

Learning Policies for Adaptive Tracking with Deep Feature Cascades

Chen Huang Simon Lucey Deva Ramanan
Robotics Institute, Carnegie Mellon University
chenh2@andrew.cmu.edu {slucey, deva}@cs.cmu.edu

Abstract

Visual object tracking is a fundamental and time-critical vision task. Recent years have seen many shallow tracking methods based on real-time pixel-based correlation filters, as well as deep methods that have top performance but need a high-end GPU. In this paper, we learn to improve the speed of deep trackers without losing accuracy. Our fundamental insight is to take an adaptive approach, where easy frames are processed with cheap features (such as pixel values), while challenging frames are processed with invariant but expensive deep features. We formulate the adaptive tracking problem as a decision-making process, and learn an agent to decide whether to locate objects with high confidence on an early layer, or continue processing subsequent layers of a network. This significantly reduces the feed-forward cost for easy frames with distinct or slow-moving objects. We train the agent offline in a reinforcement learning fashion, and further demonstrate that learning all deep layers (so as to provide good features for adaptive tracking) can lead to near real-time average tracking speed of 23 fps on a single CPU while achieving state-of-the-art performance. Perhaps most tellingly, our approach provides a **100X** speedup for almost 50% of the time, indicating the power of an adaptive approach.

1. Introduction

Visual Object Tracking (VOT) is a fundamental problem in vision. We consider the single object tracking task, where an object is identified in the first video frame and should be tracked in subsequent frames, despite large appearance changes due to object scaling, occlusion, and so on. VOT is the basic building block of many time-critical systems such as video surveillance and autonomous driving. Thus it is important for a visual tracker to meet the strict constraints of time and computational budget, especially on mobile or embedded computing architectures where real-time analysis perception is often required.

Although much progress has been made in the tracking literature, there still exist tremendous challenges in

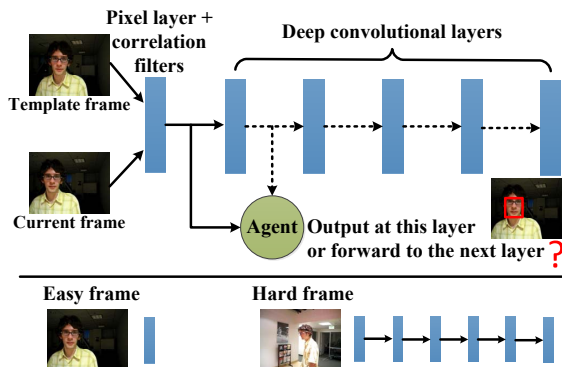


Figure 1. Learning policies for adaptive tracking with deep feature cascades. The agent decides whether we can accurately locate objects on an early layer of the cascade. This avoids the need to always wait for the last layer to make decisions, saving a substantial amount of feed-forwarding time.

designing a tracker that has both high accuracy and high speed. Real-time tracking methods like TLD [21] and correlation filters [19] usually rely on low-level features that are not descriptive enough to disambiguate target and background. Several recent works [15, 27, 11, 31, 10] overcome this limitation by learning correlation filters on hierarchical deep features, but the real-time capacity largely fades away. Other deep trackers [41, 30, 40] take full advantage of the end-to-end learning, and fine-tune the deep network online to achieve top performance. However, even on a high-end GPU, the speed of such trackers is usually around 1 fps which is too slow for practical use.

Two recent deep trackers [3, 18] significantly boost their speed by deploying fixed convolutional networks without any online learning. During tracking, the pre-trained network is simply traversed in a feed-forward pass for similarity evaluation or location regression, allowing to track objects at real-time speeds on GPU. Nevertheless, on a modern CPU, smartphone or tablet with much less computing power, such deep trackers can only process a couple of frames per second, far below the normal video frame-rate 30 fps. Obviously, the major computational burden comes from the forward pass through the entire network, and can

be larger with deeper architectures.

We aim to improve both the accuracy and speed of deep tracking, instead of trading off the quality for speed by *e.g.* using compressed models [1]. We propose to *learn* to speed up deep trackers in an adaptive manner. Our adaptive approach builds on the observation that the tracking complexity varies across frames. For example, it is usually effective to use features from the last layer of a deep network to track objects undergoing large appearance change (*e.g.* abrupt motion) - as these higher-level features are more tolerant to dramatic appearance variation [27]. However, when the object is visually distinct or barely moves, early layers are in most scenarios sufficient for precise localization - offering the potential for substantial computational savings. In the extreme case, the “zeroth” pixel-level layer may suffice for such easy frames, as evidenced by the success of pixel-based correlation filters.

Such an adaptive strategy crucially depends on making the right *decision* - should the tracker stop at the current feature layer or keep computing features on the next layer? To this end, we learn an agent to automatically achieve this goal as illustrated in Fig. 1. The agent learns to find the target at each layer, and decides if it is confident enough to output and stop there. If not, it moves forward to the next layer to continue. This is equivalent to learning a “timely” decision policy to choose the optimal layer for tracking. We will show such policy learning is much more robust than heuristically thresholding the response map of the current active layer. It is also in sharp contrast to the layer selection mechanism in [40], which only selects from two fixed levels of convolutional layers after the entire forward pass is finished. Instead, we formulate this problem as a decision-making process, making *sequential* decisions with early stopping ability. Specifically, we learn the policy in a reinforcement learning [29] fashion during the training phase, and simply apply the policy for adaptive tracking at test time.

