

Tracking as Online Decision-Making: Learning a Policy from Streaming Videos with Reinforcement Learning

James Supančić, III
University of California, Irvine
jsupanci@uci.edu

Deva Ramanan
Carnegie Mellon University
deva@cs.cmu.edu

Abstract

We formulate tracking as an online decision-making process, where a tracking agent must follow an object despite ambiguous image frames and a limited computational budget. Crucially, the agent must decide where to look in the upcoming frames, when to reinitialize because it believes the target has been lost, and when to update its appearance model for the tracked object. Such decisions are typically made heuristically. Instead, we propose to learn an optimal decision-making policy by formulating tracking as a partially observable decision-making process (POMDP). We learn policies with deep reinforcement learning algorithms that need supervision (a reward signal) only when the track has gone awry. We demonstrate that sparse rewards allow us to quickly train on massive datasets, several orders of magnitude more than past work. Interestingly, by treating the data source of Internet videos as unlimited streams, we both learn and evaluate our trackers in a single, unified computational stream.

1. Introduction

Object tracking is one of the basic computational building blocks of video analysis, relevant for tasks such as general scene understanding and perception-for-robotics. A particularly popular formalism is that of model-free tracking, where a tracker is provided with a bounding-box initialization of an unknown object. Much of the recent state-of-the-art advances make heavy use of machine learning [23, 54, 40, 56, 20], often producing impressive results by improving core components such as appearance descriptors or motion modeling.

Challenges: We see two significant challenges that limit further progress. First, the *limited quantity of annotated video data* impedes both training and evaluation. While image datasets involve millions of images for training and testing, tracking datasets have hundreds of videos. Lack of data seems to arise from the difficulty of annotating

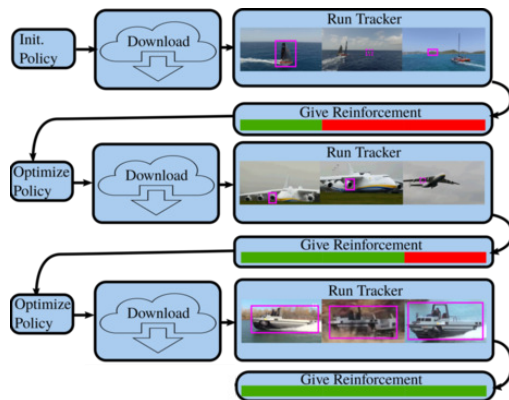


Figure 1. **Streaming interactive training:** We propose an iterative procedure for interactively training trackers from data. We *download* a new video from the Internet and run the current tracker on it, *evaluate* the tracker’s performance with interactive rewards, and then *retrain* the tracker policy (with reinforcement learning) with the reward signals. Importantly, rather than requiring interactive labeling of bounding-boxes, we require only binary (incorrect / correct) feedback from human users. This scheme allows us to train and evaluate our tracker on massive streaming datasets, 100X larger than prior work (Table 1).

videos, as opposed to images. Second, as vision (re-)integrates with robotics, video processing must be done in an online, streaming fashion. This requires a tracker to make *on-the-fly decisions* such as when to re-initialize itself [51, 16, 30, 47, 21] or update its appearance model (the so-called template-update problem [54, 21]). Such decisions are known to be crucial in terms of final performance, but are typically hand-designed rather than learned.

Contribution 1 (interactive video processing): We show that reinforcement learning (RL) can be used to address both challenges in distinct ways. In terms of data, rather than requiring videos to be labeled with detailed bounding-boxes at each frame, we interactively train trackers with far more limited supervision (specifying binary rewards/penalties only when a tracker fails). This allows us to train on massive video datasets that are 100× larger than

prior work. Interestingly, RL also naturally lends itself to streaming “open-world” evaluation: when running a tracker on a never-before-seen video, the video can be used for both evaluation of the current tracker and for training (or refining) the tracker for future use (Fig. 1). This streaming evaluation allows us to train and evaluate models in an integrated fashion seamlessly. For completeness, we also evaluate our learned models on standard tracking benchmarks.

Contribution 2 (tracking as decision-making): In terms of tracking, we model the tracker itself as an *active agent* that must make online decisions to maximize its reward, which is (as above) the correctness of a track. Decisions ultimately specify where to devote finite computational resources at any point of time: should the agent process only a limited region around the currently predicted location (e.g., “track”), or should it globally search over the entire frame (“reinitialize”)? Should the agent use the predicted image region to update its appearance model for the object being tracked (“update”), or should it be “ignored”? Such decisions are notoriously complicated when image evidence is ambiguous (due to say, partial occlusions): the agent may continue tracking an object but perhaps decide not to update its model of the object’s appearance. Rather than defining these decisions heuristically, we will ultimately use data-driven techniques to learn good *policies* for active decision-making (Fig. 2).

Contribution 3 (deep POMDPs): We learn tracker decision policies using reinforcement learning. Much recent work in this space assumes a Markov Decision Process (MDP), where the agent observes the true state of the world [34, 56], which is the true (possibly 3D) location and unoccluded appearance of the object being tracked. In contrast, our tracker only assumes that it receives *partial* image observations about the world state. The resulting partially-observable MDP (POMDP) violates Markov independence assumptions: actions depend on the entire history of observations rather than just the current one [19, 43]. As in [15, 22], we account for this partial observability by maintaining a *memory* that captures *beliefs* about the world, which we update over time (Sec. 3). In our case, beliefs capture object location and appearance, and action policies specify how and when to update those beliefs (e.g., how and when should the tracker update its appearance model) (Sec. 4). However, policy-learning is notoriously challenging because actions can have long-term effects on future beliefs. To efficiently learn policies, we introduce frame-based heuristics that provide strong clues as to the long-term effects of taking a particular action (Sec. 5).

