

We evaluate UpdateNet by combining it with two state-of-the-art Siamese trackers: SiamFC [1] and DaSiamRPN [50]. Through extensive experiments on common tracking benchmarks, such as VOT2018 [22], we demonstrate how our UpdateNet provides enhanced updating capabilities that in turn result in improved tracking performance (see Figure 1). We also present results in the recent LaSOT dataset [13], which is substantially more challenging as it contains abundant long-term sequences. Overall, we propose an efficient model to learn how to effectively update the object template during online tracking and that can be applied to different existing Siamese trackers.

## 2. Related work

**Tracking Frameworks.** Most existing tracking methods are either based on tracking-by-detection or employ template matching. Object trackers based on tracking-by-detection pose the task of target localization as a classification problem where the decision boundary is obtained by online learning a discriminative classifier using image patches from the target object and the background. Among the tracking-by-detection approaches, discriminative correlation filter based trackers [19, 48, 11, 7] have recently shown excellent performance on several tracking benchmarks [42, 43, 24, 22]. These trackers learn a correlation filter from example patches of the target appearance to discriminate between the target and background appearance.

The other main tracking framework is based on template matching, typically using Siamese networks [1, 40, 41, 18, 27, 50], that implement a similarity network by spatial cross-correlation. Bertinetto *et al.* [1] proposed a Siamese tracker which is based on a two-stream architecture. One stream extracts the object template’s features based on an *exemplar* image that contains the object to be tracked. The other stream receives as input a large *search region* in the target image. The two outputs are cross-correlated to generate a response map of the search region. Many trackers have extended the SiamFC architecture [40, 16, 41, 27, 50, 49] for tracking. The Siamese-based trackers have gained popularity since they provide a good trade-off between computational speed and tracking performance. However, most of these approaches struggle to robustly classify the target especially in the presence of distractors due to no online learning. In this work, we analyze the limitations of Siamese trackers regarding the update of the template model and propose a solution to address them.

**Updating the object template** Most trackers either use simple linear interpolation to update the template in every frame [3, 19, 9, 11, 20, 5] or do not update the initial template at all [1, 41, 27, 50]. Such update mechanisms are insufficient in most tracking situations, as the target object may suffer appearance changes given by deformations, fast motion, or occlusion. Moreover, fixed update schedules also result in object templates that are more focused on recent frames [10], while forgetting the historical appearances of the object. To address this issue, Danelljan *et al.* [10, 11] propose to include a subset of historic frames as training samples when calculating the current correlation filter, which leads to better results than the traditional linear frame-by-frame update. Nonetheless, storing multiple samples in memory results in increased computation and memory usage, which in turn heavily decreases the tracking speed. The ECO tracker [7] tries to alleviate this problem by modelling the distribution of training samples as a mixture of Gaussians, where each component represents a distinct appearance. This significantly reduces required storage and,

combined with a conservative update strategy (only every five frames), leads to increased tracking efficiency. Even with more previous samples, the correlation filter is still updated by averaging the filters of their corresponding samples (as still linear interpolation update).

Recently, Yang *et al.* [45] employed a Long Short-Term Memory (LSTM) to estimate the current template by storing previous templates in memory during on-line tracking, which is computationally expensive and a rather complex system. Choi *et al.* [6] also uses a template memory but uses reinforcement learning to select one of the stored templates. This method fails to accumulate information from multiple frames. The meta-tracker of [33] extends the initialization of the target model in the first frame by a pre-trained approach, but still need a linear update in on-line tracking. Yao *et al.* [46] propose to learn the updating coefficients for CF trackers using SGD offline. While the solution for correlation filter is still the hand-crafted manner, and these coefficients are fixed without updating during tracking.

To adapt to the object variations, Guo *et al.* [15] propose to compute a transformation matrix with respect to the initial template through regularized linear regression in the Fourier domain. Since only the initial template is considered when estimating the transformation, this approach ignores the historical object variations observed during tracking, which may result important for a smoother adaptation of the exemplar template. Moreover, they compute the transformation matrix as a closed-form solution on the Fourier domain, which suffers from issues related to the boundary effect [21]. Our work instead uses a powerful yet easily trainable model to update the object template based not only the first frame, but also on the accumulated template using all previous frames, leveraging the observed training data. Furthermore, our UpdateNet is trained to *learn* how to perform an effectively update the object template, based on the observed training tracking data.

### 3. Updating the object template

In this section, we present our approach to learn how to update the object template during online tracking. We start by revisiting the standard update mechanism in tracking and identifying its drawbacks. Then, we introduce our formulation to overcome them and describe our model and training procedure in detail. The focus of this paper is on Siamese trackers. Note, however, that our approach is not limited to Siamese trackers and the same formulation could be applied to other types of trackers, e.g. DCF [19, 11, 7].