In doing so, we are able to provide a speedup of around $10\times$ (on average) over the baseline deep tracker [3] and achieve even higher accuracies on existing OTB [43] and VOT [22] tracking benchmarks. Perhaps most tellingly, our approach provides a $100\times$ speedup almost 50% of the time, indicating the power of an adaptive approach - it turns out that most frames are quite easy to track! Accuracy is improved because each layer of the network is directly trained to be informative for adaptive tracking, similar to the past approaches for “deep supervision” [25]. Concretely, our adaptive tracker works by defining object templates across multiple layers of a network, including the “zeroth” pixel layer. Templates are evaluated across Region of Interest (ROI) with a convolutional filter. We use fast correlation filters [19] to compute response maps for lower layers of our network, where Fourier processing significantly speeds up the convolutional procedure (Fig. 1). We refer to our

approach as EARly-Stopping Tracker (EAST). On a single CPU, it has a near real-time average speed of 23.2 fps, and is about 190 fps almost 50% of the time. This makes it the first CPU-friendly deep tracker among the top benchmark performers. It is worth noting that our policy learning method is quite generic. Further, it is readily applied to train end-to-end with an existing deep network for other time-critical vision tasks besides visual tracking.

2. Related Work

Real-time tracking and correlation filters: Visual tracking methods can rely on either generative (*e.g.* [32]) or discriminative (*e.g.* [16]) models. Discriminative models are often found to outperform in accuracy by discriminating the target from background. Such trackers can usually run fast using hand-crafted features (*e.g.* HOG [9]) and various learning methods of P-N learning [21], structured SVM [16], multi-expert entropy minimization [45], and correlation filter [4].

Among them, Discriminative Correlation Filter (DCF)-based methods [4, 19] are the family of tracking methods with high efficiency and high accuracy as well. The fast speed of DCF is due to its efficient use of all spatial shifts of a training sample by exploiting the discrete Fourier transform. The pioneering MOSSE [4] and improved Kernelized Correlation Filter (KCF) [19] trackers can operate at 669 fps and 292 fps respectively on a single CPU, which far exceeds the real-time requirement. Recent advances of DCF have achieved great success by the use of multi-feature channels [14, 15, 27, 11, 31, 10], scale estimation [13, 26], long-term memory [28], and boundary effect alleviation [11, 12]. However, with increasing accuracy comes a dramatic decrease in speed (0.3–11 fps on a high-end GPU).

Tracking by deep learning: Directly applying correlation filters on the multi-dimensional feature maps of deep Convolutional Neural Networks (CNNs) is one straight-forward way of integrating deep learning for tracking. Usually the deep CNN is fixed, and the DCF trackers trained on every convolutional layer are combined by a hierarchical ensemble method [27] or an adaptive Hedge algorithm [31]. Danelljan *et al.* [15] recently introduced a continuous spatial domain formulation C-COT, to enable efficient integration of multi-resolution deep features. C-COT and its improved ECO [10] can achieve top performance in the VOT challenge [22], but the tracking speed is still slow due to the high dimension of the entire deep feature space.

Another category of deep trackers [41, 30, 40] update a pre-trained CNN online to account for the target-specific appearance at test time. Such trackers usually take a classification approach to classify many patches and choose the highest scoring one as the target object. Unfortunately, online training and exhaustive search severely hamper their speed. The top performing tracker MDNet [30] has the

GPU speed of only 1 fps or so. Recent advances include using Recurrent Neural Networks (RNNs) [20, 7] to model temporal information using an attention mechanism, but the speed remains slow.

One common reason for the slow speed of the above-mentioned deep trackers is they always conduct a complete feed-forward pass to the last CNN layer. This ignores the fact that the tracking complexity differs across varying conditions. One of our conclusions is that most frames in current video benchmarks are rather easy. For those frames, forwarding to only early layers may suffice. In principle, such an insight can be used to speed up many recent real-time deep trackers, such as GOTURN [18] (165 fps on GPU) and SiamFC [3] (86 fps on GPU), to make them more CPU-friendly with near frame-rate speed.

Feature selection in tracking: Good features are important to tracking. The initial DCF trackers were limited to a single feature channel, *e.g.* a grayscale image in MOSSE [4]. The DCF framework was later extended to multi-channel features such as HOG [13, 19], Haar-like features [16], binary patterns [21] and Color Attributes [14]. Generally, the hand-crafted features are cheap to compute, but they are not discriminative enough to handle complex tracking scenarios. Many recent deep trackers (*e.g.* [41, 30]) exploit the semantically robust features from the last CNN layer (fully-connected). However, spatial details of the tracked object are lost in the last layer which is not optimal for visual tracking. Danelljan *et al.* [11] found the first convolutional layer is very suitable for tracking. Other works [15, 27, 31, 10] choose to utilize all the hierarchical convolutional features, where early layers can preserve high spatial resolution and deep layers are more discriminative.

In this paper, we make the best of hand-crafted and deep convolutional features in a cascaded structure, and learn an agent to select a minimum sequence of feature layers for fast tracking purposes. Unlike FCNT [40] that selects features from two pre-defined layers only after a complete forward pass, our selection is sequential and can stop early at any layer with enough confidence.

Feature cascades: CNNs are a natural cascaded architecture with increasingly abstract feature representations. Contemporary works either improve the cascade’s optimality by deep layer supervision [25], or stack multiple CNNs into a deeper cascade for coarse-to-fine [35] or multi-task [8] predictions. Our work differs in learning a decision policy of using only early feature layers in a cascade, and in combining feature cascades and reinforcement learning [36] to achieve this goal. Our approach bears some similarity to the “attentional cascade” structure [39] which uses a cascade of gradually more complex classifiers. The difference is that attentional cascade aims to use early classifiers to eliminate easy negatives and reduce the burden of complex classifier evaluation, whereas we aim to use these early layers to make

a strong decision as early as possible.

