

# Long-Term Visual Object Tracking Benchmark

by

Abhinav Moudgil, Vineet Gandhi

in

*14th Asian Conference on Computer Vision  
(ACCV-2018)*

Perth, Western Australia

Report No: IIIT/TR/2018/-1



Centre for Visual Information Technology  
International Institute of Information Technology  
Hyderabad - 500 032, INDIA  
December 2018

# Long-Term Visual Object Tracking Benchmark

Abhinav Moudgil and Vineet Gandhi

Center for Visual Information Technology, Kohli Center on Intelligent Systems  
International Institute of Information Technology, Hyderabad, India  
abhinav.moudgil@research.iiit.ac.in, vgandhi@iiit.ac.in

**Abstract.** We propose a new long video dataset<sup>1</sup> (called Track Long and Prosper - TLP) and benchmark for single object tracking. The dataset consists of 50 HD videos from real world scenarios, encompassing a duration of over 400 minutes (676K frames), making it more than 20 folds larger in average duration per sequence and more than 8 folds larger in terms of total covered duration, as compared to existing generic datasets for visual tracking. The proposed dataset paves a way to suitably assess long term tracking performance and train better deep learning architectures (avoiding/reducing augmentation, which may not reflect real world behaviour). We benchmark the dataset on 17 state of the art trackers and rank them according to tracking accuracy and run time speeds. We further present thorough qualitative and quantitative evaluation highlighting the importance of long term aspect of tracking. Our most interesting observations are (a) existing short sequence benchmarks fail to bring out the inherent differences in tracking algorithms which widen up while tracking on long sequences and (b) the accuracy of trackers abruptly drops on challenging long sequences, suggesting the potential need of research efforts in the direction of long-term tracking.

## 1 Introduction

Visual tracking is a fundamental task in computer vision and is a key component in wide range of applications like surveillance, autonomous navigation, video analysis and editing, augmented reality etc. Many of these applications rely on long-term tracking, however, only few tracking algorithms have focused on the challenges specific to long duration aspect [19,31,18,36]. Although they conceptually attack the long term aspect, the evaluation is limited to shorter sequences or couple of selected longer videos. The recent correlation filter [6,9,2,45] and deep learning [43,33,3,14] based approaches have significantly advanced the field, however, their long term applicability is also unapparent as the evaluation is limited to datasets with typical average video duration of about 20-40 seconds. Not just the evaluation aspect, the lack of long term tracking datasets has been a hindrance for training in several recent state of the art approaches. These methods either limit themselves to available small sequence data [33,43] or use augmentation on datasets designed for other tasks like object detection [14].

<sup>1</sup> Dataset and tracking results are available at <https://amoudgl.github.io/tlp/>

Motivated by the above observation, we propose a new long duration dataset called Track Long and Prosper (TLP), consisting of 50 long sequences. The dataset covers a wide variety of target subjects and is arguably one of the most challenging datasets in terms of occlusions, fast motion, viewpoint change, scale variations etc. However, compared to existing generic datasets, the most prominent aspect of TLP dataset is that it is larger by more than 20 folds in terms of average duration per sequence, which makes it ideal to study challenges specific to long duration aspect. For example, drift is a common problem in several tracking algorithms and it is not always abrupt and may occur due to accumulation of error over time (which may be a slow procedure and can be difficult to gauge in short sequences). Similarly, long sequences allow us to study the consistency of a tracker to recover from momentary failures.

We select 17 recent state of the art trackers which are scalable to be evaluated on TLP dataset and provide a thorough evaluation in terms of tracking accuracy and real time performance. Testing on such a large dataset significantly reduces the overfitting problem, if any, and reflects if the tracker is actually designed to consistently recover from challenging scenarios. To present a further perspective, we provide a comprehensive attribute wise comparison of different tracking algorithms by selecting various sets of short sequences (derived from original TLP sequences), in which each set only contains sequences where a particular type of challenge is dominant (like illumination variation, occlusions, out of view etc.).

We observe that the rankings from previous short sequence datasets like OTB50 [41] significantly vary from the rankings obtained on the proposed TLP dataset. Several top ranked trackers on recent benchmarks fail to adapt to long-term tracking scenario and their performance drops significantly. Additionally, the performance margin notably widens among several trackers, whose performances are imperceptibly close in existing benchmarks. More specifically, apart from MDNet [33], performance of all other evaluated tracker drops below 25% on commonly used metric of area under the curve of success plots. Our investigation hence strongly highlights the need for more research efforts in long term tracking and to our knowledge the proposed dataset and benchmark is the first systematic exploration in this direction.

## 2 Related Work

### 2.1 Tracking Datasets

There are several existing datasets which are widely used for evaluating the tracking algorithms and are summarized in Table 2.1. The OTB50 [41], OTB100 [42] are the most commonly used ones. They include 50 and 100 sequences respectively and capture a generic real world scenario (where some videos are taken from platforms like YouTube and some are specifically recorded for tracking application). They provide per frame bounding box annotation and per sequence annotation of attributes like illumination variation, occlusion, deformation etc.