#### 3.1. Standard update

Several recent tracking approaches [4, 3, 19, 40, 41, 27, 50] use a simple averaging strategy to update the object appearance model given a new data sample. This strategy dates from early tracking methods [36] and has long been the standard for online updating given its acceptable results

and in spite of its limitations. The template is updated as a running average with exponentially decaying weights over time. The choice of exponential weights yields the following recursive formula for updating the template,

$$\tilde{T}_i = (1 - \gamma)\tilde{T}_{i-1} + \gamma T_i. \quad (1)$$

Here,  $i$  is the frame index,  $T_i$  is the new template sample computed using only the current frame, and  $\tilde{T}_i$  is the accumulated template. The update rate  $\gamma$  is normally set to a fixed small value (e.g.  $\gamma = 0.01$ ) following the assumption that the object's appearance changes smoothly and consistently in consecutive frames. In DCF trackers (e.g. [3, 19]),  $T$  corresponds to the correlation filter. In Siamese trackers, instead,  $T$  is the object appearance template extracted from a particular frame by a fully convolutional feature extractor. While the original SiamFC tracker [1] does not perform any model update, more recent Siamese trackers [1, 41, 27, 50] have adopted (1) to update their templates.

While template averaging provides a simple means of integrating new information, it has several severe drawbacks:

- It applies a constant update rate for every video, despite the possibly different updating necessities caused by multiple factors such as camera motion. Even within the same video, the required update on the object template may dynamically vary at different times.
- The update is also constant along all spatial dimensions of the template, including the channel dimension. This prevents updating only part of the template, which is desirable under partial occlusions, for example.
- The tracker cannot recover from drift. Partially, this is caused by the fact that it loses access to the appearance template  $T_0$ , which is the only template which is without doubt on the object.
- The update function is constrained to a very simple linear combination of previous appearance templates. This severely limits the flexibility of the update mechanism, important when the target undergoes complex appearance changed. Considering more complex combination functions is expected to improve results.

#### 3.2. Learning to update

We address the drawbacks listed above by proposing a model that learns an adaptive update strategy. Since the focus of this paper is on Siamese trackers,  $T$  represents here the object appearance template. To address the limitations of simple template averaging, we propose to learn a generic function  $\phi$  that updates the template according to,

$$\tilde{T}_i = \phi(T_0^{GT}, \tilde{T}_{i-1}, T_i). \quad (2)$$

The learned function  $\phi$  computes the updated template based on initial ground-truth template  $T_0^{GT}$ , the last accumulated template  $\tilde{T}_{i-1}$  and the template  $T_i$  extracted

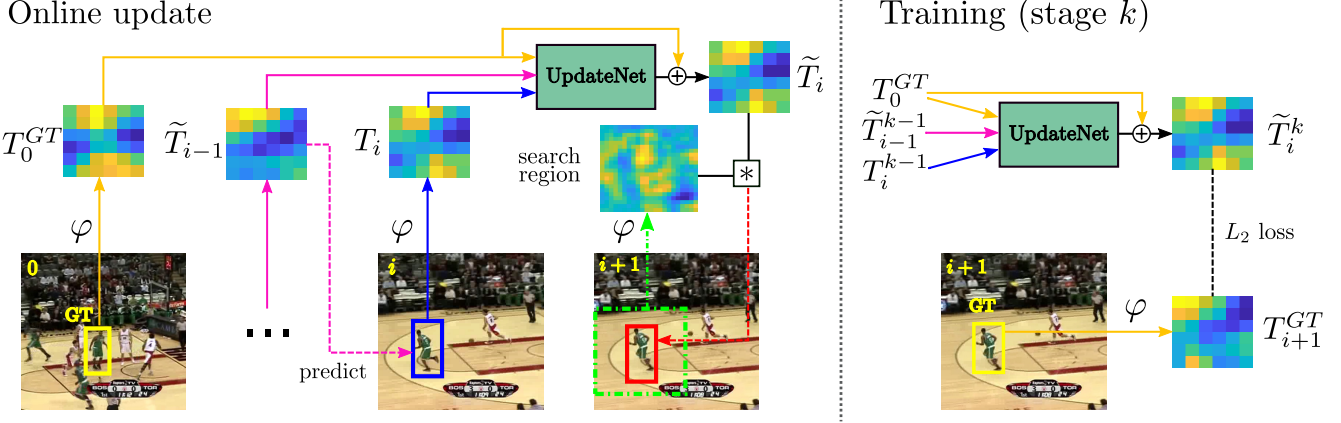


Figure 2. **Overview of our tracking framework with UpdateNet.** (Left) The online update of the object template is performed by UpdateNet, which receives as input the initial ground-truth template, last accumulated template and current predicted template, and outputs updated accumulated template. (Right) Training of UpdateNet using the distance to the ground-truth object template on next frame.

from the predicted target location in the current frame. In essence, the function updates the previous accumulated template  $\tilde{T}_{i-1}$  by integrating the new information given by the current frame  $T_i$ . Therefore,  $\phi$  can be adapted to the specific updating requirements of the current frame, based on the difference between the current and accumulated templates. Moreover, it also considers the initial template  $T_0^{GT}$  in every frame, which provides highly reliable information and increases robustness against model drift. The function  $\phi$  is implemented as a convolutional neural network, which grants great expressive power and the ability to learn from large amounts of data. We call this neural network *UpdateNet* and describe it in detail in the following section.