Reinforcement learning for tracking: Reinforcement learning (RL) [29, 36] is capable of learning good policies to take a sequence of actions based on trail and error. It has been successfully applied to vision tasks (*e.g.* object detection [5]) by treating them as a decision-making process. For visual tracking, there are two recent works that use RL to temporally attend target regions [44] and select suitable template [6]. Our work is the first one to use RL to learn an early decision policy for speeding up deep tracking.

3. Method

We use the deep Convolutional Neural Network (CNN) as a rich feature cascade for robust visual tracking. Given an input video frame at time t , the process of tracking an object with bounding box can be seen as the problem of maximizing a confidence function $f_l : H_t \rightarrow \mathbb{R}$ over the set of hypothesized object regions R_t :

$$r_t^* = \arg \max_{r \in H_t} f_l(r), \tag{1}$$

where H_t usually consists of regions around the object location in the previous frame, and l denotes the feature layer at which the object confidence is obtained. Many deep trackers [41, 30] exploit $l = L$ as the last fully-connected layer or last convolutional layer of AlexNet [24] or deeper VGG-Net [34]. Other trackers [15, 27, 31, 10] exploit a full layer set $\{l\}_{l=1}^L$ of all the convolutional layers to take the best advantage of feature hierarchies. While these methods have been successful and effective, they are still slow and may not be needed when tracking during easy frames.

In this paper, we propose a principled sequential method that accumulates confidence $f_{l=1, \dots, L_t}(r)$ from an adaptively small set of feature layers $\{l\}_{l=1}^{L_t}$ in order to track efficiently without losing accuracy. For example, $L_t = 2$ convolutional layers $C1-C2$ can suffice to track a distinct face in Fig. 3; but for a blurry face in cluttered backgrounds, we may want to gather more evidence from a deeper layer $C5$. During such sequential search, our method needs to adjust the bounding box to progressively localize the object using more and more robust features. Ideally, we would like to minimize the number of forwarded layers necessary to locate objects. A naive approach might use a heuristic for determining when to advance to the next layer: for example, one might advance if the maximum value of the current response map is below a threshold. However, defining such heuristics can be difficult when the response map is ambiguous or has multiple peaks (*e.g.* Fig. 3). Instead, we propose to train a functioning agent end to end by deep reinforcement learning [29]. The agent learns the action and search policy (including the early stopping criteria) so that it can make decisions under uncertainty to reach the target. Fig. 2 shows the framework of our policy learning.

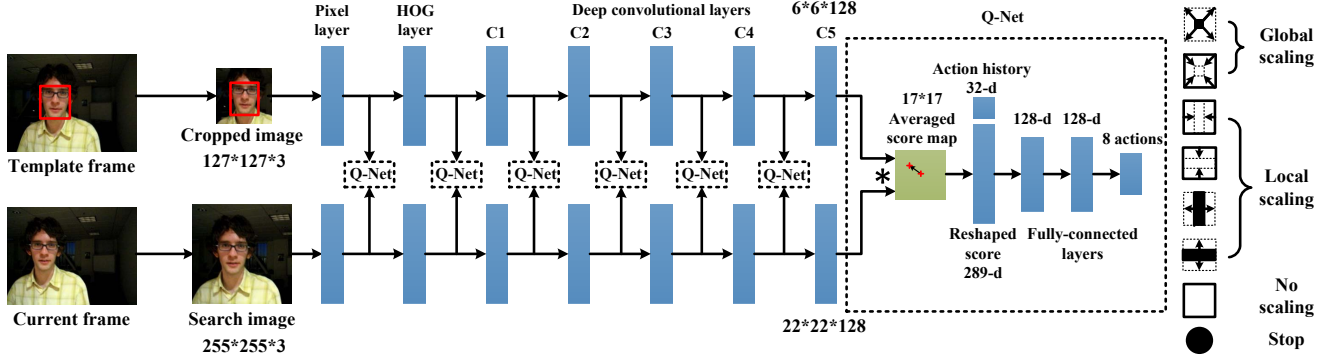


Figure 2. System framework of our EARly-Stopping Tracker (EAST) by policy learning.




Figure 3. Example sequences of actions taken to adjust the bounding box scale over deep convolutional layers. Box translation is determined by the relative position of maximum score on the score map. Note each score map is averaged with all the maps from preceding layers. Our agent learns to wisely act upon the score maps. It terminates the search early if the confidence score is highly peaked. When the score map is ambiguous (*e.g.* has two peaks for a blurry face in cluttered backgrounds), the agent postpones decision and enlarges the box according to the more unambiguous score map at the next layer. Further actions of box scaling are performed with stronger confidence.

3.1. Fully-Convolutional Siamese Network

In this section, we review the fully-convolutional Siamese tracker [3], which we make use of for its good trade-off between speed (86 fps on GPU) and accuracy in exploiting deep convolutional layers. Other trackers are either too slow [15, 27, 31, 10] as a baseline (<11 fps on GPU), or do not provide explicit response maps (direct regression instead in [18]) for our policy learning purposes.

The Siamese network [3] is trained offline to locate a 127×127 template image z within a larger 255×255 search image x . A similarity function is learned to compare the template image z to a candidate region of the same size in the search image x , so as to return a high score for the truth region and low score otherwise. Such similarity evaluation is fully-convolutional with respect to x in the network, much more efficient than exhaustive search. Specifically, a cross-correlation layer is proposed to compute similarity for all translated sub-regions in x in one pass:

$$F_l(z, x) = \varphi_l(z) * \varphi_l(x) + v\mathbb{I}, \quad (2)$$

where φ_l is the convolutional feature embedding at layer

$l = 5$ (*i.e.* C5 layer), and $v \in \mathbb{R}$ is an offset value. Here $F_l(\cdot, \cdot)$ is a confidence score map of size 17×17 , as opposed to the single confidence score $f_l(\cdot)$ in Eq. 1.