2. Related Work

Tracking datasets: Several established benchmarks exist for evaluating trackers [55, 23]. Interestingly, there is evidence to suggest that many methods tend to overfit due

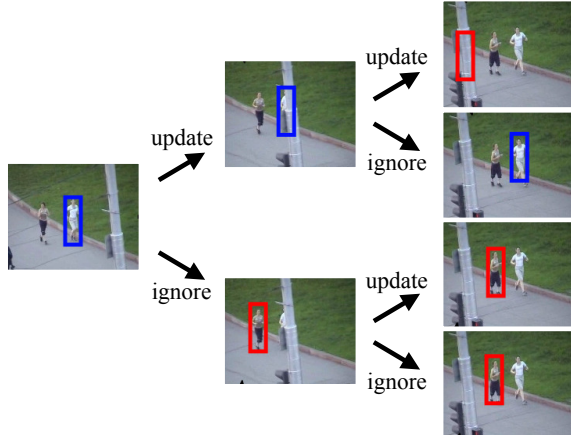


Figure 2. **Decisions in tracking:** Trackers must decide when to update their appearance and when to re-initialize. This example enumerates the four possible outcomes of updating appearance (or not) over two frames, where blue denotes a good track and red denotes an error. Given a good track (left), it is important to update appearance to track through challenging frames with occlusions (center), but equally important to not update after an occlusion to prevent drift (right). Though such decisions are typically made heuristically, we recast tracking as a sequential decision-making process, and learn a policy with reinforcement learning.

to aggressive tuning [40]. Withholding test data annotation and providing an evaluation server addresses this to some extent [26, 42]. Alternatively, we propose to evaluate on an open-world stream of Internet videos, making over-fitting impossible by design. It is well-known that algorithms trained on “closed-world” datasets (say, with centered objects against clean backgrounds [39, 4]) are difficult to generalize to “in-the-wild” footage [48]. We invite the reader to compare our videos in the supplementary material to contemporary video benchmarks for tracking.

Interactive tracking: Several works have explored interactive methods that use trackers to help annotate data. The computer first proposes a track. Then a human corrects major errors and retrains the tracker using the corrections [8, 2, 50]. Our approach is closely inspired by such active learning formalisms but differs in that we make use of minimal supervision in the form of a binary reward (rather than a bounding box annotation).

Learning-to-track: Many tracking benchmarks tend to focus on short-term tracking (< 2000 frames per video) [55, 23]. In this setting, a central issue appears to be modeling the appearance of the target. Methods that use deep features learned from large-scale training data perform particularly well [51, 52, 29, 27, 36]. *Our focus on tracking over longer time frames poses additional challenges - namely, how to reinitialize after cuts, occlusions and failures, despite changes in target appearance [14, 54].* Several trackers address these challenges with hand-designed policies for model updating and reini-

Dataset	# Videos	# Frames	Annotations	Type
OTB-2013 [55]	50	29,134	29,134	AABB
PTB [46]	50	21,551	21,551	AABB
VOT-2016 [23]	60	21,455	21,455	RBB
ALOV++ [44]	315	151,657	30,331	AABB
NUS-PRO [26]	365	135,310	135,310	AABB
<i>Ours</i>	16,384	10,895,762	108,957	Binary

Table 1. **Our interactive** learning formulation allows us to train and evaluate on dramatically more videos than prior work. We annotate binary rewards, while the other datasets provide Axis Aligned (AABB) or Rotated (RBB) Bounding Boxes.

tialization - TLD [21], ALIAN [41] and SPL [47] explicitly do so in the context of long-term tracking. On the other hand, our method takes a data-driven approach and learns policies for model-updating and reinitialization. Interestingly, such an “end-to-end” learning philosophy is often embraced by the multi-object tracking community, where strategies for online reinitialization and data association are learned from data [24, 56, 28]. Most related to us are [56], who use an MDP for multi-object tracking, and [22], who use RL for single target tracking. Both works use heuristics to reduce policy learning to a supervised learning task, avoiding the need to reason about rewards in the far future. Our experimental results show that explicit Q-learning outperforms such heuristics because it can learn to capture long-term effects of taking particular actions.

Real-time tracking through attention: An interesting (but perhaps unsurprising) phenomenon is that better trackers tend to be slower [23]. Indeed, on the VOT benchmark, most recent trackers do not run in real time. Generally, trackers that search locally [45, 49] run faster than those that search globally [47, 30, 17]. To optimize recognition efficiency, one can learn a policy to guide selective search or attention. Inspired by recent work which finds a policy for selective search using RL [20, 33, 38, 10, 18, 3, 31], we also learn a policy that decides whether to track (*i.e.*, search positions near the previous estimate) or reinitialize (*i.e.*, search globally over the entire image). But in contrast to prior work, we additionally learn a policy to decide when to update a tracker’s appearance model. To ensure that our tracker operates with a fixed computational budget, we implement reinitialization by searching over a random subset of positions (equal in number to those examined by track).

3. POMDP Tracking

We now describe our POMDP tracker, using standard notation where possible [43]. For our purposes, a POMDP is defined by a tuple of states (Ω, O, A, B): At each frame i , the world state $\omega_i \in \Omega$ generates a noisy observation $o_i \in O$ that is mapped by an agent into an action $a_i \in A$, which in turn generates a reward.

In our case, the state ω_i captures the true location and

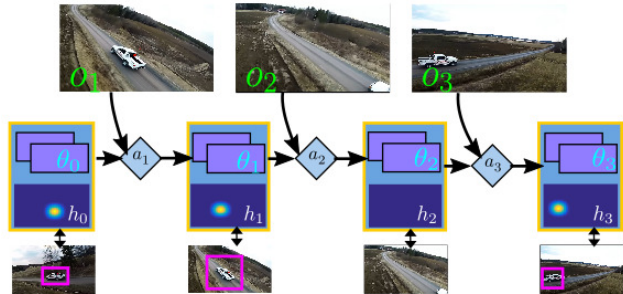


Figure 3. **Tracker architecture:** At each frame i , our tracker updates a location heatmap h_i for the target using the current image observation o_i , a location prior given by the previous frames’ heatmap h_{i-1} , and the previous appearance model θ_{i-1} . Crucially, our tracker learns a policy for actions a_i that optimally update h_i and θ_i (2).