The ALOV300++ dataset [34] focuses on diversity and includes more than 300 short sequences (average length of only about 9 seconds). The annotations

**Table 1.** Comparing TLP with other object tracking datasets.

	Frame rate (FPS)	# videos	Min Duration (sec)	Mean Duration (sec)	Max Duration (sec)	Total Duration (sec)
UAV123[32]	30	123	3.6	30.5	102.8	3752
OTB50[41]	30	51	2.3	19.3	129	983
OTB100[42]	30	100	2.3	19.6	129	1968
TC128[27]	30	129	2.3	14.3	129	1844
VOT14[23]	30	25	5.7	13.8	40.5	346
VOT15[22]	30	60	1.6	12.2	50.2	729
ALOV300[34]	30	314	0.6	9.2	35	2978
NFS[11]	240	100	0.7	16	86.1	1595
<b>TLP</b>	<b>24/30</b>	<b>50</b>	<b>144</b>	<b>484.8</b>	<b>953</b>	<b>24240</b>

in ALOV300++ dataset are made every fifth frame. A small set of challenging sequences (partially derived from OTB50, OTB100 and ALOV300++ datasets) has been used in VOT14 [23] and VOT15 [22] datasets. They extend the rectangular annotations to rotated ones and provide per frame attribute annotations, for more accurate evaluation. Both of these datasets have been instrumental in yearly visual object tracking (VOT) challenge.

Some datasets have focused on particular type of applications/aspects. TC128 [27] was proposed to study the role of color information in tracking. It consists of 128 sequences (some of them are common to OTB100 dataset) and provides per frame annotations and sequence wise attributes. Similarly, UAV [32] targets the tracking application, when the videos are captured from low-altitude unmanned aerial vehicles. The focus of their work is to highlight challenges incurred while tracking in video taken from an aerial viewpoint. They provide both real and synthetically generated UAV videos with per frame annotations.

More recently, two datasets were proposed to incorporate the benefits of advances in capture technology. The NFS [11] dataset was proposed to study the fine grained variations in tracking by capturing high frame rate videos (240 FPS). Their analysis shows that since high frame video reduces appearance variation per frame, it is possible to achieve state of the art performance using substantially simpler tracking algorithms. Another recent dataset called AMP [46], explores the utility of 360° videos to generate and study tracking with typical motion patterns (which can be achieved by varying the camera re-parametrization in omni-directional videos). Contemporary to our work, [29] and [37] also review recent trackers for long-term tracking. However, they limit the long-term tracking definition to the ability of a tracker to re-detect after object goes out of view and the quality of their long term datasets is lower than our proposed TLP dataset, in terms of resolution and per sequence length. We evaluate the trackers from a holistic perspective and show that even if there is no apparent major challenge or target disappearance, tracking consistently for a long period of time is an extremely challenging task.

Although recent advances pave the way to explore several novel and specific fine grained aspects, the crucial long term tracking aspect is still missing from most of the current datasets. The typical average length per sequence is still only about 10-30 seconds. The proposed TLP dataset takes it to about 8-9 minutes

per sequence, making it the largest densely annotated high-resolution dataset for the application of visual object tracking.

## 2.2 Tracking Methods

Most of the earlier approaches trained a model/classifier considering the initial bounding box marked by the user as “foreground” and areas farther away from the annotated box as “background”. The major challenge in most of these approaches is to properly update the model or the classifier over time to reduce drift. This problem has been tackled in several innovative ways, such as Multiple Instance Learning [1], Online Boosting [12], P-N Learning [19], or by using ensemble of classifiers trained/initiated at difference instances of time [44]. Most of the recent advances, however, have focused only in two major directions i.e Correlation Filter (CF) based tracking [6,9,2,45] and deep learning based tracking [43,33,3,14]. The CF based trackers have gained huge attention due to their computational efficiency derived by operating in Fourier domain and the capability of efficient online adaptation. The interest in CF based approaches was kindled by the MOSSE [4] tracker proposed by Bolme *et al.*, which demonstrated an impressive speed of about 700 FPS. Thereafter, several works have built upon this idea and have significantly improved tracking accuracy. The list includes ideas of using kernelized correlation filters [15]; exploiting multi-dimensional features [10,16]; combining template based features with pixel wise information for robustness to deformation and scale variations [26,2]; employing kernel ridge regression to reduce drift [31] etc. The work by Kiani *et al.* [21] identified the boundary effects in Fourier domain as one of the reasons for sub-optimal performance of CF based approaches. Solutions such as Spatially regularized CF [8] and Background aware CF [20] were later proposed to mitigate the boundary effects.

More recent efforts in CF based trackers utilize deep convolutional features instead of hand crafted ones like HOG [5]. Multiple convolutional layers in a hierarchical ensemble of independent DCF trackers was employed by Ma *et al.* [30]. Danelljan *et al.* [9] extended it by fusing multiple convolutional layers with different spatial resolutions in a joint learning framework. This combination of deep CNN features has led CF based trackers to the top of several benchmarks like OTB50; however, they come with an additional computational cost. Some recent efforts have been made to enhance running speeds by using ideas like factorized convolutions [6]; however, the speeds are still much slower than the traditional CF trackers.

Deep learning based trackers present another paradigm in visual object tracking. Data deficiency appeared to be the major limitation in early attempts [40,25]. Later approaches [39,17] learnt offline tasks like objectness or saliency offline from object detection datasets and benefited from it during online tracking. However, the gap between the two tasks turned out to be a limitation in these works. The work by Nam *et al.* [33] proposed a novel direction by posing the tracking problem as evaluating the positive and negative candidate windows randomly sampled around the previous target state. They proposed a two phase



**Fig. 1.** First frames of all the 50 sequences of TLP dataset. The sequences are sorted in ascending order on the basis of mean success rate (defined in Section 4) of all trackers at IoU threshold of 0.5. The sequences at the bottom right are more difficult to track than the ones at the top left.