During tracking, this Siamese network simply evaluates the similarity online between the template image in previous frame and the search region in current frame, leading to fast speed. The relative position of the maximum score multiplied by the stride of the network, gives the object translation from frame to frame.

Our goal is to learn an early decision policy from these confidence score maps F_l , to adaptively 1) predict the object bounding box across layers, and 2) stop early at a layer $l < 5$ when sufficiently confident about the prediction. Note the score map dimension depends on the size of input feature maps. The score map F_l on an early layer l will have a larger resolution than 17×17 , and so we downsample to this size to facilitate learning. Also, SiamFC [3] searches over multiple scales of the search image to handle scale variations. We only work on the *original* scale for high efficiency, and learn to gradually infer the box scale from heatmaps computed during the single forward-pass.

3.2. Learning Policies with Reinforcement Learning

We treat the tracking problem as a Markov Decision Process (MDP) where an agent can make a sequence of actions across feature layers, see Fig. 2. This agent learns when to stop advancing to the next layer, as well as how to gradually deform a bounding box once per layer, which is initialized to the estimated box from the previous frame. The ultimate goal is to output a tight box around the object with as few as layers as possible. The challenge is to be able to operate with rewards that rule out supervision at each step, and at the same time, to minimize the number of steps to locate the target given its changing complexity.

We train the agent in a reinforcement learning (RL) setting to learn decision policies. In the RL setting there are a set of states S and actions A , and a reward function R . At each step on layer l , the agent examines the current state S_l , and decides on the action A_l to either transform the box or *stop* with a box output. The action A_l is expected to reduce the uncertainty in localizing the object, and receives positive or negative reward R_l reflecting how well the current box covers the object and how few steps are used before action *stop*. By maximizing the expected rewards, the agent learns the best policy to take actions and can explicitly balance accuracy (search for more layers) and efficiency (stop early if highly confident).

Actions: Our action set A includes seven anisotropic scaling transformations to a box and one *stop* action to terminate search. We do not use the agent to predict the centroid of the box, and instead compute it directly from the relative position of the maximum score on score map as in Eq. 1. We also experimented with requiring the agent to report box translations, but found directly inferring them from the score maps simplified training and increased convergence (due to the smaller space of actions A).

For scaling actions, there are two global and four local (modify aspect ratio) transformations as shown in Fig. 2. Similar to [5], any of these actions makes a combined horizontal and vertical change to the box by a factor of 0.2 relative to its current size. We also introduce a *no scaling* action that does not scale the box at all. This action allows the agent to postpone decision when the current score map is ambiguous or a decision simply cannot be made. Fig. 3 exemplifies this case where two peaks exist on the first-layer score map for a cluttered scene. The agent decides not to act on this map but waits for a more unambiguous map at the next layer to act (enlarge the box).

States: The state S_l is represented as a tuple (F'_l, h_l) , where F'_l is the score map and h_l is a vector of the history of taken actions. We define $F'_l = \sum_{k=1}^l F_k/l$ as the average of the score map at current layer l and all its preceding maps from earlier layers. Thus F'_l not only encodes the currently observed confidence but also the confidence history that has

been collected. This is found to work better empirically than using F_l only, and is similar to the hypercolumn representation [17] with the benefits of simultaneously leveraging information from early layers that capture fine-grained spatial details and deeper layers that capture semantics. Also, the resulting robustness comes at a negligible cost in averaging the score maps already obtained. The history vector h_l keeps track of the past 4 actions. Each action in the vector is represented by an 8-dimensional one-hot vector or zero vector (when processing the first layer). We find that including $h_l \in \mathbb{R}^{32}$ helps stabilize action trajectories.

Rewards: The reward function $R(S_{l-1}, S_l)$ reflects the localization accuracy improvement from state S_{l-1} to S_l after taking a particular action A_l . The accuracy is measured by the Intersection-over-Union (IoU) between the predicted box b and ground-truth box g . We can formally define IoU as $IoU(b, g) = area(b \cap g)/area(b \cup g)$. Since each state S_l is associated with a box b_l , the reward function is then defined following [5]:

$$R(S_{l-1}, S_l) = \begin{cases} \text{sign}(IoU(b_l, g) - IoU(b_{l-1}, g)), & A_l \neq \text{stop} \\ +3, & IoU(b_l, g) \geq 0.6, \quad A_l = \text{stop} \\ -3, & IoU(b_l, g) < 0.6, \quad A_l = \text{stop}, \end{cases} \quad (3)$$

where the accuracy improvement is quantized to ± 1 if the current action is not *stop*. This reward scheme encourages positive transformations even with small accuracy improvement. If there is no transformation for further improvement, or if the agent already arrives at the last layer $l = L$, the action should be *stop*. In this case, the reward function will penalize the predicted box b_l with IoU less than 0.6. Note such a reward scheme implicitly penalizes a large number of layers l since Q-learning (detailed next) models the expected future rewards when deciding on an action (positive or negative).