### 3.3. Tracking framework with UpdateNet

We present here the structure of UpdateNet and describe how it is applied for online tracking. Figure 2 (left) presents an overview of our adaptive object update strategy using UpdateNet with a Siamese tracker. We extract deep features from image regions with a fixed fully convolutional network  $\phi$ , employing the same feature extractor as in the SiamFC tracker [1]. We extract  $T_0^{GT}$  from the ground-truth object location in the initial frame (number 0 in Figure 2). In order to obtain  $T_i$  for the current frame, we use the accumulated template from all previous frames  $\tilde{T}_{i-1}$  to predict the object location in frame  $i$  (dashed purple line) and extract the features from this region (solid blue line). Note that  $\tilde{T}_{i-1}$  corresponds to the output of UpdateNet for the previous time step, not shown here for conciseness. We concatenate the extracted features  $T_0^{GT}$  and  $T_i$  with the accumulated features  $\tilde{T}_{i-1}$  to form the input of UpdateNet. This input is then processed through a series of convolutional layers (sec. 4.3) and outputs the predicted new accumulated template  $\tilde{T}_i$ . For the first frame, we set  $T_i$  and  $\tilde{T}_{i-1}$  to  $T_0^{GT}$  as there have not been any previous frames.

The only ground-truth information used by UpdateNet is

the given object location in the initial frame, all other inputs are based on predicted locations. Hence,  $T_0^{GT}$  is the most reliable signal on which UpdateNet can depend to guide the update. For this reason, we employ a residual learning strategy [17], where UpdateNet learns how to modify the ground-truth template  $T_0^{GT}$  for the current frame. This is implemented by adding a skip connection from  $T_0^{GT}$  to the output of UpdateNet. This approach still takes into account the set of historical appearances of the object for updating, but pivots such update on the most accurate sample. We have also experimented with adding skip connections from other inputs as well as no residual learning at all (see sec. 4).

### 3.4. Training UpdateNet

We train our UpdateNet to predict the target template in the coming frame, i.e. the predicted template  $\tilde{T}_i$  should match the template  $T_{i+1}^{GT}$  extracted from the ground-truth location in the next frame (Figure 2, right). The intuition behind this choice is that  $T_{i+1}^{GT}$  is the optimal template to use when searching for the target in the next frame. In order to achieve this, we train UpdateNet by minimizing the Euclidean distance between the updated template and the ground-truth template of the next frame, defined as

$$\mathcal{L}_2 = \left\| \phi(T_0^{GT}, \tilde{T}_{i-1}, T_i) - T_{i+1}^{GT} \right\|_2. \quad (3)$$

In the remainder of this section we describe the procedure employed to generate the training data and introduce a multi-stage training approach for UpdateNet.

**Training samples.** In order to train UpdateNet to minimize (3), we need pairs of input triplets  $(T_0^{GT}, \tilde{T}_{i-1}, T_i)$  and outputs  $T_{i+1}^{GT}$  that reflect the updating needs of the tracker during online application. The object templates of the initial frame  $T_0^{GT}$  and target frame  $T_{i+1}^{GT}$  can be easily obtained by extracting features from ground-truth locations in corresponding frames. In case of the current frame's template



$T_i$ , however, using the ground-truth location represents a seldom encountered situation in practice, for which the predicted location in the current frame is very accurate. This unrealistic assumption biases the update towards expecting very little change with respect to  $T_i$ , and thus UpdateNet cannot learn a useful updating function. Therefore, we need to extract  $T_i$  samples for training by using an imperfect localization in the  $i$ -th frame. We can simulate such situation by using the accumulated template  $\tilde{T}_{i-1}$ , ideally presenting localization errors that occur during online tracking.

**Multi-stage training.** In theory, we could use the accumulated template  $\tilde{T}_{i-1}$  output by UpdateNet. However, this would force the training to be recurrent, making the procedure cumbersome and inefficient. To avoid this, we split our training procedure into sequential stages that iteratively refine UpdateNet. In the first stage, we run the original tracker on the training dataset using the standard linear update

$$\tilde{T}_i^0 = (1 - \gamma) \tilde{T}_{i-1}^0 + \gamma T_i^0, \quad (4)$$

which generates accumulated templates and realistically predicted locations for each frame. We set the update rate  $\gamma$  to the recommended value for the tracker. This corresponds to a first approximation to likely inputs for UpdateNet during tracking inference, albeit with the less sophisticated linear update strategy. In every posterior training stage  $k \in \{1, \dots, K\}$ , we use the UpdateNet model trained in the previous stage to get accumulated templates and object location predictions as follows

$$\tilde{T}_i^k = \phi_i^k \left( T_0^{GT}, \tilde{T}_{i-1}^{k-1}, T_i^{k-1} \right). \quad (5)$$

Such training data samples closely resemble the expected data distribution at inference time, as they have been output by UpdateNet. We investigate a suitable value for the total number of stages  $K$  in the experimental section (sec. 4).

