



# Depth-Adaptive Computational Policies for Efficient Visual Tracking

Chris Ying<sup>1(✉)</sup> and Katerina Fragkiadaki<sup>2(✉)</sup>

<sup>1</sup> Google Brain, Mountain View, CA, USA  
chrisying@google.com

<sup>2</sup> Machine Learning Department, CMU, Pittsburgh, PA, USA  
katfef@cs.cmu.edu

**Abstract.** Current convolutional neural networks algorithms for video object tracking spend the same amount of computation for each object and video frame [3]. However, it is harder to track an object in some frames than others, due to the varying amount of clutter, scene complexity, amount of motion, and object’s distinctiveness against its background. We propose a depth-adaptive convolutional siamese network that performs video tracking adaptively at multiple neural network depths. Parametric gating functions are trained to control the depth of the convolutional feature extractor by minimizing a joint loss of computational cost and tracking error. Our network achieves accuracy comparable to the state-of-the-art on the VOT2016 benchmark. Furthermore, our adaptive depth computation achieves higher accuracy for a given computational cost than traditional fixed-structure neural networks. The presented framework extends to other tasks that use convolutional neural networks and enables trading speed for accuracy at runtime.

**Keywords:** Visual tracking · Metric learning  
Conditional computation · Deep learning


## 1 Introduction

Multilayer neural networks are the defacto standard machine learning tools for many tasks in computer vision, including visual tracking [3]. Current visual trackers use a fixed amount of computation for every object and video frame [3, 14, 20–22]. However, different video scenes have varying amount of complexity, background clutter, object motion, camera motion, or frame rate. Fixed compute architectures do not adapt to the difficulty of the input and are can be suboptimal computation-wise.

In this work, we propose neural architectures whose computation adapts to the difficulty of the task from frame to frame, rather than being fixed at runtime. We focus on the task of visual tracking in videos. We present learning

---

C. Ying—Work done as student at the Machine Learning Department, CMU.



**Fig. 1.** Adaptive neural computational policies for visual tracking. At each frame, we match the key frame depicting the object of interest to a search region cropped around the location of the detected object in the previous frame. Our controllers (gates) decide how many blocks of the VGG net [18] (divided into 5 blocks of layers) to compute before computing a cross-correlation map and determining the target location. In general, deeper layers yield more accurate predictions but also require more computational power. Our depth-adaptive model picks the first depth for the uncluttered scene on the left and the fourth for the cluttered scene on the right. Red gate color denotes halting of computation at that gate. (Color figure online)

algorithms for training computational policies that control the depth of Siamese convolutional networks [9] for video tracking. Siamese convolutional networks track by computing cross-correlation maps of deep features, between a *key frame*, where the object is labeled, and a *search frame*, where the object needs to be localized, as shown in Fig. 2. The peak of the cross-correlation map denotes the presence of the target object. We observe that in many search frames, the tracked object is similar to the object in the key frame and distinct from its background, and it would be computationally wasteful to compute elaborate features for its detection. To address this, we propose *conditional* computation controlled by gating functions that dynamically determines how many layers of our convolutional feature extractor should be computed before computing the cross-correlation map and thus the target’s location. In Fig. 1, our model uses only the first block of convolutional layers to find the motorcyclist against road, but uses 4 convolutional blocks to find the correct drummer among visually-similar peers. Our gate controllers are trained in an end-to-end differentiable framework without the need for sample-intense reinforcement learning. We test our model on the challenging VOT2016 dataset [11] and demonstrate that we perform video tracking with close to state-of-the-art accuracy at real-time speeds.