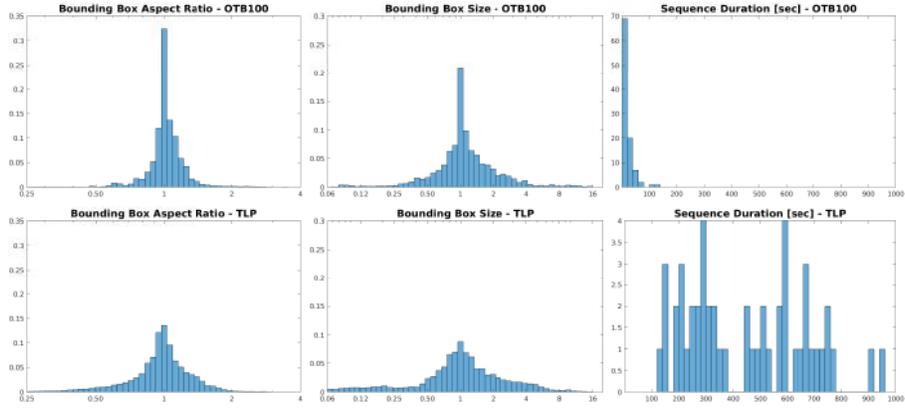
training, first domain dependent and offline to fix initial layers and second online phase to update the last fully connected layer.

Siamese frameworks [3,14] have been proposed to regress the location of the target, given the previous frame location. They either employ data augmentation using affine transformations on individual images [14] or use video detection datasets for large scale offline training [3]. In both these approaches, the network is evaluated without any fine tuning at the test time, which significantly increases their computational efficiency. However, this comes at the cost of losing ability to update appearance models or learn target specific information, which may be crucial for visual tracking. More recently, Yun *et al.* [43] proposed a tracker controlled by action-decision network (ADNet), which pursues the target object by sequential actions iteratively, utilizing reinforcement learning for visual tracking.

### 3 TLP Dataset

The TLP dataset consists of 50 videos collected from YouTube. The dataset was carefully curated with 25 indoor and 25 outdoor sequences covering a large variety of scene types like sky, sea/water, road/ground, ice, theatre stage, sports arena, cage etc. Tracking targets include both rigid and deformable/articulated objects like vehicle (motorcycle, car, bicycle), person, face, animal (fish, lion, puppies, birds, elephants, polar bear), aircraft (helicopter, jet), boat and other generic objects (e.g sports ball). The application aspect was also kept into account while selecting the sequences, for example we include long sequences from theatre performances, music videos and movies, which are rich in content, and tracking in them may be useful in context of several recent applications like virtual camera simulation or video stabilization [24,13]. Similarly, long term tracking in sports videos can be quite helpful for automated analytics [28]. The large variation in scene type and tracking targets can be observed in Figure 1. We further compare the TLP dataset with OTB in Figure 2, to highlight that the variation in bounding box size and aspect ratio with respect to the initial frame is significantly larger in TLP and the variations are also well balanced. The significant differences in duration of sequences in OTB and TLP are also apparent.

The per sequence average length in TLP dataset is over 8 minutes. Each sequence is annotated with rectangular bounding boxes per frame, which were done using the VATIC [38] toolbox. The annotation format is similar to OTB50 and OTB100 benchmarks to allow for easy integration with existing toolboxes. We have 33/50 sequences (amounting to 4% frames in total) in TLP dataset where the target goes completely out of view and thus, we provide absent label for each frame in addition to the bounding box annotation. All the selected sequences are single shot (do not contain any cut) and have a resolution of  $1280 \times 720$ . Similar to VOT [22], we choose the sequences without any cuts, to be empirically fair in evaluation, as most trackers do not explicitly model a re-detection policy. However, the recovery aspect of trackers still gets thoroughly



**Fig. 2.** Column 1 and 2: Proportional change of the targets aspect ratio and bounding box size (area in pixels) with respect to the first frame in OTB100 and TLP. Results are compiled over all sequences in each dataset as a histogram with log scale on the x-axis. Column 3: Histogram of sequence duration (in seconds) across the two datasets.

evaluated on the TLP dataset, due to presence of full occlusions and out of view scenarios in several sequences.

**TinyTLP and TLPattr:** We further derive two short sequence datasets from TLP dataset. The TinyTLP dataset consists of first 600 frames (20 sec) in each sequence of the TLP dataset to compare and highlight the challenges incurred due to long-term tracking aspect. The length of 20 sec was chosen to align the average per sequence length with OTB100 benchmark. The TLPattr dataset consists of total 90 short sequences focusing on different attributes. Six different attributes were considered in our work i.e (a) fast motion of target object or camera, (b) illumination variation around target object between consecutive frames, (c) large scale variation of the target object, (d) partial occlusions of the target object by other objects or background, (e) out of view or full occlusions, where object leaves the camera view or it is not visible at all and (f) background clutter. The TLPattr dataset includes 15 short sequences corresponding to each of the attribute.