## 4. Experiments

### 4.1. Training dataset

We use the recent Large-scale Single Object Tracking (LaSOT) [13] to train our UpdateNet. LaSOT has 1,400 sequences in 70 categories, which amounts to a total of 3.52M frames. Each category contains exactly twenty sequences, making the dataset balanced across classes. It also provides longer sequences that contain more than 1,000 frames (2,512 frames on average) in order to satisfy the current long-term trend in tracking. We used the official training and test splits, which preserve the balanced class distribution. In fact, we only employ a subset containing 20 training sequences from 20 randomly selected categories, with a total of 45,578 frames. We have found experimentally that this suffices to learn an effective updating strategy, and that additional data brings only a small performance boost while increasing training time.

### 4.2. Evaluation datasets and protocols

We evaluate results on standard tracking benchmarks: VOT2018/16 [23], LaSOT [13] and TrackingNet [31].

**VOT2018/16 [23].** VOT2018 dataset has 60 public testing sequences, with a total of 21,356 frames. It is used as the most recent edition of the VOT challenge. The VOT protocol establishes that when the evaluated tracker fails, i.e. when the overlap with the ground-truth is below a given threshold, it is re-initialized in the correct location five frames after the failure. The main evaluation measure used to rank the trackers is Expected Average Overlap (EAO), which is a combination of accuracy (A) and robustness (R). We also use VOT2016 [24] for comparison purposes, which has 10 different sequences with VOT2018 [22]. We compute all results using the provided toolkit [22].

**LaSOT [13].** LaSOT is a much larger and more challenging dataset including long-term sequences. Following recent works that use this dataset [26, 14], we report results on protocol II, i.e. LaSOT testing set. The testing subset has 280 sequences with 690K frames in total. LaSOT dataset [13] follows the OPE criterion of OTB [42]. It consists of precision plot which is measured by the center location error, and success plot which is measured through the intersection over union (IoU) between the predicted bounding box and the ground-truth. Besides precision plot and success plot, LaSOT also uses normalized precision plot to counter the situation that target size and image resolution have large discrepancies for different frames and videos, which heavily influences the precision metric. We use success plot and normalized precision plot to evaluate the trackers in the article. We use their code [13] to create all plots.

**TrackingNet [31].** This is a large-scale tracking dataset consisting of videos in the wild. It has a total of 30,643 videos split into 30,132 training videos and 511 testing videos, with an average of 470,9 frames. It uses precision, normalized precision and success as evaluation metrics.

### 4.3. Implementation details

We use SiamFC [1] and DaSiamRPN [50] as our base trackers, and the backbone Siamese network adopts the modified AlexNet. We do not perform any changes except for in the updating component. The original implementation of SiamFC did not perform any object update. We borrow a linear update rate from CFNet [40] as  $\gamma = 0.0102$  for templates generating in the training stage 1. We use the original version of DaSiamRPN, which does not employ any update strategy. We analyze the effect of the linear update rate in the tracking performance in sec. 4.7. To train UpdateNet, we set a group of templates, including  $T_0^{GT}$ ,  $\tilde{T}_{i-1}$ ,  $T_i$  and  $T_{i+1}^{GT}$  as the input. They are all sampled sequentially from a same video. It is noteworthy that  $\tilde{T}_{i-1}$  and  $T_i$  are generated by the real tracking pro-

Update for SiamFC	Skip	EAO ( $\uparrow$ )	A ( $\uparrow$ )	R ( $\downarrow$ )
Linear	-	0.188	0.50	0.59
UpdateNet ( $K = 1$ )	-	0.205	0.48	0.58
UpdateNet ( $K = 1$ )	$T_i$	0.207	0.47	0.57
UpdateNet ( $K = 1$ )	$\tilde{T}_{i-1}$	0.214	0.49	0.58
UpdateNet ( $K = 1$ )	$T_0^{GT}$	0.250	0.50	0.53
UpdateNet ( $K = 2$ )	$T_0^{GT}$	0.257	0.51	0.50
UpdateNet ( $K = 3$ )	$T_0^{GT}$	<b>0.262</b>	<b>0.52</b>	<b>0.49</b>

Table 1. **Ablation study on VOT2018 [22].** We present several update strategies for SiamFC [1]. The results are reported in terms of EAO, normalized weighted mean of accuracy (A), and normalized weighted mean of robustness score (R). ‘Skip’ column indicates the origin of the skip connection, if any. Here,  $K$  is the number of stages UpdateNet is trained for.

cedure, while  $T_0^{GT}$  and  $T_{i+1}^{GT}$  are ground-truth templates. We store all training object templates on disk, extracted using either linear/no update (stage  $k = 1$ ) or a previous version of UpdateNet ( $k > 1$ ). Let the template size be  $H \times W \times C$ . UpdateNet is a two-layer convolutional neural network: one  $1 \times 1 \times 3 \cdot C \times 96$  convolutional layer, followed by a ReLU and a second convolutional layer of dimensions  $1 \times 1 \times 96 \times C$ . For SiamFC,  $H = W = 6$  and  $C = 256$ , whereas DaSiamRPN  $C = 512$ . In the first stage, the weights are initialized from scratch and the learning rate is decreased logarithmically at each epoch from  $10^{-6}$  to  $10^{-7}$ . In next stage, the weights are initialized by the best model from the last stage, and the learning rate is decreased logarithmically at each epoch from  $10^{-7}$  to  $10^{-8}$ . We train the model for 50 epochs with mini-batches of size 64. We use Stochastic Gradient Descent (SGD) with momentum of 0.9 and weight decay of 0.0005.