## 2 Related Works

### 2.1 Metric Learning for Visual Tracking

Metric learning approaches for visual tracking learn an appearance distance function between image box crops, so that the distance is large between image crops depicting different objects, and the distance is small between image crops depicting deformations of the same object instance. An accurate distance function then can be used to localize an object by computing the distances between the key frame and various crops within the search frame. Such a distance function can be learned using (a) Siamese networks [9], which use the same neural network

weights to extract features from a pair of images before using a single fully connected layer to predict the distance, they are trained using contrastive loss function that minimizes distance between same instance examples and requires distances to be above a certain margin for dissimilar examples. (b) Triplet networks, [8] trained with a ranking loss that ensures distance of positive pairs is lower than the distance of negative pairs, and obviates the need of a margin hyper-parameter.

## 2.2 Conditional Computation

Conditional computation refers to activating different network components depending on the input and serves as a promising way to reduce computational cost without sacrificing representational power. In [2], conditional computation is implemented by selectively activating different weights in each layer and is trained via reinforcement learning. [17] uses a sparse gating function to determine which sub-networks to execute (each of which are “experts” for different inputs), and shows that it is possible to train the gating and network weights jointly via back-propagation. Graves [5] proposed an adaptive computation model for Recurrent Neural Networks (RNNs), that determines (depending on the input) the number of computational (pondering) steps required before producing an accurate output. Recent work [4], adapts this model to convolutional networks for object detection in static images, where the network is trained to learn the number of convolutional layers to be evaluated per image location, e.g., “easy” image regions (e.g., sky) should require less computation than more feature-rich ones (e.g., a car). Our work differs in that (a) we use conditional computation in videos, rather than static images, and (b) the input we predicate computational decisions on is the quality of a cross-correlation tracking map, as opposed to image classification accuracy.


## 2.3 Estimating or Back-Propagating Gradients

A central question in all works that learn adaptive computation policies is how to train discrete *gates/controllers*, the discrete elements that determine how computation should be scheduled. Researchers have typically used non-differentiable score function estimators (a.k.a. REINFORCE [23]) for estimating the gradient with respect to such binary thresholds. REINFORCE has been shown to yield gradients with very high variance and requires too many samples for informative gradients to be estimated. High sample complexity is an attribute of many other model-free RL methods, e.g., Q-learning [15]. Indeed, recent works that use such RL techniques for neural net architectural search [25] or conditional computation [13], only scale to small networks and datasets [13], or use large computational resources for training, e.g., in recent work 800 GPUs were used concurrently [25], as opposed to a single GPU in our case. Alternatively, researchers have used soft, differentiable gates with carefully designed differentiable architectures for image generation (DRAW [7]), accessing an external memory (Neural Turing Machines [6]), deciding halting of a recurrent networks [5], etc. Our training scheme, which

similarly uses soft and differentiable gates during training to provide meaningful gradients, allows us to scale our policies to controlling deep neural architectures.

### 3 Depth-Adaptive Fully-Convolutional Siamese Networks

Our model builds upon fully-convolutional Siamese networks from [3], a state-of-the-art model for visual object tracking, which uses the same convolutional neural network to extract deep features from the key and search frames. The model then uses 2D cross-correlation to efficiently calculate the similarity score of the object in the key frame to every spatial location in the search frame, as shown in Fig. 2. This implicitly implements a triplet network-like loss by penalizing all the negative locations and increasing the similarity at the true location. Unlike [3], which uses an AlexNet-like [12] architecture and trains from scratch using ImageNet Video dataset [16], we use a VGG feature extractor [18] pretrained from the ImageNet static image classification dataset.



**Fig. 2.** Siamese network with 2D cross-correlation for key-search frame pairs. The deep feature maps for the key and search frames are extracted by the same convolutional neural network. \* denotes 2D cross-correlation.