Deep Q-learning: The optimal policy of selecting actions should maximize the sum of expected rewards on a given frame. Since we do not have *a priori* knowledge about the correct layer or action to choose, we address the learning problem through deep Q-learning [29]. This approach learns an action-value function $Q(S_l, A_l)$ to select the action A_{l+1} that gives the highest reward at each layer. The learning process iteratively updates the action-selection policy by:

$$Q(S_l, A_l) = R + \gamma \max_{A'} Q(S', A'), \quad (4)$$

where $Q(S', A')$ is the future reward and γ the discount factor. The function $Q(S, A)$ is learned by a deep Q-Network as illustrated in Fig. 2. It takes as input the state representation S , *i.e.* the reshaped score vector and action history vector. The network consists of two 128-dimensional fully-connected layers, finally mapping to 8 actions. Each fully-

connected layer is randomly initialized, and is followed by ReLU and dropout regularization [24].

Note during training, we not only update the weights of the Q-Network, but also the pre-trained convolutional layer when the agent receives rewards on that layer. Similar to the deeply-supervised net [25], our approach provides a direct target signal for learning the feature representation at each layer, so as to improve performance of our adaptive tracker.

Testing with learned policies: During testing, the agent does not receive rewards or update the Q-function. It just follows the decision policy to deform the box and output it when a *stop* action is performed. Our agent takes only 2.1 steps on average to locate the target between frames on the OTB-50 dataset [42]. This means we can correctly track most objects by using 2 deep layers. Only for those hard frames, the search degenerates to a full forward pass. The overall tracking algorithm follows SiamFC [3] to search over candidate regions around the estimated location from previous frame. As a result, one order of magnitude speedups are achieved over traditional non-adaptive deep trackers. Set aside efficiency, the policies are still appealing in that they mimic the dynamic *attention* mechanism by progressively attending to the target region in feature cascades.

Implementation details: We use the AlexNet [24]-like convolutional architecture as in SiamFC [3]. The whole network including Q-Net is trained on the ImageNet Video [33] trainval set (4417 videos) for 50 epochs, each completed after the agent has interacted with all training images. We make use of an ϵ -greedy [36] optimization during Q-learning, taking a random action with probability ϵ to encourage exploration of diverse action policies. We anneal ϵ linearly from 1 to 0.1 over the first 30 epochs, and fix ϵ to 0.1 in the remaining 20 epochs. We use a learning rate $1e-3$, discount factor $\gamma = 0.9$ and batch size 64. The network parameters are updated with direct stochastic gradient descent using MatConvNet [38] on a single NVIDIA GeForce Titan X GPU and an Intel Core i7 CPU at 4.0GHz.

3.3. Learning with cheap features

Our policy learning can be applied to a feature cascade with any type of feature layers. We explore the use of additional cheap feature layers after the pixel-layer and before the more expensive deep layers. Inspired by the success of correlation filters defined on multi-channel HOG layers [9], we explore an optional HOG layer. In theory, other cheap features such as Color Attributes [14] might apply. When processing our pixel and HOG layers, we make use of fast correlation filters. We specifically make use of the Dual Correlation Filter (DCF) [19], which exhibits a good trade-off in CPU speed (270+ fps on CPU) and accuracy compared to alternatives such as the Kernelized Correlation Filter (2× slower), STC [46] (350 fps but lower accuracy) and SRDCF [12] (5 fps).

4. Results

Before comparing our EARly-Stopping Tracker (EAST) with prior works, we first conduct an ablation study of some of its important variants. We compare EAST to variants using different feature cascades—EAST-Pixel-HOG, EAST-HOG, EAST-Pixel, as well as the baseline SiamFC [3]. Fig. 4 shows the probability of stopping at different feature layers and the associated speed on OTB-50 [42]. EAST indeed learns to use only early layers over 70% of the time, as they suffice when tracking easy frames. Hard frames are processed with additional layers, degenerating to a full forward-pass (as in SiamFC) only when needed. Fig. 5 illustrates some examples of easy and hard frames (and their stopped layers) on particular video sequences.

Obviously, the earlier layer we stop at, the greater the speedup. EAST-Pixel-HOG (which lacks a pixel or HOG layer) is about $4.5\times$ faster than SiamFC on average, running at 10.7 fps on a CPU and 467.3 fps on a GPU. Recall that Pixel and HOG layers can be processed by fast correlation filters [19] that run at 278 fps and 292 fps, respectively. By adding such layers, EAST achieves a CPU speedup of $10\times$ on average and $100\times$ for those easy frames. Even though HOG may incur an additional feature computation cost compared to the Pixel layer, it produces a greater speedup because it enables more accurate pruning, and so is selected more often by EAST. Note that the average GPU speed is sacrificed to some extent due to larger reliance on CPU computations. Nonetheless, EAST still produces a near-real-time CPU rate of 23.2 fps, making it quite practical for CPU-bound tracking (required on many embedded devices).

Table 1 (top) summarizes the speed and accuracy of these EAST variants. We use the Overlap Success (OS) rate as a strict evaluation metric for accuracy. One-Pass Evaluation (OPE) is employed to compare accuracies in terms of Area Under the Curve (AUC) of OS rates for different threshold values. Table 1 shows that the use of *more* and *stronger* feature layers systematically improves the accuracy of the tracker, reaching to the AUC score of 0.638 for our full EAST model. Speed is also improved due to the larger computational savings on cheap layers.

Table 1 (middle) further compares other EAST variants:

- EAST_{last}: tracking by always forwarding to the last feature layer.
- EAST_{th}: feed-forwarding if the maximum value of the current response map is lower than 0.9. Scale is determined by the size of thresholded region.
- EAST_{ch}: policy learning without confidence history, *i.e.* we use F_l instead of $F'_l = \sum_{k=1}^l F_k/l$.
- EAST_{ah}: policy learning without action history h_l .