Each sequence in TLPattr is carefully selected in such a way that the only dominant challenge present in it is a particular attribute, it is assigned to. For example, for fast motion, we first select all instances in entire TLP dataset where the motion of the center of the ground truth bounding box between consecutive frames is more than 20 pixels. We temporally locate every such fast motion event and curate a short sequence around it by selecting 100 frames before and after the fast motion event. We then sort the short sequences based on the amount of motion (with the instance with most movement between two frames as the top sequence) and manually shortlist 15 sequences (starting from the top), where fast motion is the only dominant challenge present and simultaneously avoiding selection of multiple short sequences from the same long video. For attributes

like illumination variation and background clutter the selection was fully manual. The rationale behind curating the TLPattr dataset was the following: (a) Giving a single attribute to entire sequence (as in previous works like OTB50) is ill posed on long sequences as in TLP. Any attribute based analysis with such an annotation would not capture the correct correlation between the challenge and the performance of the tracking algorithm. (b) Using per frame annotation of attributes is also difficult for analysis in long videos, as the tracker may often fail before reaching the particular frame where attribute is present and (c) The long sequences and variety present in TLP dataset allows us to single out a particular attribute and choose subsequences where that is the only dominant challenge. This paves the way for accurate attribute wise analysis.

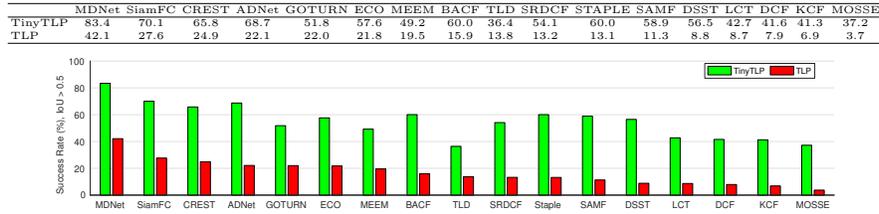
## 4 Evaluation

### 4.1 Evaluated Algorithms

We evaluated 17 recent trackers on the TLP and TinyTLP datasets. The trackers were selected based on three broad guidelines i.e.: (a) they are computationally efficient for large scale experiments; (b) their source codes are publicly available and (c) they are among the top performing trackers in existing benchmarks. Our list includes CF trackers with hand crafted features, namely SRDCF [8], MOSSE [4], DCF [16], DSST [7], KCF [16], SAMF [26], Staple [2], BACF [20] and LCT [31]; CF trackers with with deep features: ECO [6] and CREST [35] and deep trackers i.e. GOTURN [14], MDNet [33], ADNet [43] and SiamFC[3]. We also included TLD [19] and MEEM [44] as two older trackers based on PN learning and SVM ensemble, as they specifically target the drift problem for long-term applications. We use default parameters on the publicly available version of the code when evaluating all the tracking algorithms.

### 4.2 Evaluation Methodology

We use precision plot, success plot and longest subsequence measure for evaluating the algorithms. The precision plot [1,41] shows the percentage of frames whose estimated location is within the given threshold distance of the ground truth. A representative score per tracker is computed, by fixing a threshold over the distance (we use the threshold as 20 pixels). The success metric [41] computes the intersection over union (IoU) of predicted and ground truth bounding boxes and counts the number of successful frames whose IoU is larger than a given threshold. In out of view scenarios, if the tracking algorithm explicitly predicts the absence, we give it an overlap of 1 otherwise 0. The success plot shows the ratio of successful frames as the IoU threshold is varied from 0 to 1. A representative score for ranking the trackers is computed as the area under curve (AUC) of its success plot. We also employ the conventional success rate measure, counting frames above the threshold of 0.50 (IoU > 0.50).



**Fig. 3.** Success rate of individual trackers on TinyTLP and TLP datasets. The algorithms are sorted based on their performance on TLP.

**LSM metric:** We further propose a new metric called Longest Subsequence Measure (LSM) to quantify the long term tracking behaviour. The LSM metric computes the ratio of the length of the longest successfully tracked continuous subsequence to the total length of the sequence. A subsequence is marked as successfully tracked, if  $x\%$  of frames within it have  $\text{IoU} > 0.5$ , where  $x$  is a parameter. LSM plot shows the variation in the normalized length of longest tracked subsequence per sequence, as  $x$  is varied. A representative score per tracker can be computed by fixing the parameter  $x$  (we use the threshold as 0.95).

The LSM metric captures the ability of a tracker to track continuously in a sequence within a certain bound on failure tolerance (parameter  $x$ ) and bridges the gap over existing metrics which fail to address the issue of frequent momentary failures. For example, it often happens in long sequences that tracker loses the target at some location and freezes there. If coincidentally the target passes the same location (after a while), the tracker starts tracking it again. LSM penalizes such scenarios by considering only the longest continuous tracked subsequences.

### 4.3 Per Tracker Evaluation

Table 4.3 presents the success rate of each individual tracker on TinyTLP and TLP datasets. The MDNet tracker is the best performing tracker on both the datasets. TLD is the worst performing tracker on TinyTLP and MOSSE performs worst on TLP dataset. The performance significantly drops for each tracker on TLP dataset, when compared to TinyTLP dataset, which clearly brings out the challenges incurred in long-term tracking. The relative performance drop is minimum in MDNet where the success rate reduces from 83.4% to 42.1% (roughly by a factor of 2) and is most in MOSSE tracker, which reduces from 37.2% in TinyTLP to 3.7% in TLP (reduction by more than a factor of 10).