We extend the fully-convolutional Siamese network for depth-adaptive computation by first dividing the convolutional layers into 5 “blocks” of convolutions and adding intermediate cross-correlations after every convolutional block. To finetune the convolutional weights, we calculate the softmax cross-entropy loss  $\mathcal{L}_i$  between each of the computed cross-correlation maps  $c_i$  and a ground-truth map  $G$  for  $i = 1, \dots, 5$  in our tracking training set. The ground-truth map is a 2D Gaussian centered at the true location of the object in the search frame.

$$\mathcal{L}_i = \text{softmax-cross-entropy}(c_i, G) \quad (1)$$

Furthermore, we introduce parametric gating functions between each of the convolutional blocks, which act as controllers for the depth of the VGG feature extractor at runtime. These gating functions take as input the cross-correlation map computed using the deep features at the current depth, and output a *confidence score for halting computation* at that particular depth. In theory, we

could use a convolutional neural network to extract the confidence score from each cross-correlation map. However, we would like our gating functions to be computationally inexpensive, so instead we use a small set of intuitive features to capture the “quality” (certainty) of each cross-correlation map, such as, kurtosis (measures “peakiness”), entropy, top-5 max peak values, and the first 5 moments. Let  $f$  denote the shallow feature extractor that given a cross-correlation map outputs the features above. We then learn a linear predictor, parameterized by  $\phi_i$  to output the confidence score  $\mathbf{g}_i$  for the gate at each depth, re-scaled to  $(0, 1)$  via a sigmoid function, as follows:

$$\mathbf{g}_i(c_i; \phi_i) = \text{sigm}(f(c_i)^T \cdot \phi_i) \in (0, 1). \quad (2)$$

Our full depth-adaptive model is depicted in Fig. 3. At training time, we use soft gates in order to use back-propagation for learning the gate weights, and at test time, we use hard gate thresholding, to halt computation at a particular network depth.

To train the model effectively and achieve a satisfactory trade-off of tracking accuracy and computational savings, we found the following two design choices to be crucial:

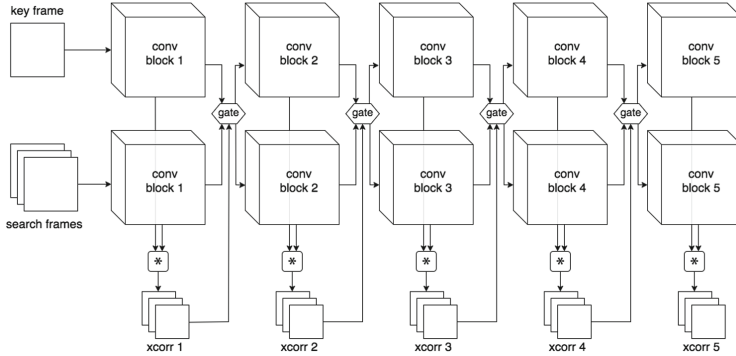
1. **Intermediate supervision:** Rather than training using the loss at the deepest layer only, like [3], we use a sum of tracking losses at all layers. This introduces intermediate supervision, which has been found to be useful in non-adaptive computational architectures such as [19, 24].
2. **Budgeted gating:** We found that directly using the confidence scores  $\mathbf{g}_i$  is insufficient for learning a depth-adaptive policy since each score does not affect the scores at other depths, which leads to polarized policies (either always use the shallowest depth or always use the deepest depth). Instead, we use a “budgeted” confidence score  $\mathbf{g}_i^*$ , in Eq. 3, where the scores sum to 1.0 and we have the desired behavior that a higher confidence score in a shallower depth corresponds to less need for deeper layers and vice-versa.

$$\mathbf{g}_i^*(c_i; \phi_i) = \begin{cases} (1 - \sum_{j=1}^{i-1} \mathbf{g}_j^*(c_j; \phi_j)) \mathbf{g}_i(c_i; \phi_i), & i \in [1, 4] \\ 1 - \sum_{j=1}^4 \mathbf{g}_j^*(c_j; \phi_j), & i = 5. \end{cases} \quad (3)$$