#### 4.4. Ablation study

We start our evaluation by ablating our approach on different components in order to analyze their contribution to the final performance. Table 1 shows the results using VOT2018 [22] dataset under the EAO measure. In the middle of the table, it shows that updating the object template with the first stage of UpdateNet results beneficial when it is residually trained with respect to  $T_0^{GT}$ , as the learned update strategy is grounded on a reliable object sample. Moreover, our multi-stage training further increases the performance achieved by UpdateNet, reaching a total improvement of 7.4% with respect to the original SiamFC with no update. For the remainder of the paper, we use UpdateNet trained on 3 stages and with skip connection from  $T_0^{GT}$ .

#### 4.5. Analysis on representation update

This section attempts to provide insights regarding the performance improvements achieved by UpdateNet. Siamese networks are trained to project the images into a




Figure 3. **Visualization of accumulated templates for SiamFC.** ‘Frame’ column shows the search region and ground-truth box used to extract the templates, whose top four channels we show in ‘GT’. ‘No-update’ presents the response map resulting from applying the initial template to the search region. For ‘Linear’ and ‘UpdateNet’ strategies we also show their accumulated templates.

feature space in which spatial correlation is maximal. Update strategies operate on the learned features, possibly interfering with their correlation abilities and potentially damaging the tracking performance. In order to study the interference of the update strategies on the features we visualize the accumulated templates of SiamFC for both linear update and UpdateNet in Figure 3. We also include the ground-truth template extracted from the annotated bounding-box. For each template we show the feature maps of the four most dynamic channels in the ground-truth template, arranged in a  $2 \times 2$  grid. For comparison reasons, the accumulated templates are generated using the ground-truth object locations instead of the predicted locations during tracking. Moreover, next to each accumulated template we also show the response map generated when correlating the template with the search region. We observe several interesting properties that support the performance gains seen in practice. First, the accumulated templates using UpdateNet resemble the ground-truth more closely than those in linear update (see e.g. highlight on bottom-right channel for frame 38, first example). Second, response maps tend to be sharper on the object location for UpdateNet, which shows how our strategy does not negatively interfere with the desired correlation properties of the learned features. Finally, the accumulated templates of the linear update change at a very slow rate and are clearly deficient in keeping up with the appearance variation exhibited in videos.

In order to further study this observation, we propose quantifying the change rate between templates of contiguous frames. For each  $i \in \{1, \dots, N\}$  we compute the aver-




Figure 4. **Change rate between contiguous frames.** We present individual results for two example videos (top, middle) and average results for all videos in VOT2018.

age difference in the template as  $\delta_i = \frac{1}{|E|} \sum_E |T_i - T_{i-1}|$ , where  $N$  is the number of frames in a video and the sum runs over each element of the feature maps (e.g.  $E = 6 \times 6 \times 256$ ). We present the results in Figure 4. The bottom row contains the average change rate  $\delta$  of all 60 videos in VOT2018 [22]. It is clear that the linear update strategy cannot deliver the updating rate required by the change in the features of the ground-truth template. UpdateNet, on the other hand, provides a much more adaptive strategy that is substantially closer in magnitude to the change rate of the ground-truth template. The top and middle rows also show the change rate of the two individual sequences in Figure 3, 'book' and 'soccer1'. We can see UpdateNet mimics the real template in high change periods, as indicated by the high correlation on their extremes. This leads to predicting better response map as shown in Figure 3.

#### 4.6. Generality and tracking speed

In this section, we evaluate the generality of our UpdateNet by applying it to other Siamese trackers as shown in Figure 5. It presents results on VOT2018 in terms of EAO with respect to the tracking speed. We measure tracking speed in frames per second (FPS) and use a logarithmic scale on its axis. We observe that we improve on Siamese trackers, e.g. SiamFC [1] and DaSiamRPN [50] by adding a very small temporal overhead. Finally the top-performance trackers are shown in Figure 6. We compare with trackers including DRT [37], DeepSTRCF [28], LSART [38], R\_MCPF [48], SRCT [25], CSRDCF [30], LADCF [44], MFT [22], UPDT [2] and ATOM [8] among others as in [22]. Among the top trackers, our approach achieves superior performance while maintaining a very high efficiency. Further, our tracker obtains a performance relative gain of 2.8% over the base tracker DaSiamRPN [50].