appearance of the object being tracked in frame i . To help build intuition, one can think of the location as 2D pixel coordinates and appearance as a 2D visual template. Instead of directly observing this world state, the tracking agent maintains a belief over world states, written as

$$b_i = (\theta_i, h_i), \quad \text{where } \theta_i \in R^{h \times w \times f}, h_i \in R^{H \times W}$$

where θ_i is a distribution over appearances (we use a point-mass distribution encoded by a single $h \times w$ filter defined on f convolutional features), and h_i is a distribution over pixel positions (encoded as a spatial heatmap of size $H \times W$). Given the previous belief b_{i-1} and current observed video frame o_i , the tracking agent updates its beliefs about the current frame b_i . Crucially, *tracker actions* a_i specify how to update beliefs, that is, whether to update the appearance model and whether to reinitialize by disregarding previous heatmaps. From this perspective, our POMDP tracker is a memory-based agent that learns a *policy* for when and how to update its own memory (Fig. 3).

Specifically, beliefs are updated as follows:

$$h_i = \begin{cases} TRACK(h_{i-1}, \theta_{i-1}, o_i) & \text{if } a_i^{(1)} = 1 \\ REINIT(\theta_{i-1}, o_i) & \text{otherwise} \end{cases} \quad (1)$$

$$\theta_i = \begin{cases} UPDATE(\theta_{i-1}, h_i, o_i) & \text{if } a_i^{(2)} = 1 \\ \theta_{i-1} & \text{otherwise} \end{cases}$$

where $a_i = (a_i^{(1)}, a_i^{(2)})$. Object heatmaps h_i are updated by running the current appearance model θ_{i-1} on image regions o_i near the target location previously predicted using h_{i-1} (“tracking”). Alternatively, if the agent believes it has lost track, it may globally evaluate its appearance model (“reinit”). The appearance model is then “updated” with the currently-predicted image region or possibly unchanged. In our framework, the tracking, reinitialization, and appearance-update modules can be treated as black-

boxes given the above functional form. We will discuss particular implementations shortly (with copious details in the supplementary material), but first, we focus on the heart of our RL approach: a principled framework for learning to take appropriate actions. To do so, we begin by reviewing standard approaches for decision-making.

Online heuristics: The simplest approach to picking actions might be to pre-define heuristics functions which estimate the correct action. For example, many trackers reinitialize whenever the maximum confidence of the heatmap is below a threshold. Let us summarize the information available to the tracker at frame i as a “state” s_i , which includes the previous belief b_{i-1} (the previous heatmap and appearance model) and current image observation o_i .

$$a_i^* = \text{Heur}_{on}(s_i), \quad \text{where } s_i = (b_{i-1}, o_i) \quad (2)$$

Offline heuristics: A generalization of the above is to use *offline* training data to build better heuristics. Crucially, one can now make use of ground-truth training annotations as well as future frames to better gauge the impact of possible actions. We can write this knowledge as the true world state $\{\omega_i, \forall i\}$. For example, a natural heuristic may be to reinitialize whenever the predicted object location does not overlap the ground-truth position for that frame. Similarly, one may update appearance whenever doing so improves the confidence of ground-truth object locations across future frames:

$$a_i^* = \text{Heur}_{off}(s_i, \{\omega_i : \forall i\}) \quad (3)$$

Crucially, these heuristics cannot be applied at test time because ground truth is not known! However, they can generate per-frame target action labels a_i^* on training data, effectively reducing policy learning to a *supervised-learning problem*. Though offline heuristics appear to be a simple and intuitive approach to policy learning, we have not seen them widely used for learning tracker action policies.

Q-functions: Offline heuristics can be improved by unrolling them forward in time: the benefit of a putative action can be better modeled by applying that action, processing the next frame, and using the heuristic to score the “goodness” of possible actions in that next frame. This intuition is formalized through the well-known *Bellman equations* that recursively define Q-functions to return a goodness score (the expected future reward) for each putative action a :

$$Q(s_i, a_i) = R(s_i) + \gamma \max_{a_{i+1}} Q(s_{i+1}, a_{i+1}), \quad (4)$$

where s_i includes both the tracker belief state and image observation, and $R(s_i)$ is the reward associated with the reporting the estimated object heatmap h_i . We let $R(s_i) = 1$ for a correct prediction and 0 otherwise. Finally, $\gamma \in [0, 1]$

is a discount factor that trades off immediate vs future per-frame rewards. Given a tracker state and image observation s_i , the optimal action is readily computed from the Q-function:

$$a^* = \arg \max_a Q(s_i, a)$$

Q-learning: Traditionally, Q-functions are iteratively learned with Q-learning [43]:

$$Q(s_i, a_i) \leftarrow Q(s_i, a_i) + \dots \quad (5)$$

$$\alpha \left(R(s_i) + \gamma \max_{a_{i+1}} Q(s_{i+1}, a_{i+1}) - Q(s_i, a_i) \right)$$

where α is a learning rate. To handle continuous belief states, we approximate the Q-function with a CNNs:

$$Q(s_i, a_i) \approx \text{CNN}(s_i, a_i)$$

that processes states s_i and binary actions a_i to return a scalar value capturing the expected future reward for taking that action. Recall that a state s_i encodes a heatmap and an appearance model from previous frames and an image observation from the current frame.

4. Interactive Training and Evaluation

In this section, we describe our procedure for interactively learning CNN parameters w (that encode tracker action policies) from streaming video datasets. To do so, we gradually build up a database of *experience replay memories* [1, 34], which are a collection of state-action-reward-nextstate tuples $\mathcal{D} = \{(s_i, a_i, r_i, s_{i+1})\}$:

$$L(w, \mathcal{D}) = \left(r_i + \max_{a_{i+1}} Q_w(s_{i+1}, a_{i+1}) - Q_w(s_i, a_i) \right)^2 \quad (6)$$

Gradient descent on the above objective is performed as follows: given a training sample (s_i, a_i, r_i, s_{i+1}) , first perform a forward pass to compute the current estimate $Q_w(s_i, a_i)$ and the target: $r_i + \max_{a_{i+1}} Q_w(s_{i+1}, a_{i+1})$. Then backprop through the weights w to reduce $L(w)$.