We cannot train the gate parameters  $\phi_i, i = 1 \dots 5$  and VGG weights jointly, since the gate feature extractor  $f$  is non-differentiable. Thus, we train in two phases. In the first phase, we finetune the VGG weights by minimizing the non-gated loss at all depths  $\mathcal{L}^{\text{conv}}$ :

$$\mathcal{L}^{\text{conv}} = \sum_{i=1}^5 \mathcal{L}_i \quad (4)$$

In the second phase, we fix the VGG weights and train the gate parameters by minimizing a loss function  $\mathcal{L}^{\text{gate}}$  that combines tracking loss and computational cost in all depths:



**Fig. 3.** Depth adaptive Siamese convolutional networks. Convolutional weights are shared between the two network stacks. Each conv block includes 2–4 convolutions with ReLU activation. The feature maps of the key and search frames at the end of each block are cross-correlated to yield 5 cross-correlation ( $\text{xcorr}$ ) maps. A gating function is added at the end of each convolutional block that controls whether the network stops at this layer, or continues computation to a higher depth.

$$\mathcal{L}^{\text{gate}} = \underbrace{\sum_{i=1}^5 \mathbf{g}_i^* \mathcal{L}_i}_{\text{tracking loss}} + \lambda \underbrace{\sum_{i=1}^5 p_i \mathbf{g}_i^*}_{\text{computational cost}}, \quad (5)$$

where the hyper-parameter  $\lambda$  trades off tracking accuracy and computational efficiency. We found that  $\lambda \in [0.5, 1.0]$  resulted in a diverse set of depths of computation (greater or less than that range generally led to “polarized” results, either all deepest layer or all shallowest layer). The parameter  $p_i$  encodes the relative incremental computational cost of each successive layer. For our experiments, we set each  $p_i$  to the incremental additional cost as reported in Table 1 with the  $p_1 = 1.0$ . For example,  $p_2 = 2.43 - p_1 = 1.43$ .

Much like [17], our training method provides balanced updates to all gates, meaningful gradients, and requires less training data. Though soft gates are used during training to enable back-propagation, at runtime, we threshold the budgeted confidence score and halt computation at a gate if the score exceeds some tune-able value. The score is a value in  $[0, 1]$  and in our experiments we set the threshold at 0.25, 0.5, or 0.75 for increasing degrees of strictness (i.e. higher threshold means we are less likely to accept the tracking result at a shallower layer).

### 3.1 Implementation Details

The base architecture we use for the convolutional layers is the 19-layer VGG architecture [18]. We remove the fully connected layers of the architecture and treat the remaining convolutional and max-pool layers as the feature extractor.

The VGG architecture is divided into 5 blocks of convolutions with 2–4 convolutional layers each, each ending in a max-pool layer. We remove the last max-pool layer in order to keep the deep feature maps as large as possible (i.e.  $16 \times 16$  in the last layer). To improve training, we normalize the key and search feature maps via batch normalization and rescale the output cross-correlation map to  $[0, 1]$ . Cross-correlation is an expensive operation so to keep the computational costs low, we downsample the feature maps to  $16 \times 16$  before cross-correlation.

Since training is performed with a single key frame and a batch of search frames, cross-correlation can be efficiently implemented on GPU by performing 2D convolution on the search feature maps with the key feature maps as the filter, treating the feature channels as the input channel size (the output channel size is 1).

Our model is implemented in TensorFlow v1.0.0 [1] using pretrained VGG network weights on ImageNet [16] for image classification. All training and evaluation was performed on a single NVIDIA TITAN X GPU, an Intel Xeon E5-2630 v3 CPU, and 16 GB of RAM.

To efficiently implement hard-gating, we use TensorFlow’s control flow operators (`tf.cond`). Hard-gating is only fully efficient when the batch size of the search frames is 1 since the computation is bottlenecked by the deepest cross-correlation map that is required by a sample in a batch. In practice, this is not as much of an issue since consecutive frames tend to use similar depths for prediction.