Figure 5. **EAO vs. speed on VOT2018.** We compare our UpdateNet combined with two different Siamese trackers against the state-of-the-art methods. UpdateNet can substantially improve the tracking performance without significantly affecting the speed.




Figure 6. **EAO performance on VOT2018.** We compare our method with the state-of-the-art methods on VOT2018. Our proposed approach achieves superior performance.

#### 4.7. Fine-tuning the linear update rate

The linear update in the previous section uses the update rate for SiamFC recommended by the authors [40] ( $\gamma = 0.0102$ ) and for DaSiamRPN from the original tracker [50] ( $\gamma = 0$ ). We now investigate whether the linear update strategy can bring higher performance gains when fine-tuning the update rate on the test set. We test several update rates uniform sampled from the  $[0, 0.2]$  interval. Figure 7 shows the performance of the linear update for DaSiamRPN (light green) and SiamFC (dark green). The red dashed line at the top and pink dashed line in the middle are the performance of our UpdateNet applied on DaSiamRPN and SiamFC respectively. We can see for SiamFC the peak performance is indeed achieved between 0.01 and 0.05. For DaSiamRPN, the original tracker with no-update performs best, which proves that for more complex Siamese trackers trained off-line, on-line linear update may even damage the performance. This shows that even a fine-tuned linear update cannot improve its results any further. Moreover, our UpdateNet outperforms all update rate values without requiring any manual fine-tuning. Despite the need of a higher update rate for some videos, we can see how the per-




Figure 7. **The linear update rate evaluation for DaSiamRPN and SiamFC on VOT2018 [22].** The x-axis is the update rate values. The y-axis is the EAO scores on VOT protocol [22]. The red and pink dashed lines are our UpdateNet performances with DaSiamRPN and SiamFC, respectively.

	DSiam [15]	MemTrack [45]	SiamFC		DaSiamRPN	
			Linear	UpdateNet	Linear	UpdateNet
EAO	0.181	0.273	0.235	0.289	<b>0.439</b>	<b>0.481</b>
A	0.492	0.533	0.529	0.543	<b>0.619</b>	<b>0.610</b>
R	2.934	1.441	1.908	1.320	<b>0.262</b>	<b>0.206</b>

Table 2. **Results for other updating strategies on VOT2016.** DSiam [15] and MemTrack [45] use SiamFC as base tracker. The best two results are highlighted in red and blue fonts, respectively.

formance continuously and rapidly decreases as the update rate increases, evidencing the unsuitability of a fixed and general update rate for all videos.

#### 4.8. Comparison with other updating strategies

Some recent approaches [15, 45] proposed alternative updating strategies for Siamese trackers. Table 2 presents a comparison with DSiam [15] and MemTrack [45] on VOT2016, as [45] only reports results on this VOT edition (see Figure 5 for DSiam results on VOT2018). Our UpdateNet leads to a more effective update and higher tracking performance, while also being applicable to different Siamese architectures. Despite the already excellent performance of DaSiamRPN, UpdateNet brings an improvement of 4.2%, reaching state-of-the-art. Moreover, our approach yields a substantial absolute gain of 5.6% in terms of robustness, which is a common weak point of Siamese trackers.

#### 4.9. LaSOT dataset

We test our model on the recent LaSOT dataset [13]. Since long-term sequences are common in LaSOT, the updating component of the tracker is crucial, as more sudden variations may appear and object appearance may depart further from the input object template. We show the top-10 trackers, including MDNet [32], VITAL [35], StructSiam [49], DSiam [15], SINT [39], STRCF [28], ECO [7], SiamFC [1] and DaSiamRPN [50] in Figure 8. The results are presented following the official protocol. We can see




Figure 8. **Evaluation on LaSOT testing set.** Normalized precision and success plots following the OPE protocol II.

	ATOM [8]	ECO [7]	CFNet [40]	MDNet [32]	SiamFC		DaSiamRPN	
					Linear	UpdateNet	Linear	UpdateNet
Precision (%)	<b>64.8</b>	49.2	53.3	56.5	53.3	53.1	59.1	<b>62.5</b>
Norm. Prec. (%)	<b>77.1</b>	61.8	65.4	70.5	66.3	67.4	73.3	<b>75.2</b>
Success (%)	<b>70.3</b>	55.4	57.8	60.6	57.1	58.4	63.8	<b>67.7</b>

Table 3. **State-of-the-art comparison on TrackingNet.** Our UpdateNet significantly improves DaSiamRPN [50] with an absolute gain of 3.4% and 3.9%, in terms of precision and success. The best two results are highlighted in red and blue fonts, respectively.

how UpdateNet enhances the updating capabilities of DaSiamRPN and leads to a significant performance boost on all measures. As a result, our tracker with learned update strategy surpasses all state-of-the-art trackers on this dataset. This brings further evidence to the advantages of adaptive update strategy in terms of accurate object localization.