The complete training algorithm is written in Alg. 1. We choose random videos from the Internet by sampling phrases using WordNet [32]. Given the sampled phrase and video, an annotator provides an initialization bounding box and begins running the existing tracker. After tracking, the annotator marks those frames (in strides of 50) where the tracker was incorrect using a standard 50% intersection-over-union threshold. Such binary annotation (“correct” or “failed”) requires far less time per frame than bounding-box annotation: we design a real-time interface that simply requires a user to depress a button during tracker failures. By playing back videos at a (user-selected) sped-up frame rate, users annotate 1200 frames per minute on average (versus 34 for bounding boxes). Annotating our entire dataset of 10

```

1 while True do
2   Download random video;
3    $\theta_1 \leftarrow \text{UPDATE}(h_1, o_1)$ ; /* manually init. */
4   forall  $i \in \text{video}$  do
5      $s_i = ((\theta_{i-1}, h_{i-1}), o_i)$ ;
6     /* track or reinitialize? */
7     if  $Q_w(s_i, \mathbf{t}) > Q_w(s_i, \mathbf{r})$  then
8        $h_i \leftarrow \text{TRACK}(h_{i-1}, \theta_{i-1}, o_i)$ ;  $a_i^{(1)} \leftarrow 1$ ;
9     else  $h_i \leftarrow \text{REINIT}(\theta_{i-1}, o_i)$ ;  $a_i^{(1)} \leftarrow 0$ ;
10    /* update or ignore? */
11    if  $Q_w(s_i, \mathbf{m}) > Q_w(s_i, \mathbf{i})$  then
12       $\theta_i \leftarrow \text{UPDATE}(\theta_{i-1}, h_i, o_i)$ ;  $a_i^{(2)} \leftarrow 1$ ;
13    else  $\theta_i \leftarrow \theta_{i-1}$ ;  $a_i^{(2)} \leftarrow 0$ ;
14    /* manually evaluate the performance */
15     $r_i \leftarrow \text{annotated frame correct?}$ ;
16    /* update experience database */
17     $D \leftarrow D \cup (s_i, a_i, r_i, s_{i+1})$ 
18   $w \leftarrow \arg \min L(w; D)$ 

```

Algorithm 1: Our final learning algorithm interactively labels a streaming dataset of videos while learning a tracker action policy Q_w . Given a video, steps 5 through 11 run a tracker according to the current policy. An annotator then assesses the binary reward r_i (correctness) for the highest-scoring bounding box extracted from the heatmap h_i by using an intersection-over-union threshold. Annotated frames (and their associated state-action-reward tuples) are added to our experience replay database D . We then sample a minibatch of replay memories and update the action policy w with backprop (Eq. 6).

million frames (Table 1) requires a little under two-days of labor, versus the months required for equivalent bounding-box annotation. After running our tracker and interactively marking failures, we use the annotation as a reward signal to update the policy parameters for the next video. Thus each video is used to both evaluate the current tracker and train it for future videos.

5. Implementation

TRACK/REINIT: Our TRACK function (Eq. 2) takes as input the previous heatmap h_{i-1} , appearance model θ_{i-1} , image observation o_i , and produces a new heatmap for the current frame. We make use of the state-of-the-art fully-convolutional tracker FCNT [51] and refer to the reader to that work for precise implementation details, but summarize them here: TRACK crops the current image o_i to a region of interest (ROI) around the most-likely object location from the previous frame (the argmax of h_{i-1}). This ROI is resized to a canonical image size (e.g., 224×224) and processed with a CNN (VGG16) to produce a convolutional feature map. The object appearance model θ_{i-1}

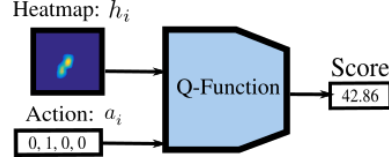


Figure 4. **Q-CNNs:** A Q-function predicts a score (the expected future reward) as a function of (1) the localization heatmap and (2) an action encoded using a one-hot encoding.

is represented as a filter on this feature map, allowing one to compute a new heatmap with a convolution. When the tracker believes it has lost track, the REINIT model simply processes a random ROI.

UPDATE: We update the current filter θ using positive and negative patches extracted from the current frame i . We extract a positive patch from the maximal location in the reported heatmap h_i , and extract negative patches from adjacent regions with less than 30% overlap. We update θ following the default scheme in the underlying tracker: for FCNT, θ is a two-layer convolutional template that is updated with a fixed number of gradient descent iterations (10). For CCOT, θ is a multi-resolution set of convolutional templates that is fit through conjugate gradients.

Q-function CNN: Recall that our Q-functions process a tracker state $s_i = ((h_{i-1}, \theta_{i-1}), o_i)$ and a candidate action a_i , to return a scalar representing the expected future reward of taking that action (Fig. 4 and Eq. 6). In practice, we define two separate Q-functions for our two binary decisions (TRACK/REINIT and UPDATE/IGNORE). To plug into standard learning approaches, we formally define a single Q function as the sum of the two functions, implying that the optimal decisions can be computed independently for each. We found it sufficed to condition on the heatmap h_{i-1} and implemented each function as a CNN, where the first two hidden layers are shared across the two functions. Each shared hidden layer consists of $16 \times 16 \times 4$ convolution followed by ReLU and 4×4 max pooling. For each decision, an independent fully-connected layer ultimately predicts the expected future reward. When training the Q-function using experience-replay, we use $\gamma = .95$, a learning rate of $1e-4$, a momentum of .9 and $1e-8$ weight decay.