EAST_{last} is similar to the baseline SiamFC [3] in that all layers are always processed, but differs in that it is trained




Figure 4. Stopping probability at different feature layers and the resulting speed on OTB-50 dataset [42]. For each of the 5 models with different feature cascades, we show the average speed on both CPU and GPU (top), and the CPU speedup ratio over the baseline SiamFC [3] (1.9 fps) at each layer (bottom). SiamFC searches over multi-scaled images to handle scale variations, while we predict the scale in a single forward-pass, leading to a constant tracking speedup. Our early stopping policy further accelerates tracking (4.5 \times) by adaptively using early layers as compared to SiamFC that always uses the last layer C5. When the early layer is the cheap HOG or pixel layer with fast CPU speed (270+ fps), we are able to increase the average CPU speed by one order of magnitude and operate at around 100 \times faster speed for almost 50% of the time. Our full model EAST operates at near real-time speed 23.2 fps on CPU. On the other hand, it is found the more reliance on CPU computations will generally increase the CPU speed, but also lose the benefits of GPU speed to some acceptable extent.

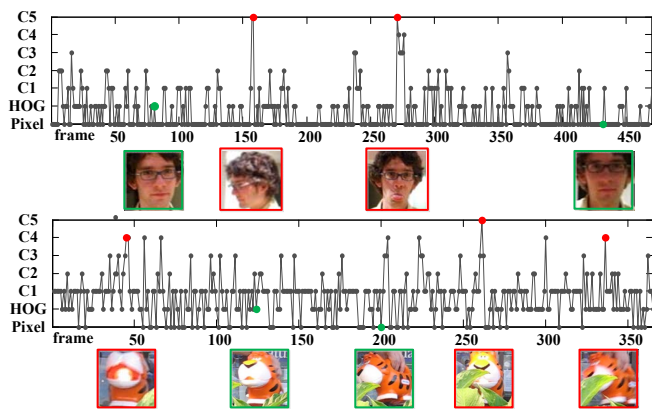


Figure 5. Stopped feature layers for the frames in videos *David* (frames 300-770 of the original video) and *Tiger2*. Easy and hard frames are marked in green and red.

with deep supervision and does not require an image pyramid to model scale. As a result, it both works better and faster (in CPU terms) than SiamFC. Interestingly, EAST_{last} outperforms EAST but is much slower. This suggests that our adaptive strategy is slightly hurting accuracy. EAST_{th} makes use of manually-designed heuristics for stopping, avoiding the need to predict actions with a Q-Net, making it 2 \times faster on CPU. However, heuristic policies fail to work as well as the one learned by Q-learning. Finally, eliminating scoremap and action histories also hurts performance, likely because such histories stabilize the search process.

OTB-50 results: The OTB-50 [42] benchmark contains 50 video sequences for evaluation. Table 2 compares the AUC scores of our EAST and the state-of-the-art trackers: TLD [21], GOTURN [18], Struck [16], KCF [19], DSST [13], MEEM [45], RTT [7], FCNT [40], Staple [2], HDT [31], HCF [27], LCT [28], SiamFC [3] and SINT [37].

Table 1. One-Pass Evaluation (OPE) results of Area Under the Curve (AUC) score and speed (CPU/GPU fps) on OTB-50 dataset.

Method	AUC	Speed
EAST-Pixel-HOG	0.619	10.7 / 467.3
EAST-HOG	0.628	14.5 / 197.6
EAST-Pixel	0.631	19.9 / 223.6
EAST _{last}	0.645	3.4 / 76.3
EAST _{th}	0.597	45.7 / 172.0
EAST _{ch}	0.621	21.1 / 137.2
EAST _{ah}	0.610	17.4 / 127.6
SiamFC [3]	0.612	1.9 / 86
EAST	0.638	23.2 / 158.9

CPU/GPU speeds (fps) are also reported.

EAST achieves the highest AUC of any method. It does so while being significantly faster. For example, the runner-up SINT is 4 fps on GPU, while other GPU-based real-time deep trackers (GOTURN and SiamFC) are significantly slower on a CPU (2-3 fps). When compared to fast correlation trackers *e.g.* KCF defined on cheap features, EAST is significantly more accurate while still maintaining near-real-time speeds. This is in contrast to correlation filters defined on deep features (HCF and HDF), who are not real-time even with a GPU. The Staple tracker combines HOG and color features by a ridge regression, while FCNT tracks by selecting features from deep layers. EAST outperforms both in terms of accuracy. The Recurrently Target-attending Tracker (RTT) trains Recurrent Neural Networks (RNNs) to capture attentions as a regularization on correlation filter maps. However, it is noticeably slower and less accurate than EAST (3 fps on CPU, and an AUC of 0.588).

OTB-100 results: The OTB-100 [43] dataset is the extension of OTB-50 and is more challenging. We test on the full 100 videos to compare with recent related trackers: RDT [6], SRDCF [12], MDNet [30], C-COT [15], ECO [10]. Table 3 summarizes their AUC scores and

Table 2. The Area Under the Curve (AUC) score for One-Pass Evaluation (OPE), and speed (fps, * indicates GPU speed, otherwise CPU speed) on the OTB-50 dataset. The best results are shown in bold.

Method	TLD [21]	GOTURN [18]	Struct [16]	KCF [19]	DSST [13]	MEEM [45]	RTT [7]	FCNT [40]	Staple [2]	HDT [31]	HCF [27]	LCT [28]	SiamFC [3]	SINT [37]	EAST
AUC	0.438	0.450	0.474	0.516	0.554	0.572	0.588	0.599	0.600	0.603	0.605	0.612	0.612	0.625	0.638
Speed	22	165*	10	172	24	10	3	3*	80	10*	11*	27	86*	4*	23/159*

Table 3. The Area Under the Curve (AUC) score for One-Pass Evaluation (OPE), and speed (fps, * indicates GPU speed, otherwise CPU speed) on the OTB-100 dataset. The best results are shown in bold.