#### 4.10. TrackingNet dataset

We evaluate our UpdateNet-DaSiamRPN on the testing set of TrackingNet [31] using their three evaluation metrics (Table 3). Compared with DaSiamRPN, our UpdateNet+DaSiamRPN obtains absolute gains of 3.4%, 1.9% and 3.9% in terms of precision, normalized precision and success. UpdateNet leads to a significant performance improvement on all three metrics. This shows how a learning the model update can greatly benefit Siamese trackers on several datasets and under different measures.

### 5. Conclusions

Siamese trackers usually update their appearance template using a simple linear update rule. We identify several shortcomings of this linear update and propose to learn the updating step as an optimization problem. We employ a neural network, coined UpdateNet, that learns how to update the current accumulated template given the appearance template of the first frame, the current frame, and the accumulated template of the previous step. The proposed UpdateNet is general and can be integrated into all Siamese trackers. Comparable results on four benchmark datasets (VOT2016, VOT2018, LaSOT, and TrackingNet) show that the proposed approach to updating does significantly improve the performance of the trackers with respect to the standard linear update (or with respect to no-update at all).

**Acknowledgements.** We thank the Spanish project TIN2016-79717-R and mention Generalitat de Catalunya CERCA Program.



## References

- [1] Luca Bertinetto, Jack Valmadre, João F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *ECCV workshop*, 2016. 1, 2, 3, 4, 5, 6, 7, 8
- [2] Goutam Bhat, Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. Unveiling the power of deep tracking. In *ECCV*, 2018. 7
- [3] David S Bolme, J Ross Beveridge, Bruce A Draper, and Yui Man Lui. Visual object tracking using adaptive correlation filters. In *CVPR*, 2010. 1, 2, 3
- [4] David S Bolme, Bruce A Draper, and J Ross Beveridge. Average of synthetic exact filters. In *CVPR*, 2009. 3
- [5] Jongwon Choi, Hyung Jin Chang, Tobias Fischer, Sangdoo Yun, Kyuewang Lee, Jiyeoup Jeong, Yiannis Demiris, and Jin Young Choi. Context-aware deep feature compression for high-speed visual tracking. In *CVPR*, 2018. 2
- [6] Janghoon Choi, Junseok Kwon, and Kyoung Mu Lee. Visual tracking by reinforced decision making. *CoRR*, abs/1702.06291, 2017. 3
- [7] Martin Danelljan, Goutam Bhat, F Shahbaz Khan, and Michael Felsberg. Eco: efficient convolution operators for tracking. In *CVPR*, 2017. 1, 2, 3, 8
- [8] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *CVPR*, 2019. 7, 8
- [9] Martin Danelljan, Gustav Hager, Fahad Shahbaz Khan, and Michael Felsberg. Learning spatially regularized correlation filters for visual tracking. In *ICCV*, 2015. 2
- [10] Martin Danelljan, Gustav Hager, Fahad Shahbaz Khan, and Michael Felsberg. Adaptive decontamination of the training set: A unified formulation for discriminative visual tracking. In *CVPR*, 2016. 2
- [11] Martin Danelljan, Andreas Robinson, Fahad Shahbaz Khan, and Michael Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *ECCV*, 2016. 1, 2, 3
- [12] Ali Emami, Farhad Dadgostar, Abbas Bigdeli, and Brian C Lovell. Role of spatiotemporal oriented energy features for robust visual tracking in video surveillance. In *AVSS*, 2012. 1
- [13] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. *CoRR*, abs/1809.07845, 2018. 2, 5, 8
- [14] Heng Fan and Haibin Ling. Siamese cascaded region proposal networks for real-time visual tracking. *CoRR*, abs/1812.06148, 2018. 5
- [15] Qing Guo, Wei Feng, Ce Zhou, Rui Huang, Liang Wan, and Song Wang. Learning dynamic siamese network for visual object tracking. In *ICCV*, 2017. 3, 8
- [16] Anfeng He, Chong Luo, Xinmei Tian, and Wenjun Zeng. A twofold siamese network for real-time object tracking. In *CVPR*, 2018. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [18] David Held, Sebastian Thrun, and Silvio Savarese. Learning to track at 100 fps with deep regression networks. In *ECCV*, 2016. 2
- [19] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *TPAMI*, 37(3):583–596, 2015. 1, 2, 3
- [20] Hamed Kiani Galoogahi, Ashton Fagg, and Simon Lucey. Learning background-aware correlation filters for visual tracking. In *ICCV*, 2017. 2
- [21] Hamed Kiani Galoogahi, Terence Sim, and Simon Lucey. Correlation filters with limited boundaries. In *CVPR*, 2015. 3
- [22] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Luka Čehovin Zajc, Tomáš Vojir, Goutam Bhat, Alan Lukežič, Abdelrahman Eldesokey, et al. The sixth visual object tracking vot2018 challenge results. In *ECCV workshop*, 2018. 2, 5, 6, 7, 8
- [23] Matej Kristan, Jiri Matas, Aleš Leonardis, Tomas Vojir, Roman Pflugfelder, Gustavo Fernandez, Georg Nebehay, Fatih Porikli, and Luka Čehovin. A novel performance evaluation methodology for single-target trackers. *TPAMI*, 38(11):2137–2155, Nov 2016. 5
- [24] Matej Kristan, Roman Pflugfelder, Ale Leonardis, Jiri Matas, Fatih Porikli, Luka Cehovin, Georg Nebehay, Gustavo Fernandez, Toma Vojir, Adam Gatt, et al. The visual object tracking vot2016 challenge results. In *ECCV workshop*, 2016. 2, 5
- [25] Hyemin Lee and Daijin Kim. Salient region-based online object tracking. In *WACV*, 2018. 7
- [26] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. *CoRR*, abs/1812.11703, 2018. 5
- [27] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *CVPR*, 2018. 1, 2, 3
- [28] Feng Li, Cheng Tian, Wangmeng Zuo, Lei Zhang, and Ming-Hsuan Yang. Learning spatial-temporal regularized correlation filters for visual tracking. In *CVPR*, 2018. 7, 8
- [29] Liwei Liu, Junliang Xing, Haizhou Ai, and Xiang Ruan. Hand posture recognition using finger geometric feature. In *ICPR*, 2012. 1
- [30] Alan Lukežic, Tomás Vojir, L Cehovin Zajc, Jiri Matas, and Matej Kristan. Discriminative correlation filter with channel and spatial reliability. In *CVPR*, 2017. 7
- [31] Matthias Muller, Adel Bibi, Silvio Giancola, Salman Alsubaihi, and Bernard Ghanem. Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In *ECCV*, 2018. 5, 8
- [32] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 1, 8
- [33] Eunbyung Park and Alexander C Berg. Meta-tracker: Fast and robust online adaptation for visual object trackers. In *ECCV*, 2018. 3
- [34] Benjamin Renoust, Duy-Dinh Le, and Shinichi Satoh. Visual analytics of political networks from face-tracking of news