## 4 Experiments

We train and test our model on the Visual Object Tracking dataset VOT2016 [11]. The dataset consists of 60 videos with a total of 21455 frames of various resolutions. Each frame is labelled with the box corner coordinates of a bounding box that corresponds to a single object being tracked in the video. The videos have noisy backgrounds, the object can change shape or orientation, and there is occlusion in some frames. Since the VOT2016 dataset does not include a train-validation split, we randomly pick 25% of the videos to hold out as the test set. Note that the VOT dataset is designed for an evaluation-only competition so our results are not directly comparable to existing benchmarks. Our goal is not necessarily to beat state-of-the-art methods, but rather to present a useful technique for fast video tracking which can improve nearly any convolutional model.

We preprocess the videos by selecting a key frame every 10 frames and the subsequent up-to-100 frames as the search frames. We resize and crop the key frames to  $128 \times 128$  centered at the tracked object such that there is at least 25% padding around the bounding box. Each of the search frames are resized with the same scale and cropped to  $256 \times 256$  such that the frame is centered at the object at the previous frame. If the cropped search frame extends beyond the edge of the image, we pad the extra pixels with the mean RGB value of the dataset. The predicted object box is found using the position of the maximum

value in the cross-correlation map as the offset and the bounding box dimensions are the same as the key frame reference box.

#### 4.1 Evaluation Metrics

We measure tracking accuracy using Intersection-Over-Union (IOU) between the predicted object box  $b^{pred}$  and the ground truth box  $b^{gt}$ :

$$\text{IOU}(b^{pred}, b^{gt}) = \frac{|b^{pred} \cap b^{gt}|}{|b^{pred} \cup b^{gt}|}. \quad (6)$$

We measure IOU at up-to 1, 5, and 25 frames ahead of the key frame, e.g., for IOU@25, we take the average IOU of the tracker with key frame  $t$  and search frames  $t + 1, t + 2, \dots, t + 25$ . The larger the frame gap between key frame and search frame, the more the tracking target deforms and the harder it is to track.

We measure computational cost by computing the number of floating point operations (FLOPs) required to perform tracking on a batch of 25 search frames. Experimentally, we find that FLOPs is a good proxy for true computational cost as measured in frames-per-second (FPS). The reason FLOPs is the preferable metric is that FPS is heavily tied to hardware and software constraints, which may prevent the architecture from achieving the theoretical speedup.

#### 4.2 Siamese Tracker Performance

We finetune our VGG feature extractor starting from pretrained weights using the tracking loss of Eq. 4. The tracking performance during finetuning is shown in Fig. 4. Using pretrained weights allows the model to reach peak performance after only a few epochs. For this evaluation we use the full network depth.

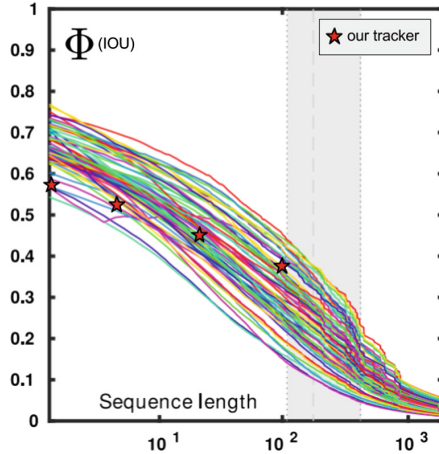


**Fig. 4.** Finetuning. Training and testing IOU curves during metric learning for different frame gaps. Starting with weights pretrained on ImageNet image classification, our VGG feature extractor fast reaches top performance.

Our implementation is comparable to the top submissions to the VOT2016 competition [10] in IOU over sequence length, as seen in Fig. 5. Though it is not



as accurate as the best trackers at small sequence lengths, it is competitive with many trackers at around 100 frames. Note also that the top four trackers from VOT2016 run at under 1 FPS while our system runs at over 54 FPS using the shallowest layer (`xcorr1`) and over 37 FPS using the deepest layer (`xcorr5`). Note that our code was not optimized so FPS is not a very good metric for computational cost. Additional software engineering (e.g. multithreaded inputs, better GPU utilization) should bring the `xcorr1` FPS to well over 100.