Method	RDT [6]	SRDCF [12]	MDNet [30]	C-COT [15]	ECO [10]	EAST
AUC	0.603	0.605	0.685	0.686	0.694	0.629
Speed	43*	5	1*	0.3	6	23/159*

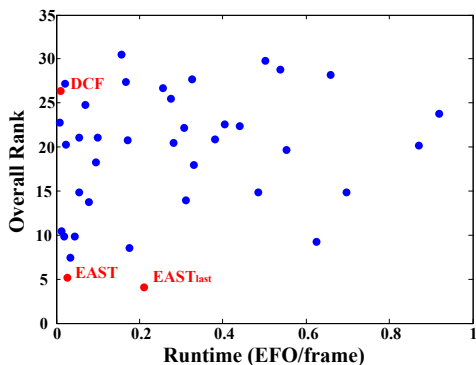


Figure 6. Overall rank and runtime of our tracker (red) in comparison to the 38 trackers from VOT-14 Challenge. We show the results of EAST and its two variants that always forward to the first (*i.e.* DCF on pixel layer) and last (*i.e.* EAST_{last}) feature layer.

CPU/GPU speeds. EAST is close to state-of-the-art in terms of accuracy and is the fastest among the top performers on OTB-100. The MDNet and correlation-filter-based advances SRDCF, C-COT and ECO all suffer from low speed, while EAST does not sacrifice the run-time performance for accuracy. RDT is a related Reinforcement Learning (RL)-based method that selects the best template to track a given frame. EAST (which also makes use of RL) proves more accurate while being 4× faster on a GPU.

VOT-14 results: We test EAST on the 25 videos from the VOT-14 [23] Tracking Challenge. Fig. 6 compares the speed and overall rank of EAST to that of the 38 tracker submissions. For the rank computation, the challenge first evaluates the average accuracy and robustness ranking (refer to the supplementary material for details) for each tracker, and then averages the two rank values to get the overall rank. As can be seen from Fig. 6, EAST achieves the best accuracy-speed trade-off, outperforming all competitors in the overall rank. We also show two EAST variants at the opposite extreme: always forwarding to the first (*i.e.* DCF [19] on pixel layer) and last (*i.e.* EAST_{last}) feature layer. Our EAST is able to adaptively take advantage of the speed and accuracy benefits of two variants by policy learning.

Table 4. Raw scores and speed for our method and the top 4 trackers of the VOT-15 Challenge. The * indicates the speed in EFO units. The CPU/GPU speeds for our EAST are given.

Tracker	Accuracy	Robustness	Overlap	Speed (fps)
MDNet [30]	0.60	0.69	0.38	1
EAST	0.57	1.03	0.34	21/148
DeepSRDCF [11]	0.56	1.05	0.32	<1*
EBT [47]	0.47	1.02	0.31	5
SRDCF [12]	0.56	1.24	0.29	5

VOT-15 results: The VOT-15 [22] Tracking Challenge has 60 testing videos chosen from a pool of 356. Trackers are automatically re-initialized five frames after failure (zero overlap). Table 4 compares our EAST with the top 4 trackers in terms of accuracy and speed (using the vot2015-challenge toolkit). Our testing speed on this benchmark are 21 fps on CPU and 148 fps on GPU, making EAST the fastest and most CPU-friendly tracker among the top performers. We achieve comparable accuracy scores to MDNet, while providing a 148× speedup on GPU, indicating the power of our adaptive policy learning approach.

5. Conclusion

This paper proposes an adaptive approach to tracking with deep feature cascades. Our fundamental insight is that most frames in typical tracking scenarios turn out to be easy, in that simple features (such as pixels or HOG) suffice. That said, some challenging frames do require “heavy-duty” invariant feature processing. The challenge is in determining which is which! By formulating the tracking problem as a decision making process, we learn a reinforcement-learning agent that can make such distinctions. Importantly, the agent learns to do so in an iterative manner, making efficient use of a feature cascade that proceeds to deeper layers only when the current one does not suffice. This dramatically reduces the feed-forwarding cost for those easy frames (by 100X), leading to an overall significant speedup for tracking. Such a policy learning method is appealing in that it is trained end-to-end and can be applied to any deep network designed for time-critical tasks.

Acknowledgment: This research was supported in part by NSF grants CNS-1518865, IIS-1618903 and IIS-1526033 and the DARPA grant HR001117C0051. Additional support was provided by the Intel Science and Technology Center for Visual Cloud Systems (ISTC-VCS), Google, and Autel. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the view(s) of their employers or the above-mentioned funding sources.