Offline heuristics: Deep Q-learning is known to be unstable, and we found good initialization was important for reliable convergence. We initialize the Q-functions in Eq. 6 (which specify the goodness of particular actions) with the offline heuristics from Eq. 3. Specifically, the heuristic action a_i^* has a goodness of 1 (scaled by future discount rewards), while other actions have a goodness of 0:

$$Q_{\text{init}}(s_i, a_i) \Leftarrow \mathbf{I}[a_i = a_i^*] \sum_{j \geq i} \gamma^{j-i} \quad (7)$$

where \mathbf{I} denotes the identity function. In practice, we found it useful to minimize a weighted average of the true

Online Tracking Benchmark [55]

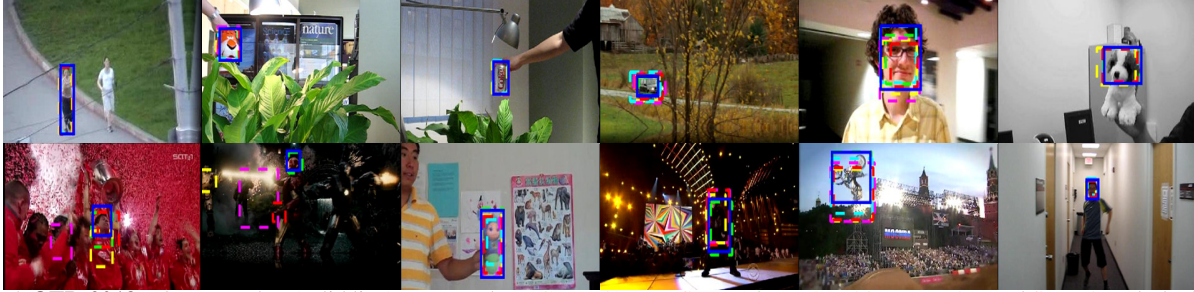


Figure 5. **OTB-2013**: Our p-tracker (solid line) compared to FCNT [51], MUSTer [16], LTCT [30], TLD [21] and SPL [47] (dashed lines). In general, many videos are easy for modern trackers, implying that a method’s rank is determined by a few challenging videos (such as the confetti celebration and fireworks on the bottom left). Our tracker learns an UPDATE and REINIT policy that does well on such videos.




Figure 6. **Internet videos** contain new challenges, such as cuts, strange and interesting behaviors, fast motion and complex illumination. We show select results for our p-tracker plus FCNT [51], MUSTer [16], LTCT [30], TLD [21] and SPL [47].

loss in Eq. 6 and a supervised loss $(Q_w - Q_{init})^2$, an approach related to heuristically-guided Q-learning [20, 7]. Defining a heuristic for TRACK vs REINIT is straightforward: a_i^* should TRACK whenever the peak of the reported heatmap overlaps the ground-truth object on frame i . Defining a heuristic for UPDATE vs IGNORE is more subtle. Intuitively, a_i^* should UPDATE appearance with frame i whenever doing so improves the confidence of future ground-truth object locations in that video. To operationalize this, we update the current appearance model θ_i on samples from frame i and compute Δ_+ , the number of future frames where the updated appearance increases confidence of ground-truth locations (and similarly Δ_- , the number of frames where the update decreases confidence of track errors). We set a_i^* to update when $\Delta_+ + \Delta_- > .5N$, where N is the total number of future frames.

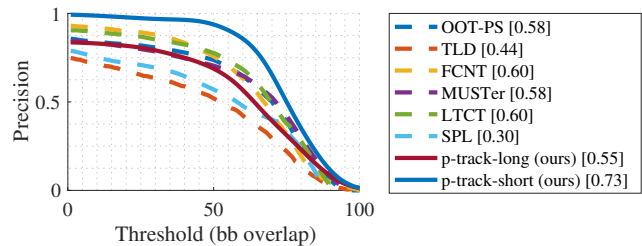


Figure 7. **OTB-2013 [55] results**: Our learned policy tracker (p-track) performs competitively on standard short-term tracking benchmarks. We find that a policy learned for long-term tracking (p-track-long) tends to select the IGNORE action more often (appropriate during occlusions, which tend to be more common in long videos). Learning a policy from short-term videos significantly improves performance, producing state-of-the-art results: compare our p-track-short vs OOT-PS [17], TLD [21], FCNT [51], MUSTer [16], LTCT [30], and SPL [47]

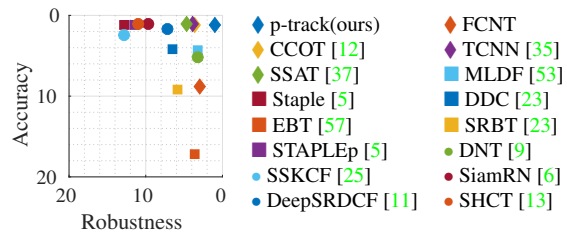


Figure 8. **VOT-2016 [23] results**: Our learned policy tracker (p-track-short) is as accurate as the state-of-the-art but is considerably more robust. Robustness is measured by a ranking of trackers according to the number of times they fail, while accuracy is the rank of a tracker according to its average overlap with the ground truth. Notably, p-track significantly outperforms FCNT [51] and CCOT [12] in terms of robustness, even though its TRACK and UPDATE modules follow directly from those works.

6. Experiments

Evaluation metrics: Following established protocols for long-term tracking [21, 47], we evaluate $F1 = \frac{2pr}{p+r}$, where precision (p) is the fraction of predicted locations that are correct and recall (r) is the fraction of ground-

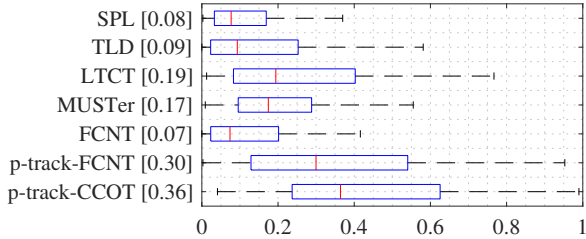


Figure 9. **Vs state-of-the-art:** Our learned policy (*p-track*) performs better than state-of-the-art baselines [51, 16, 30, 21, 47] on a held-out test set. See Sec. 6.1 for discussion.

truth locations that are correctly predicted. Because Internet videos vary widely in difficulty, we supplement averages with boxplots to better visualize performance skew. When evaluating results on standard benchmarks, we use the default evaluation criteria for that benchmark.