In general, the relative performance decrease is more in CF trackers with hand crafted features as compared to CF+deep trackers. For instance, trackers like BACF, SAMF, Staple give competitive or even better performance than CREST and ECO over TinyTLP dataset, however, their performance steeply decreases on TLP dataset. Although all the CF based trackers (hand crafted or CNN based) are quite susceptible to challenges such as long term occlusions or fast appearance changes, our experiments suggest that using learnt deep features

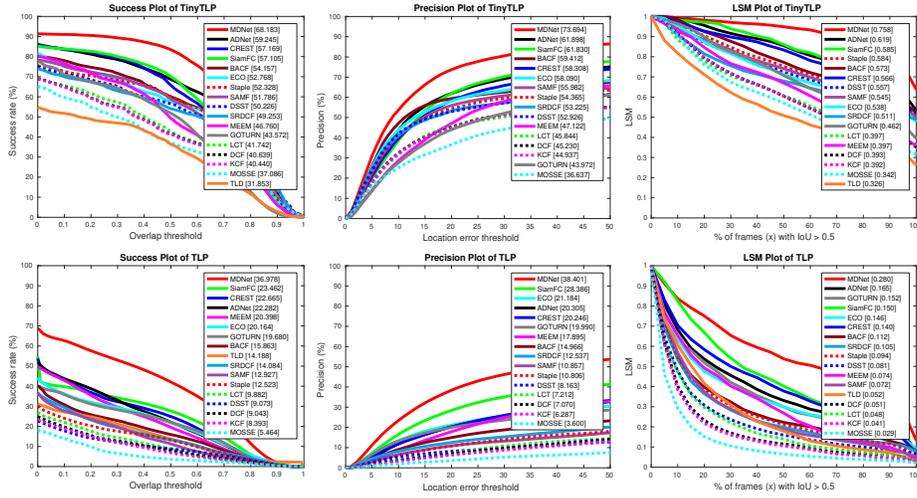
reduces accumulation of error over time and reduces drift. Such accumulation of error is difficult to quantify in short sequences and the performance comparison may not reflect the true ability of the tracker. For example, BACF outperforms ECO on TinyTLP by about 2%, however it is 6% worse than ECO on TLP. Similarly, the performance difference of SAMF and ECO is imperceptible on TinyTLP, which differs by almost a factor of 2 on TLP.

The deep trackers outperform other trackers on TLP dataset, with MDNet and SiamFC being the top performing ones. ADNet is third best tracker on TinyTLP, however, its performance significantly degrades on TLP dataset. It is interesting to observe that both MDNet and ADNet refine last fully connected layer during online tracking phase, however, MDNet appears to be more consistent and considerably outperforms ADNet on TLP. The offline trained and frozen SiamFC and GOTURN perform relatively well (both appearing in top five trackers on TLP), however SiamFC outperforms GOTURN, possibly because it is trained on larger amount of video data. Another important observation is that the performance of MEEM surpasses all state of the art CF trackers with hand crafted features on TLP dataset. The ability to recover from failures also allows TLD tracker (giving lowest accuracy on TinyTLP) to outperform several recent CF trackers on TLP.

#### 4.4 Overall Performance

The overall comparison of all trackers on TinyTLP and TLP using Success plot, Precision plot and LSM plot are demonstrated in Figure 4.4. In success plots, MDNet clearly outperforms all the other trackers on both TinyTLP and TLP datasets with AUC measure of 68.1% and 36.9% respectively. It is also interesting to observe that the performance gap significantly widens up on TLP and MDNet clearly stands out from all other algorithms. This suggests that the idea of separating domain specific information during training and online fine tuning of background and foreground specific information, turns out to be an extremely important one for long term tracking. Furthermore, analyzing MDNet and ADNet both of which employ the strategy of online updates on last FC layers during tracking, it appears that learning to detect instead of learning to track gives a more robust performance in long sequences. The performance drop of SiamFC and GOTURN on TLP also suggests a similar hypothesis.

The steeper success plots in TLP as compared to TinyTLP dataset, suggest that accurate tracking gets more and more difficult in longer sequences, possibly due to accumulation of error. The lower beginning point on TLP (around 40-50% for most trackers compared to 80-90% on TinyTLP), indicates that most trackers entirely drift away before reaching halfway through the sequence. The rankings in success plot on TLP are also quite contrasting to previous benchmarks. For instance, ECO is the best performing tracker on OTB100 closely followed by MDNet (with almost imperceptible difference), and its performance significantly slides on TLP. Interestingly, MEEM breaks into top five trackers in AUC measure of success plot on TLP (ahead of ECO). In general there is striking drop of performance between TinyTLP and TLP for most CF based trackers (more so

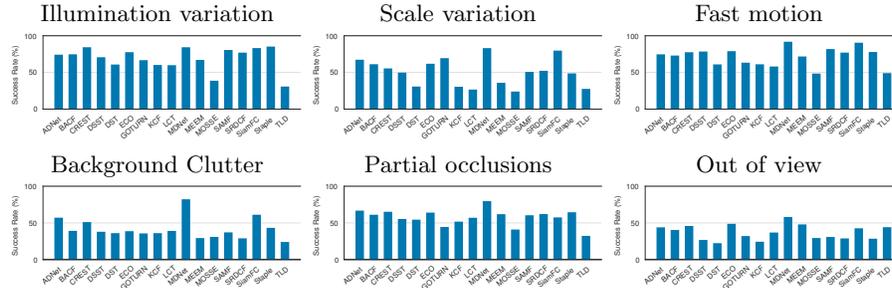


**Fig. 4.** Overall Performance of evaluated trackers on TinyTLP and TLP with success plot, precision plot and LSM plot respectively (each column). For each plot, ranked trackers are shown with corresponding representative measure i.e. AUC in success plots; 20 pixel threshold in precision plots and 0.95 as length ratio in LSM plots.

for hand crafted ones). CREST is most consistent among them and ranks in top 5 trackers for both TinyTLP and TLP.