- video. *IEEE Transactions on Multimedia*, 18(11):2184–2195, 2016. 1
- [35] Yibing Song, Chao Ma, Xiaohe Wu, Lijun Gong, Linchao Bao, Wangmeng Zuo, Chunhua Shen, Rynson WH Lau, and Ming-Hsuan Yang. Vital: Visual tracking via adversarial learning. In *CVPR*, 2018. 8
  - [36] Chris Stauffer and W Eric L Grimson. Adaptive background mixture models for real-time tracking. In *CVPR*, 1999. 2, 3
  - [37] Chong Sun, Dong Wang, Huchuan Lu, and Ming-Hsuan Yang. Correlation tracking via joint discrimination and reliability learning. In *CVPR*, 2018. 7
  - [38] Chong Sun, Dong Wang, Huchuan Lu, and Ming-Hsuan Yang. Learning spatial-aware regressions for visual tracking. In *CVPR*, 2018. 7
  - [39] Ran Tao, Efstratios Gavves, and Arnold WM Smeulders. Siamese instance search for tracking. In *CVPR*, 2016. 8
  - [40] Jack Valmadre, Luca Bertinetto, João Henriques, Andrea Vedaldi, and Philip HS Torr. End-to-end representation learning for correlation filter based tracking. In *CVPR*, 2017. 2, 3, 5, 7, 8
  - [41] Qiang Wang, Zhu Teng, Junliang Xing, Jin Gao, Weiming Hu, and Stephen Maybank. Learning attentions: residual attentional siamese network for high performance online visual tracking. In *CVPR*, 2018. 1, 2, 3
  - [42] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Online object tracking: A benchmark. In *CVPR*, 2013. 2, 5
  - [43] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Object tracking benchmark. *TPAMI*, 37(9):1834–1848, 2015. 2
  - [44] 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 tracking. *CoRR*, abs/1807.11348, 2018. 7
  - [45] Tianyu Yang and Antoni B Chan. Learning dynamic memory networks for object tracking. In *ECCV*, 2018. 3, 8
  - [46] Yingjie Yao, Xiaohe Wu, Lei Zhang, Shiguang Shan, and Wangmeng Zuo. Joint representation and truncated inference learning for correlation filter based tracking. In *ECCV*, 2018. 3
  - [47] Kaihua Zhang, Lei Zhang, Qingshan Liu, David Zhang, and Ming-Hsuan Yang. Fast visual tracking via dense spatio-temporal context learning. In *ECCV*, 2014. 1
  - [48] Tianzhu Zhang, Changsheng Xu, and Ming-Hsuan Yang. Learning multi-task correlation particle filters for visual tracking. *TPAMI*, 41(2):365–378, 2019. 1, 2, 7
  - [49] Yunhua Zhang, Lijun Wang, Jinqing Qi, Dong Wang, Mengyang Feng, and Huchuan Lu. Structured siamese network for real-time visual tracking. In *ECCV*, 2018. 2, 8
  - [50] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware siamese networks for visual object tracking. In *ECCV*, 2018. 1, 2, 3, 5, 7, 8