**Fig. 5.** VOT2016. Comparison of our full-computation model against top submissions to VOT2016 [10]. The stars represent our tracker’s accuracy at selected sequence lengths.

### 4.3 Effect of Intermediate Supervision

We compare the performance of the tracker when trained with loss only at the deepest layer against the tracker when trained with losses added in all depths. The IOU@25 comparison is shown in Fig. 6.


Training with all intermediate losses yields increased accuracy as depth increases, while training with only the deepest yields a big jump in accuracy at depth 5 but lower accuracy at shallower depths. If computational cost is not a factor, using only the deepest loss gives around 0.01 IOU benefit over using all depths. Since our goal is to use depth-adaptive feature extraction, intermediate supervision is essential.

### 4.4 Depth-Adaptive Computational Policies

The computational policies we compare are:

- Fixed-depth: always use the cross-correlation map at a fixed depth (i.e. `xcorr1`, `...`, `xcorr5`).

- Soft-gating: use a sum of the cross-correlation maps at all the depths, weighted by the budgeted confidence score. This policy does not save computational cost but serves as a baseline that also utilizes the gating functions.
- Hard-gating: we halt computation if the budgeted confidence score exceeds a tune-able threshold, which is a model hyper-parameter. We report performance while varying hyper-parameters in order to obtain accuracy/computation trade offs across the whole spectrum.




**Fig. 6.** Intermediate supervision. Supervision at intermediate layers (as opposed to the top layer only) increases the accuracy of the intermediate cross-correlation maps, and makes depth-adaptation at runtime worthwhile.

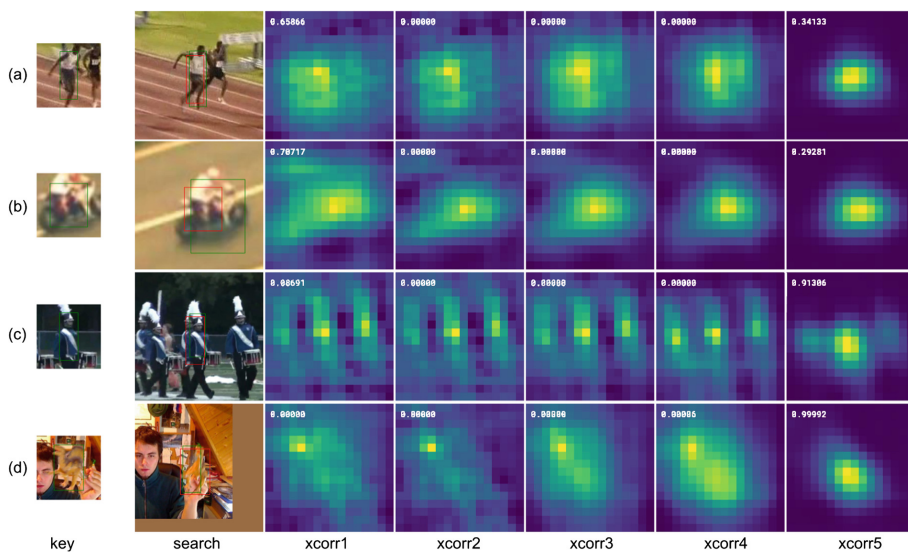
**Table 1.** Theoretical FLOPs for varying network depth.  $\text{xcorr}_i$  denotes network evaluation up until the  $i$ -th cross-correlation map. Soft-gating uses all five cross-correlation maps weighted by the gating confidence. The gating feature computation is negligible in comparison to convolutional feature extraction.