6.1. Comparative Evaluation

Short-term benchmarks: While our focus is long-term tracking, we begin by presenting results on existing benchmarks that tend to focus on the short-term setting – the Online Tracker Benchmark (OTB-2013) [55] and Visual Object Tracking Benchmark (VOT-2016) [23]. Our policy tracker (*p-track-long*), trained on Internet videos, performs competitively (Fig. 7), but tends to over-predict occlusions (which rarely occur in short-term videos). But fortunately, we can learn a dedicated policy for short-term tracking (*p-track-short*) by applying reinforcement learning (Alg. 1) on short-term training videos. For each test video in OTB-2013, we learn a policy using the 40 most dissimilar videos in VOT-2016 (and vice-versa). We define similarity between videos to be the correlation between the average (ground-truth) object image in RGB space. This ensures that, for example, we do not train using *Tiger1* when testing on *Tiger2*. Even under this controlled scenario, *p-track-short* significantly outperforms prior work on both OTB-2013 (Figs. 5 and 7) and VOT-2016 (Fig. 8).

Long-term baselines: For the long-term setting, we compare to two classic long-term trackers: TLD [21] and SPL [47]. Additionally, we also compare against short-term trackers with public code that we were able to adapt: FCNT [51], MUSTer [16], and LTCT [30]. Notably, all these baselines use hand-designed heuristics for deciding when to appearance update and reinitialize.

Long-term videos: Qualitatively speaking, long term videos from the internet are much more difficult than standard benchmarks (*c.f.* Fig 5 and Fig. 6). First, many standard benchmarks tend to contain videos that are easy for most modern tracking approaches, implying that a method’s rank is largely determined by performance on a small number of challenging videos. The easy videos focus on iconic [39, 4] views with slow motion and stable light-

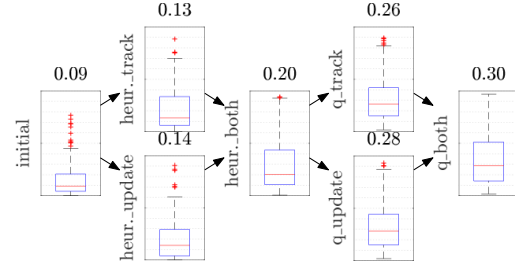


Figure 10. **System diagnostics:** Beginning with the *initial* policy of FCNT [51], we evolve towards our final data-driven policy objective. As shown, each component of our objective measurably improves performance on the hold-out test-data. See Sec. 6.2 for discussion.

ing conditions [26], featuring no cuts or long-term occlusions [14]. Internet videos are significantly more complex. One major reason is the presence of frequent cuts. We think that Internet videos with multiple cuts provide a valuable proxy for occlusions, particularly since they are a scalable data source. In theory, long-term trackers must be able to re-detect the target after an occlusion (or cut), but there is still much room for improvement. Also, many strange things happen in the wild world of Internet videos. For example, in *Transform* the car transforms into a robot, confusing all trackers (including ours). In *SnowTank* the tracker must contend with many distractors (tanks of different colors and type) and widely varying viewpoint and scale. Meanwhile, *JohnWick* contains poor illumination, fast motion, and numerous distractors.

Long-term results: To evaluate results for “in-the-wild” long-term tracking, we define a new 16-video held-out test set of long-term Internet videos that is *never used for training*. Each of our test videos contains at least 5,000 frames, a common definition of “long-term” [47, 30, 16, 14]. We compare our method to various baselines in Fig. 9. Comparisons to FCNT [51] and CCOT [12] are particularly interesting since we can make use of their TRACK and UPDATE modules. While FCNT performs quite well in the short-term (Fig. 7), it performs poorly on long-term sequences (Fig. 9). However, by learning a policy for updating and reinitialization, we produce a state-of-the-art long-term tracker. We visualize the learned policy in Fig. 12.

6.2. System Diagnostics

We now provide a diagnostic analysis of various components of our system. We begin by examining several alternative strategies for making sequential decisions (Fig. 10).

Online vs offline heuristics: We begin by analyzing the online heuristic actions of our baseline tracker, FCNT. FCNT updates an appearance model when the predicted heatmap location is above a threshold, and always tracks without reinitialization. This produces a F1 score of .09. Next, we use offline heuristics to learn the best action to

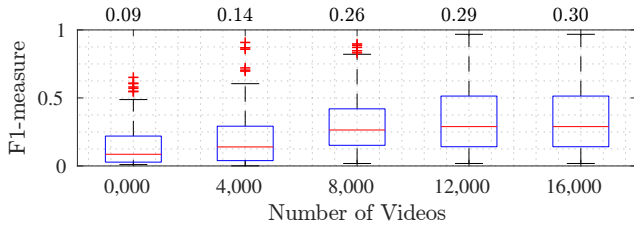


Figure 11. **Our p-tracker’s performance increases** as it learns its policy using additional Internet videos. Above, we plot the distribution of F1 scores on our hold-out test data, at various stages of training. At initialization, the average F1 score was 0.08. After seeing 16,000 videos, it achieves an average F1 score of 0.30. Our results suggest that *large-scale training*, made possible through interactive annotation, is crucial for learning good decision policies.

take. These correspond to tracking when the predicted object location is correct, and updating if the appearance model trained on the new patch produces higher scores for ground-truth locations. We train a classifier to predict these actions using the current heatmap. When this offline trained classifier is run at test-time, F1 improves to .13 with the track heuristic and .14 with the update heuristic, and .20 if both are used.

FCNT vs CCOT: We use FCNT for our ablative analysis; initializing p-track using CCOT’s more complex online heuristics proved difficult. However our final system uses only our proposed offline heuristics, so we can nonetheless train it using CCOT’s TRACK, REINIT, and UPDATE functions. In Fig. 9 we compare the final p-trackers built using FCNT’s functions against those built using CCOT’s functions. As consistent with prior work, we find that CCOT improves overall performance (from .30 to .36).