The precision plots also demonstrate similar trends as success plots, however they bring couple of additional subtle and interesting perspectives. The first observation is that SiamFC’s performance moves closer to performance of MDNet on TLP dataset. Since SiamFC is fully trained offline and does not make any online updates, it is not accurate in scaling the bounding box to the target in long term, which brings down its performance in IoU measure. However, it still hangs on to the target due to the large scale training to handle challenges in predicting the consecutive bounding boxes, hence the numbers improve in the precision plot (again precision plot on TinyTLP does not capture this observation). The ADNet tracker is ranked two on TinyTLP using precision measure, however, it drops to 4th position on TLP. The GOTURN tracker also brings minor relative improvement in precision measure and moves ahead of MEEM on TLP.

The LSM plots show the ratio of longest successfully tracked continuous subsequence to the total length of the sequence. The ratios are finally averaged over all the sequences for each tracker. A sequence is successfully tracked if  $x\%$  of frames in it have IoU > 0.5. We vary the value  $x$  to draw the plots and the representative number is computed by keeping  $x = 95\%$ . This measure explicitly quantifies the ability to continuously track without failure. MDNet performs the best on this measure as well. The relative performance of CREST drops in LSM measure, as it partially drifts away quite often, however is able to recover from it



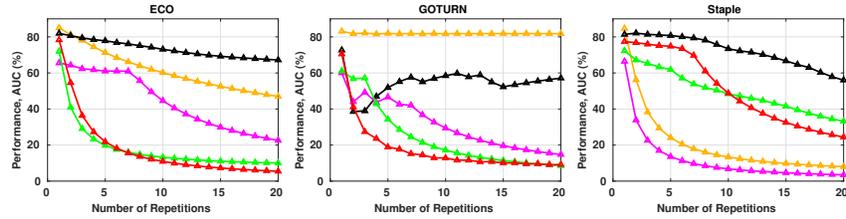
**Fig. 5.** Attribute wise performance evaluation on TLPattr dataset. Results are reported as success rate (%) with IoU > 0.5.

as well. So its overall success rate is higher, however, the average length of longest continuous set of frames it can track in a video is relatively low. In general, the ratio of largest continuously tracked subsequence to sequence length (with success rate > 0.95) averaged over all sequences is about 1/4th for MDNet and lower than 1/6th for other trackers. This indicates the challenge in continuous accurate tracking without failures.

#### 4.5 Attribute wise Performance Evaluation

The average attribute wise success rates of all the trackers on TLPattr dataset are shown in Figure 5. Each attribute in TLPattr dataset includes 15 short sequences corresponding to it (dominantly representing the particular challenge). Out of view appears to be the most difficult challenge hindering the performance of the trackers followed by background clutter, scale variation and partial occlusions. Most of the trackers seem to perform relatively better on sequences with illumination variation and fast motion. On individual tracker wise comparison, MDNet gives best performance across all the attributes, clearly indicating the tracker’s reliable performance across different challenges.

Another important perspective to draw from this experiment is that the analysis on short sequences (even if extremely challenging) is still not a clear indicator of their performance on long videos. For example, Staple and CREST are competitive in performance across all the attributes, however their performance on full TLP dataset differs by almost a factor of two in success rate measure (CREST giving a value 24.9 and Staple is only 13.1). Similarly comparison can be drawn between DSST and GOTURN, which are competitive in per attribute evaluation (with DSST performing better than GOTURN on fast motion, partial occlusions, background clutter and illumination variation). However, in long terms setting, their performance varies by a large margin (GOTURN giving success rate of 22.0, while DSST is much inferior with a value of 8.8).



**Fig. 6.** Results of three different trackers on 20 times elongated TinyTLP sequences (by reversing and concatenating the sequence in iterative way). Each color represents a different sequence and each triangle represents a repetition.

#### 4.6 Evaluation on repeated TinyTLP sequences

The essence of our paper is the need to think “long term” in object tracking, which is crucial for most practical applications of tracking. However, it remains unclear if there exists a “long term challenge in itself” and one can always argue that the performance drop in long videos is just because of “more challenges” or “frequent challenges”. To investigate this further, we conduct a small experiment where we take a short sequence and repeat it 20 times to make a longer video out of it, by iteratively reversing and attaching it at the end to maintain the continuity. This increases the length of the sequence without introducing any new difficulty or challenges. In Figure 6, we present such an experiment with three different trackers ECO (deep+CF, best performing tracker on OTB), GOTURN (pure deep) and Staple (pure CF) on 5 TinyTLP sequences for each tracker, where the tracker performs extremely well in the first iteration. We can observe that the tracking performance degrades for all three algorithms (either gradually or steeply) as the sequences get longer, which occurs possibly due to error accumulated over time. This again highlights the fact the tracking performance not just depends on the challenges present in the sequence but also gets affected by the length of the video. Hence, a dataset like TLP, interleaving the challenges and the long term aspect, is necessary for comprehensive evaluation of tracking algorithms.