References

- [1] A. Alfvén. Object tracking with compressed networks. *Master's Theses in Mathematical Sciences*, Mathematics (Faculty of Technology) and Numerical Analysis, Lund University, 2016.
- [2] L. Bertinetto, J. Valmadre, S. Golodetz, O. Miksik, and P. H. S. Torr. Staple: Complementary learners for real-time tracking. In *CVPR*, 2016.
- [3] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. S. Torr. Fully-convolutional siamese networks for object tracking. In *ECCV*, 2016.
- [4] D. S. Bolme, J. R. Beveridge, B. A. Draper, and Y. M. Lui. Visual object tracking using adaptive correlation filters. In *CVPR*, 2010.
- [5] J. C. Caicedo and S. Lazebnik. Active object localization with deep reinforcement learning. In *ICCV*, 2015.
- [6] J. Choi, J. Kwon, and K. M. Lee. Visual tracking by reinforced decision making. *arXiv preprint*, 1702.06291, 2017.
- [7] Z. Cui, S. Xiao, J. Feng, and S. Yan. Recurrently target-attending tracking. In *CVPR*, 2016.
- [8] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *CVPR*, 2016.
- [9] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005.
- [10] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg. ECO: efficient convolution operators for tracking. In *CVPR*, 2017.
- [11] M. Danelljan, G. Häger, F. S. Khan, and M. Felsberg. Convolutional features for correlation filter based visual tracking. In *ICCV Workshop*, 2015.
- [12] M. Danelljan, G. Häger, F. S. Khan, and M. Felsberg. Learning spatially regularized correlation filters for visual tracking. In *ICCV*, 2015.
- [13] M. Danelljan, G. Hger, F. Shahbaz Khan, and M. Felsberg. Accurate scale estimation for robust visual tracking. In *BMVC*, 2014.
- [14] M. Danelljan, F. S. Khan, M. Felsberg, and J. v. d. Weijer. Adaptive color attributes for real-time visual tracking. In *CVPR*, 2014.
- [15] M. Danelljan, A. Robinson, F. Shahbaz Khan, and M. Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *ECCV*, 2016.
- [16] S. Hare, A. Saffari, and P. H. S. Torr. Struck: Structured output tracking with kernels. In *ICCV*, 2011.
- [17] B. Hariharan, P. Arbelaz, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *CVPR*, 2015.
- [18] D. Held, S. Thrun, and S. Savarese. Learning to track at 100 fps with deep regression networks. In *ECCV*, 2016.
- [19] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. High-speed tracking with kernelized correlation filters. *TPAMI*, 37(3):583–596, 2015.
- [20] S. E. Kahou, V. Michalski, and R. Memisevic. RATM: recurrent attentive tracking model. In *CVPRW*, 2017.
- [21] Z. Kalal, K. Mikolajczyk, and J. Matas. Tracking-learning-detection. *TPAMI*, 34(7):1409–1422, 2012.
- [22] M. Kristan, J. Matas, A. Leonardis, M. Felsberg, L. Čehovin, G. Fernandez, T. Vojir, G. Häger, G. Nebehay, and R. Pflugfelder. The visual object tracking VOT2015 challenge results. In *ICCV 2015 Workshop*, 2015.
- [23] M. Kristan, R. Pflugfelder, A. Leonardis, J. Matas, L. Čehovin, G. Nebehay, T. Vojir, G. Fernandez, A. Lukežič, and A. Dimitriev. The visual object tracking VOT2014 challenge results. In *ECCV 2014 Workshop*, 2014.
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [25] C.-Y. Lee, S. Xie, P. W. Gallagher, Z. Zhang, and Z. Tu. Deeply-supervised nets. In *AISTATS*, 2015.
- [26] Y. Li and J. Zhu. A scale adaptive kernel correlation filter tracker with feature integration. In *ECCV Workshop*, 2014.
- [27] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang. Hierarchical convolutional features for visual tracking. In *ICCV*, 2015.
- [28] C. Ma, X. Yang, C. Zhang, and M.-H. Yang. Long-term correlation tracking. In *CVPR*, 2015.
- [29] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, and A. K. Fidjeland et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [30] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016.
- [31] Y. Qi, S. Zhang, L. Qin, H. Yao, Q. Huang, J. Lim, and M.-H. Yang. Hedged deep tracking. In *CVPR*, 2016.
- [32] D. A. Ross, J. Lim, R.-S. Lin, and M.-H. Yang. Incremental learning for robust visual tracking. *IJCV*, 77(1-3):125–141, 2008.
- [33] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015.
- [34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [35] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *CVPR*, 2013.
- [36] R. Sutton and A. Barto. *Reinforcement Learning: An introduction*. MIT Press, 1998.
- [37] R. Tao, E. Gavves, and A. W. M. Smeulders. Siamese instance search for tracking. In *CVPR*, 2016.
- [38] A. Vedaldi and K. Lenc. Matconvnet: Convolutional neural networks for matlab. In *ACM MM*, 2015.
- [39] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR*, 2001.
- [40] L. Wang, W. Ouyang, X. Wang, and H. Lu. Visual tracking with fully convolutional networks. In *ICCV*, 2015.
- [41] N. Wang, S. Li, A. Gupta, and D. Yeung. Transferring rich feature hierarchies for robust visual tracking. *arXiv preprint*, 1501.04587, 2015.
- [42] Y. Wu, J. Lim, and M.-H. Yang. Online object tracking: A benchmark. In *CVPR*, 2013.
- [43] Y. Wu, J. Lim, and M. H. Yang. Object tracking benchmark. *TPAMI*, 37(9):1834–1848, 2015.

- [44] D. Zhang, H. Maei, X. Wang, and Y.-F. Wang. Deep reinforcement learning for visual object tracking in videos. *arXiv preprint*, 1701.08936, 2017.
- [45] J. Zhang, S. Ma, and S. Sclaroff. MEEM: robust tracking via multiple experts using entropy minimization. In *ECCV*, 2014.
- [46] K. Zhang, L. Zhang, Q. Liu, D. Zhang, and M.-H. Yang. Fast visual tracking via dense spatio-temporal context learning. In *ECCV*, 2014.
- [47] G. Zhu, F. Porikli, and H. Li. Tracking randomly moving objects on edge box proposals. *arXiv preprint*, 1507.08085, 2015.