Q-learning: Finally, we use Q-learning to refine our heuristics (Eq. 6), noticeably improving the F1 score to .30. Learning the appearance update action seems to have the most significant effect on performance, producing an F1 score of .28 by itself. During partial occlusions, the tracker learns to delicately balance between appearance update and drift while accepting a few failures to avoid the cost and risk of reinitialization. Overall, the learned policy dramatically outperforms the default online heuristics, **tripling** the F1 score from 9% to 30%!

Training iterations: In theory, our tracker can be interactively trained on a never-ending stream. However, in our experiments, Q-learning appeared to converge after seeing between 8,000 and 12,000 videos. Thus, we choose to stop training after seeing 16,000 videos. In Fig. 11, we plot performance vs training iteration.

Computation: As mentioned previously, comparatively slower trackers typically perform better [23]. On a Tesla K40 GPU, our tracker runs at approximately 10 fps. While computationally similar to [51], we add the ability to recover from tracking failures by reinitializing through



Figure 12. **What does p-track learn?** We show the actions taken by our tracker given four heatmaps. P-track learns to track and update appearance even in cluttered heatmaps with multiple modes (a). However, if the confidence of other modes becomes high, p-track learns not to update appearance to avoid drift due to distractors (b). If the target mode is heavily blurred, implying the target is difficult to localize (because of a transforming robot), p-track also avoids model update (c). Finally, the lack of mode suggests p-track will reinitialize (d).

detection. To do so, we learn an attention policy that efficiently balances tracking vs reinitialization. Tracking is fast because only a small region of interest (ROI) need be searched. Rather than searching over the whole image during reinitialization, we select a random ROI (which ensures that our trackers operate at a fixed frame rate). In practice, we find that target is typically found in ≈ 15 frames.

Conclusions: We formulate tracking as a sequential decision-making problem, where a tracker must update its beliefs about the target, given noisy observations and a limited computational budget. While such decisions are typically made heuristically, we bring to bear tools from POMDPs and reinforcement learning to learn decision-making strategies in a data-driven way. Our framework allows trackers to learn action policies appropriate for different scenarios, including short-term and long-term tracking. One practical observation is that offline heuristics are an effective and efficient way to learn tracking policies, both by themselves and as a regularizer for Q-learning. Finally, we demonstrate that reinforcement learning can be used to leverage massive training datasets, which will likely be needed for further progress in data-driven tracking.

Acknowledgments: This research was supported in part by the National Science Foundation (NSF) under grants CNS-1518865 and IIS-1618903, and the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001117C0051. Additional support was provided by the Intel Science and Technology Center for Visual Cloud Systems (ISTC-VCS). 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] S. Adam, L. Busoniu, and R. Babuska. Experience replay for real-time reinforcement learning control. *IEEE Transactions on Systems, Man, and Cybernetics*, 42(2):201–212, 2012.
- [2] A. Agarwala, A. Hertzmann, D. H. Salesin, and S. M. Seitz. Keyframe-based tracking for rotoscoping and animation. *ACM Transactions on Graphics (ToG)*, 23(3):584–591, 2004.
- [3] L. Bazzani, N. de Freitas, and J.-A. Ting. Learning attentional mechanisms for simultaneous object tracking and recognition with deep networks. In *NIPS 2010 Deep Learning and Unsupervised Feature Learning Workshop*, volume 32, 2010.
- [4] T. Berg and A. C. Berg. Finding iconic images. *CVPR*, 2009.
- [5] L. Bertinetto, J. Valmadre, S. Golodetz, O. Miksik, and P. H. Torr. Staple: Complementary learners for real-time tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1401–1409, 2016.
- [6] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr. Fully-convolutional siamese networks for object tracking. In *European Conference on Computer Vision*, pages 850–865. Springer, 2016.
- [7] R. A. Bianchi, C. H. Ribeiro, and A. H. R. Costa. Heuristically accelerated reinforcement learning: Theoretical and experimental results. *ECAI*, pages 169–174, 2012.
- [8] A. Buchanan and A. Fitzgibbon. Interactive feature tracking using kd trees and dynamic programming. In *CVPR*, volume 1, pages 626–633, 2006.
- [9] Z. Chi, H. Li, H. Lu, and M.-H. Yang. Dual deep network for visual tracking. *IEEE Transactions on Image Processing*, 26(4):2005–2015, 2017.
- [10] L. Chukoskie, J. Snider, M. C. Mozer, R. J. Krauzlis, and T. J. Sejnowski. Learning where to look for a hidden target. *Proceedings of the National Academy of Sciences*, 110(Supplement 2):10438–10445, 2013.
- [11] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg. Learning spatially regularized correlation filters for visual tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4310–4318, 2015.
- [12] M. Danelljan, A. Robinson, F. S. Khan, and M. Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *European Conference on Computer Vision*, pages 472–488. Springer, 2016.
- [13] D. Du, H. Qi, W. Li, L. Wen, Q. Huang, and S. Lyu. Online deformable object tracking based on structure-aware hypergraph. *IEEE Transactions on Image Processing*, 25(8):3572–3584, 2016.
- [14] M. Edoardo Maresca and A. Petrosino. The matrioska tracking algorithm on LTDT2014 dataset. *CVPR workshop on LTDT*, 2014.
- [15] M. Hausknecht and P. Stone. Deep recurrent q-learning for partially observable MDPs. *arXiv:1507.06527*, 2015.
- [16] Z. Hong, Z. Chen, C. Wang, X. Mei, D. Prokhorov, and D. Tao. Multi-store tracker (muster): a cognitive psychology inspired approach to object tracking. *CVPR*, pages 749–758, 2015.
- [17] Y. Hua, K. Alahari, and C. Schmid. Online object tracking with proposal selection. *ICCV*, pages 3092–3100, 2015.
- [18] M. Jiang, X. Boix, G. Roig, J. Xu, L. Van Gool, and Q. Zhao. Learning to predict sequences of human visual fixations. *IEEE transactions on neural networks and learning systems*, 27(6):1241–1252, 2016.
- [19] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.
- [20] S. E. Kahou, V. Michalski, and R. Memisevic. RATM: Recurrent attentive tracking model. *CVPR*, 2016.
- [21] Z. Kalal, K. Mikolajczyk, and J. Matas. Tracking-learning-detection. *TPAMI*, 34(7):1409–1422, 2012.
- [22] S. Khim, S. Hong, Y. Kim, and P. kyu Rhee. Adaptive visual tracking using the prioritized q-learning algorithm: Mdp-based parameter learning approach. *Image and Vision Computing*, 32(12):1090–1101, 2014.
- [23] M. Kristan, A. Leonardis, J. Matas, M. Felsberg, R. Pflugfelder, L. Čehovin, T. Vojir, G. Häger, A. Lukežič, and G. Fernandez. The visual object tracking vot2016 challenge results. Springer, Oct 2016.
- [24] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler. MOTchallenge 2015: Towards a benchmark for multi-target tracking. *arXiv:1504.01942*, 2015.
- [25] J.-Y. Lee and W. Yu. Visual tracking by partition-based histogram backprojection and maximum support criteria. In *Robotics and Biomimetics (ROBIO), 2011 IEEE International Conference on*, pages 2860–2865. IEEE, 2011.
- [26] A. Li, M. Lin, Y. Wu, M.-H. Yang, and S. Yan. NUS-PRO: A new visual tracking challenge. 2015.
- [27] H. Li, Y. Li, F. Porikli, et al. Deeptrack: Learning discriminative feature representations by convolutional neural networks for visual tracking. *BMVC*, 1(2):3, 2014.
- [28] Y. Li, C. Huang, and R. Nevatia. Learning to associate: Hybridboosted multi-target tracker for crowded scene. *CVPR*, pages 2953–2960, 2009.
- [29] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang. Hierarchical convolutional features for visual tracking. *ICCV*, pages 3074–3082, 2015.
- [30] C. Ma, X. Yang, C. Zhang, and M.-H. Yang. Long-term correlation tracking. *CVPR*, pages 5388–5396, 2015.
- [31] S. Mathe, A. Pirinen, and C. Sminchisescu. Reinforcement learning for visual object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2894–2902, 2016.
- [32] G. A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995.
- [33] V. Mnih, N. Heess, A. Graves, et al. Recurrent models of visual attention. *NIPS*, pages 2204–2212, 2014.
- [34] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [35] H. Nam, M. Baek, and B. Han. Modeling and propagating cnns in a tree structure for visual tracking. *arXiv preprint arXiv:1608.07242*, 2016.