Gating policy	FLOPs ( $\times 10^9$ )	Relative to $\text{xcorr}_1$
$\text{xcorr}_1$	2.78	1.00 $\times$
$\text{xcorr}_2$	67.70	2.43 $\times$
$\text{xcorr}_3$	160.75	5.78 $\times$
$\text{xcorr}_4$	253.79	9.12 $\times$
$\text{xcorr}_5$	280.37	10.07 $\times$
Soft-gating	280.53	10.08 $\times$
Hard-gating	Varies	Varies

Table 1 shows the theoretical FLOPs required to compute the cross-correlation maps for a single key-search batch of 1 key frame and 25 search frames. For simplicity of calculation, the values only include the floating point multiplication operations, which comprise the overwhelming majority of the computation. As mentioned earlier, FLOPs is a better metric than FPS because it is independent of implementation details of the algorithm.



**Fig. 7.** Accuracy versus computation curves for our model and baselines. We generated the curve for our model (*hard gating*) by varying the relative weight  $\lambda$  of computational cost and tracking accuracy. For our fixed depth baseline, we obtain five points by varying the number of convolutional blocks from 1 to 5. Top left of the diagram is more desirable. Our model clearly outperforms the non-learned fixed-depth policies.



**Fig. 8.** Tracking results. Green box is ground truth, red box is prediction, red numbers are confidence weights. In (a), the tracker learns that *xcorr1* is sufficient for tracking. In (b), the tracker learns that it needs to compute *xcorr5* in order to confidently track the object. (Color figure online)

Figure 7 compares the accuracy and computational cost of each of the policies. Both soft and hard-gating can achieve accuracy values that exceed any fixed-depth policy and furthermore hard-gating uses significantly less computational cost to achieve the same or better accuracy. By varying hyper-parameter  $\lambda$ , we achieve different trade-offs of accuracy and computational cost depending on the requirements of the task.

The cross-correlation maps for selected video frames can be viewed in Fig. 8.

## 5 Discussion

Our experimental results show that our proposed depth-adaptive fully convolutional siamese network successfully tracks at accuracies comparable to state-of-the-art submissions to VOT2016. We demonstrate that our learned depth-adaptive policies can outperform fixed-depth networks while using significantly less computational power. Furthermore, we show that we can easily trade accuracy for computational cost as necessary by changing the model hyper-parameters.

Our work is limited by the following factors:

- Our model is finetuned purely on VOT2016 data unlike most other submissions which use ImageNet VID, or other tracking datasets for training the Siamese network. VOT2016 dataset is considerably smaller than ImageNet VID and the videos are considered more difficult to track.
- Fully-convolutional Siamese networks do not support updating the tracking model (i.e. key frame object appearance) since it does not naively support bounding box rescaling.
- Our final cross-correlation map resolution is  $16 \times 16$ , which is too coarse for fine-grained tracking. We experimented with higher resolution cross-correlation maps and obtained 15% improvement to IOU. However, this method is incompatible with hard-gating since it requires cross-correlation at multiple depths.

## 6 Conclusion

We have presented a conditional computation model for visual tracking, where computational depth is allocated based on a frame’s tracking difficulty. Our model is comprised of continuous weight filter variables and discrete learned controllers that dynamically manage the depth of the network at runtime, and balance accuracy with computational cost. The proposed model saves computation on “easy” frames without sacrificing representational power on difficult frames that require deeper features to track. We show our model outperforms naive non-adaptive policies by a significant margin, as measured by accuracy at a various computational costs. Paths for future work include multi-scale tracking

and more complex gating features, potentially using the history of the cross-correlation maps to determine computation. Though this work investigates policies that control the depth of the network, other promising “actions” for conditional computation in visual tracking include adaptive spatial computation e.g., by using motion to focus attention to the moving parts of the scene, or frame skipping e.g., by using only a subset of frames to localize the target without sacrificing tracking accuracy.

The methods presented in this work can be extended to any task which uses deep neural networks. By treating neural networks as a series of composable feature extractors, we have the ability to select feature embeddings at various degrees of complexity, which can both reduce computational cost and potentially improve performance.