#### 4.7 Run time comparisons

The run time speeds of all the evaluated algorithms are presented in Figure 7. For fair evaluation, we tested all the CPU algorithms on a 2.4GHz Intel Xeon CPU with 32GB RAM and we use a NVIDIA GeForce GTX 1080 Ti GPU for testing GPU algorithms. The CF based trackers clearly are most computationally efficient and even CPU algorithms run several folds faster than real time. The deep CF and deep trackers are computationally more expensive. MDNet gives lowest tracking speeds and runs at 1 FPS even on GPU. Among deep trackers GOTURN is the fastest tracker, however SiamFC and ADNet bring a good trade off in terms of overall success rate and run time speeds on GPU.

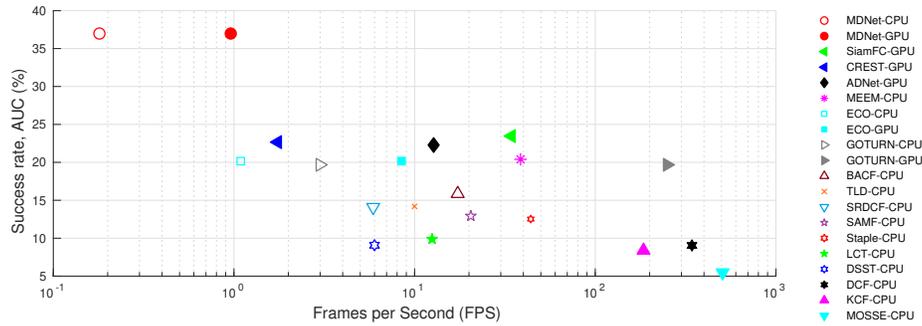


Fig. 7. Runtime comparison of different tracking algorithms.

## 5 Conclusion

This work aims to emphasize the fact that tracking on large number of tiny sequences, does not clearly bring out the competence or potential of a tracking algorithm. Moreover, even if a tracking algorithm works well on extremely challenging small sequences and fails on moderately difficult long sequences, it will be of limited practical importance. To this end, we propose the TLP dataset, focusing on the long term tracking application, with notably larger average duration per sequence, a factor which is of extreme importance and has been neglected in the existing benchmarks. We evaluate 17 state of the art algorithms on the TLP dataset, and the results clearly demonstrate that almost all state of the art tracking algorithms do not generalize well on long sequence tracking, MDNet being the only algorithm achieving more than 25% on the AUC measure of success plots. However, MDNet is also the slowest among the evaluated 17 trackers in terms of run time speeds.

Interestingly, if we only select the first 20 seconds of each sequence for evaluation (calling it TinyTLP dataset), the performance of all the trackers increases by multiple folds across different metrics. Another important observation is that the evaluations on small datasets fail to efficiently discriminate the performances of different tracking algorithms, and closely competing algorithms on TinyTLP result in quite different performance on TLP. The dominant performance of MDNet suggests that the ideas of online updating the domain specific knowledge and learning a classifier cum detector instead of a tracker (which regresses the shift), are possibly some cues to improve the performance in long term setting. Our evaluation on repeated TinyTLP sequences shows that temporal depth indeed plays an important role in the performance of evaluated trackers and appropriately brings out their strengths and weaknesses. To the best of our knowledge, TLP benchmark is the first large-scale evaluation of the state of the art trackers, focusing on long duration aspect and makes a strong case for much needed research efforts in this direction, in order to track long and prosper.

## References

1. Babenko, B., Yang, M.H., Belongie, S.: Robust object tracking with online multiple instance learning. *IEEE transactions on pattern analysis and machine intelligence* **33**(8), 1619–1632 (2011)
2. Bertinetto, L., Valmadre, J., Golodetz, S., Miksik, O., Torr, P.H.S.: Staple: Complementary learners for real-time tracking. In: *CVPR* (June 2016)
3. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. *arXiv preprint arXiv:1606.09549* (2016)
4. Bolme, D.S., Beveridge, J.R., Draper, B.A., Lui, Y.M.: Visual object tracking using adaptive correlation filters. In: *CVPR* (2010)
5. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *CVPR* (2005)
6. Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M.: Eco: Efficient convolution operators for tracking. In: *CVPR* (2017)
7. Danelljan, M., Häger, G., Khan, F., Felsberg, M.: Accurate scale estimation for robust visual tracking. In: *BMVC* (2014)
8. Danelljan, M., Hager, G., Shahbaz Khan, F., Felsberg, M.: Learning spatially regularized correlation filters for visual tracking. In: *ICCV* (2015)
9. Danelljan, M., Robinson, A., Shahbaz Khan, F., Felsberg, M.: Beyond correlation filters: Learning continuous convolution operators for visual tracking. In: *ECCV* (2016)
10. Danelljan, M., Shahbaz Khan, F., Felsberg, M., Van de Weijer, J.: Adaptive color attributes for real-time visual tracking. In: *CVPR* (2014)
11. Galoogahi, H.K., Fagg, A., Huang, C., Ramanan, D., Lucey, S.: Need for speed: A benchmark for higher frame rate object tracking. *arXiv:1703.05884* (2017)
12. Grabner, H., Leistner, C., Bischof, H.: Semi-supervised on-line boosting for robust tracking. *ECCV* (2008)
13. Grundmann, M., Kwatra, V., Essa, I.: Auto-directed video stabilization with robust l1 optimal camera paths. In: *CVPR* (2011)
14. Held, D., Thrun, S., Savarese, S.: Learning to track at 100 fps with deep regression networks. In: *ECCV* (2016)
15. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: Exploiting the circulant structure of tracking-by-detection with kernels. In: *ECCV* (2012)
16. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. *TPAMI* **37**(3), 583–596 (2015)
17. Hong, S., You, T., Kwak, S., Han, B.: Online tracking by learning discriminative saliency map with convolutional neural network. In: *ICML* (2015)
18. Hua, Y., Alahari, K., Schmid, C.: Occlusion and motion reasoning for long-term tracking. In: *ECCV* (2014)
19. Kalal, Z., Mikolajczyk, K., Matas, J.: Tracking-learning-detection. *IEEE transactions on pattern analysis and machine intelligence* **34**(7), 1409–1422 (2012)
20. Kiani Galoogahi, H., Fagg, A., Lucey, S.: Learning background-aware correlation filters for visual tracking. In: *CVPR* (2017)
21. Kiani Galoogahi, H., Sim, T., Lucey, S.: Correlation filters with limited boundaries. In: *CVPR* (2015)
22. Kristan, M., Matas, J., Leonardis, A., Felsberg, M., Cehovin, L., Fernández, G., Vojir, T., Hager, G., Nebehay, G., Pflugfelder, R.: The visual object tracking vot2015 challenge results. In: *ICCV workshops*. pp. 1–23 (2015)