- [36] H. Nam and B. Han. Learning Multi Domain Convolutional Neural Networks for Visual Tracking. *arXiv:1510.07945*, 2015.
- [37] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4293–4302, 2016.
- [38] L. Paletta, G. Fritz, and C. Seifert. Q-learning of sequential attention for visual object recognition from informative local descriptors. *ICML*, pages 649–656, 2005.
- [39] S. Palmer, E. Rosch, and P. Chase. Canonical perspective and the perception of objects. *Attention and performance*, 1(4), 1981.
- [40] Y. Pang and H. Ling. Finding the Best from the Second Bests - Inhibiting Subjective Bias in Evaluation of Visual Tracking Algorithms. *ICCV*, pages 2784–2791, dec 2013.
- [41] F. Pernici and A. Del Bimbo. Object tracking by oversampling local features. *TPAMI*, 36(12):2538–2551, 2014.
- [42] 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.
- [43] S. J. Russell, P. Norvig, J. F. Canny, J. M. Malik, and D. D. Edwards. *Artificial intelligence: a modern approach*. Prentice Hall, 2003.
- [44] A. W. Smeulders, D. M. Chu, R. Cucchiara, S. Calderara, A. Dehghan, and M. Shah. Visual tracking: An experimental survey. *TPAMI*, 36(7):1442–1468, 2014.
- [45] A. Solis Montero, J. Lang, and R. Laganiere. Scalable kernel correlation filter with sparse feature integration. *ICCV Workshops*, December 2015.
- [46] S. Song and J. Xiao. Tracking revisited using RGBD camera: Unified benchmark and baselines. *ICCV*, pages 233–240, 2013.
- [47] J. Supancic and D. Ramanan. Self-paced learning for long-term tracking. *CVPR*, pages 2379–2386, 2013.
- [48] A. Torralba and A. A. Efros. Unbiased look at dataset bias. *CVPR*, pages 1521–1528, 2011.
- [49] T. Vojir, J. Nuskova, and J. Matas. Robust scale-adaptive mean-shift for tracking. *Image Analysis*, pages 652–663, 2013.
- [50] C. Vondrick and D. Ramanan. Video annotation and tracking with active learning. In *NIPS*, pages 28–36, 2011.
- [51] L. Wang, W. Ouyang, X. Wang, and H. Lu. Visual Tracking with Fully Convolutional Networks. *CVPR*, 2015.
- [52] L. Wang, W. Ouyang, X. Wang, and H. Lu. STCT: Sequentially training convolutional networks for visual tracking. *CVPR*, 2016.
- [53] L. Wang, W. Ouyang, X. Wang, and H. Lu. Stct: Sequentially training convolutional networks for visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1373–1381, 2016.
- [54] N. Wang, J. Shi, D.-Y. Yeung, and J. Jia. Understanding and diagnosing visual tracking systems. *ICCV*, pages 3101–3109, 2015.
- [55] Y. Wu, J. Lim, and M.-H. Yang. Online object tracking: A benchmark. *CVPR*, pages 2411–2418, 2013.
- [56] Y. Xiang, A. Alahi, and S. Savarese. Learning to track: Online multi-object tracking by decision making. *ICCV*, pages 4705–4713, 2015.
- [57] G. Zhu, F. Porikli, and H. Li. Beyond local search: Tracking objects everywhere with instance-specific proposals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 943–951, 2016.