## References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: large-scale machine learning on heterogeneous systems. Software: [tensorflow.org](https://www.tensorflow.org) (2015)
2. Bengio, E., Bacon, P., Pineau, J., Precup, D.: Conditional computation in neural networks for faster models. CoRR, abs/1511.06297 (2015)
3. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.S.: Fully-convolutional siamese networks for object tracking. In: Hua, G., Jégou, H. (eds.) ECCV 2016. LNCS, vol. 9914, pp. 850–865. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-48881-3\\_56](https://doi.org/10.1007/978-3-319-48881-3_56)
4. Figurnov, M., Collins, M.D., Zhu, Y., Zhang, L., Huang, J., Vetrov, D.P., Salakhutdinov, R.: Spatially adaptive computation time for residual networks. In: CVPR (2017)
5. Graves, A.: Adaptive computation time for recurrent neural networks. CoRR, abs/1603.08983 (2016)
6. Graves, A., Wayne, G., Danihelka, I.: Neural Turing machines. CoRR, abs/1410.5401 (2014)
7. Gregor, K., Danihelka, I., Graves, A., Rezende, D.J., Wierstra, D.: DRAW: a recurrent neural network for image generation. In: ICML, pp. 1462–1471 (2015)
8. Hoffer, E., Ailon, N.: Deep metric learning using triplet network. CoRR, abs/1412.6622 (2014)
9. Koch, G.: Siamese neural networks for one-shot image recognition. Ph.D. thesis, University of Toronto (2015)
10. Kristan, M., et al.: The visual object tracking VOT2016 challenge results. In: Hua, G., Jégou, H. (eds.) ECCV 2016. LNCS, vol. 9914, pp. 777–823. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-48881-3\\_54](https://doi.org/10.1007/978-3-319-48881-3_54)
11. Kristan, M., Matas, J., Leonardis, A., Felsberg, M., Cehovin, L., Fernandez, G., Vojir, T., Hager, G., Nebehay, G., Pflugfelder, R.: The visual object tracking VOT2015 challenge results. In: ICCV, pp. 1–23 (2015)

12. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*, pp. 1097–1105 (2012)
13. Liu, L., Deng, J.: Dynamic deep neural networks: optimizing accuracy-efficiency trade-offs by selective execution. [arXiv:1701.00299](https://arxiv.org/abs/1701.00299) (2017)
14. Ma, C., Huang, J.-B., Yang, X., Yang, M.-H.: Hierarchical convolutional features for visual tracking. In: *ICCV*, pp. 3074–3082 (2015)
15. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.A.: Playing atari with deep reinforcement learning. [arXiv:1312.5602](https://arxiv.org/abs/1312.5602) (2013)
16. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. *IJCV* **115**(3), 211–252 (2015)
17. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q.V., Hinton, G.E., Dean, J.: Outrageously large neural networks: the sparsely-gated mixture-of-experts layer. *CoRR*, abs/1701.06538 (2017)
18. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556 (2014)
19. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: *CVPR* (2015)
20. Wang, L., Ouyang, W., Wang, X., Lu, H.: Visual tracking with fully convolutional networks. In: *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 3119–3127, December 2015
21. Wang, N., Yeung, D.-Y.: Learning a deep compact image representation for visual tracking. In: Burges, C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K. (eds.), *Advances in Neural Information Processing Systems*, vol. 26, pp. 809–817 (2013)
22. Weng, S.-K., Kuo, C.-M., Tu, S.-K.: Video object tracking using adaptive kalman filter. *J. Vis. Commun. Image Represent.* **17**(6), 1190–1208 (2006)
23. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.* **8**(3), 229–256 (1992)
24. Xie, S., Tu, Z.: Holistically-nested edge detection. *CoRR*, abs/1504.06375 (2015)
25. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. [arXiv:1611.01578](https://arxiv.org/abs/1611.01578) (2016)