23. Kristan, M., Matas, J., Leonardis, A., et al.: The visual object tracking vot2014 challenge results. In: ECCV Workshop (2014)
24. Kumar, M., Gandhi, V., Ronfard, R., Gleicher, M.: Zooming on all actors: Automatic focus+ context split screen video generation. In: Computer Graphics Forum. vol. 36, pp. 455–465. Wiley Online Library (2017)
25. Li, H., Li, Y., Porikli, F.: Robust online visual tracking with a single convolutional neural network. In: ACCV (2014)
26. Li, Y., Zhu, J.: A scale adaptive kernel correlation filter tracker with feature integration. In: ECCV Workshops (2). pp. 254–265 (2014)
27. Liang, P., Blasch, E., Ling, H.: Encoding color information for visual tracking: Algorithms and benchmark. *IEEE Transactions on Image Processing* **24** (2015)
28. Lu, W.L., Ting, J.A., Little, J.J., Murphy, K.P.: Learning to track and identify players from broadcast sports videos. *IEEE transactions on pattern analysis and machine intelligence* **35**(7), 1704–1716 (2013)
29. Lukežič, A., Zajc, L.Č., Vojfir, T., Matas, J., Kristan, M.: Now you see me: evaluating performance in long-term visual tracking. arXiv:1804.07056 (2018)
30. Ma, C., Huang, J.B., Yang, X., Yang, M.H.: Hierarchical convolutional features for visual tracking. In: ICCV (2015)
31. Ma, C., Yang, X., Zhang, C., Yang, M.H.: Long-term correlation tracking. In: CVPR (2015)
32. Mueller, M., Smith, N., Ghanem, B.: A benchmark and simulator for uav tracking. In: ECCV (2016)
33. Nam, H., Han, B.: Learning multi-domain convolutional neural networks for visual tracking. In: CVPR (2016)
34. Smeulders, A.W., Chu, D.M., Cucchiara, R., Calderara, S., Dehghan, A., Shah, M.: Visual tracking: An experimental survey. *TPAMI* **36**(7), 1442–1468 (2014)
35. Song, Y., Ma, C., Gong, L., Zhang, J., Lau, R., Yang, M.H.: Crest: Convolutional residual learning for visual tracking. In: ICCV (2017)
36. Supancic, J.S., Ramanan, D.: Self-paced learning for long-term tracking. In: CVPR (2013)
37. Valmadre, J., Bertinetto, L., Henriques, J.F., Tao, R., Vedaldi, A., Smeulders, A., Torr, P., Gavves, E.: Long-term tracking in the wild: A benchmark. arXiv:1803.09502 (2018)
38. Vondrick, C., Patterson, D., Ramanan, D.: Efficiently scaling up crowdsourced video annotation. *IJCV* pp. 1–21 (2012)
39. Wang, N., Li, S., Gupta, A., Yeung, D.Y.: Transferring rich feature hierarchies for robust visual tracking. arXiv preprint arXiv:1501.04587 (2015)
40. Wang, N., Yeung, D.Y.: Learning a deep compact image representation for visual tracking. In: NIPS. pp. 809–817 (2013)
41. Wu, Y., Lim, J., Yang, M.H.: Online object tracking: A benchmark. In: CVPR (2013)
42. Wu, Y., Lim, J., Yang, M.H.: Object tracking benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **37**(9), 1834–1848 (2015)
43. Yun, S., Choi, J., Yoo, Y., Yun, K., Young Choi, J.: Action-decision networks for visual tracking with deep reinforcement learning. In: CVPR (2017)
44. Zhang, J., Ma, S., Sclaroff, S.: Meem: robust tracking via multiple experts using entropy minimization. In: ECCV (2014)
45. Zhang, T., Xu, C., Yang, M.H.: Multi-task correlation particle filter for robust object tracking. In: CVPR (2017)
46. Cehovin Zajc, L., Lukežic, A., Leonardis, A., Kristan, M.: Beyond standard benchmarks: Parameterizing performance evaluation in visual object tracking (2017)