

Measuring the Accuracy of Object Detectors and Trackers

Tobias Böttger^(✉), Patrick Follmann, and Michael Fauser

MVTec Software GmbH, Munich, Germany
{boettger,follmann,fauser}@mvttec.com

Abstract. The accuracy of object detectors and trackers is most commonly evaluated by the Intersection over Union (IoU) criterion. To date, most approaches are restricted to axis-aligned or oriented boxes and, as a consequence, many datasets are only labeled with boxes. Nevertheless, axis-aligned or oriented boxes cannot accurately capture an object's shape. To address this, a number of densely segmented datasets has started to emerge in both the object detection and the object tracking communities. However, evaluating the accuracy of object detectors and trackers that are restricted to boxes on densely segmented data is not straightforward. To close this gap, we introduce the relative Intersection over Union (rIoU) accuracy measure. The measure normalizes the IoU with the optimal box for the segmentation to generate an accuracy measure that ranges between 0 and 1 and allows a more precise measurement of accuracies. Furthermore, it enables an efficient and easy way to understand scenes and the strengths and weaknesses of an object detection or tracking approach. We display how the new measure can be efficiently calculated and present an easy-to-use evaluation framework. The framework is tested on the DAVIS and the VOT2016 segmentations and has been made available to the community.

1 Introduction

Visual object detection and tracking are two rapidly evolving research areas with dozens of new algorithms being published each year. To compare the performance of the many different approaches, a vast amount of evaluation datasets and schemes are available. They include large detection datasets with multiple object categories, such as PASCAL VOC [8], smaller, more specific detection datasets with a single category, such as cars [10], and sequences with multiple frames that are commonly used to evaluate trackers such as VOT2016 [11], OTB-2015 [24], or MOT16 [14]. Although very different in their nature, all of the benchmarks use axis-aligned or oriented boxes as ground truth and estimate the accuracy with the Intersection over Union (IoU) criterion.

Nevertheless, boxes are very crude approximations of many objects and may introduce an unwanted bias in the evaluation process, as is displayed in Fig. 1. Furthermore, approaches that are not restricted to oriented or axis-aligned boxes




Fig. 1. In image (a), both oriented boxes have an identical IoU with the ground truth segmentation. Nevertheless, their common IoU is only 0.71. Restricting the ground truth to boxes may introduce an undesired bias in the evaluation. In image (b), the best possible IoU of an axis-aligned box is only 0.66. Hence, for segmented data, it is difficult to use the absolute value of the IoU as an accuracy measure since it generally does not range from 0 to 1. Furthermore, although the object detection (green) in image (c) has an overlap of 0.62 with the ground truth segmentation, its IoU with the ground truth axis-aligned bounding box is only 0.45 and would be considered a false detection in the standard procedure. The proposed rIoU is the same for both boxes in (a) and 1.0 for the green boxes in (b) and (c). (Color figure online)

will not necessarily have higher accuracy scores in the benchmarks [3–5]. To address these problems, a number of densely segmented ground truth datasets has started to emerge [13, 16, 23].

Unfortunately, evaluating the accuracy of object detectors and trackers that are restricted to boxes on densely segmented data is not straightforward. For example, the VOT2016 Benchmark [11] generates plausible oriented boxes from densely segmented objects and the COCO 2014 Detection challenge [13] uses axis-aligned bounding boxes of the segmentations to simplify the evaluation protocol. Hence, approaches may have a relatively low IoU with the ground truth, although their IoU with the actual object segmentation is the same (or even better) than that of the ground truth box (see Fig. 1(c)).

To enable a fair evaluation of algorithms restricted to axis-aligned or oriented boxes on densely segmented data we introduce the relative Intersection over Union accuracy (rIoU) measure. The rIoU uses the best possible axis-aligned or oriented box of the segmentation to normalize the IoU score. The normalized IoU ranges from 0 to 1 for an arbitrary segmentation and allows to determine the true accuracy of a scheme. For tracking scenarios, the optimal boxes have further advantages. By determining three different optimal boxes for each sequence, the optimal oriented box, the optimal axis-aligned box and, the optimal axis-aligned box for a fixed scale, it is possible to identify scale changes, rotations, and occlusion in a sequence without the need of by-frame labels.

The optimal boxes are obtained in a fast and efficient optimization process. We validate the quality of the boxes in the experiments section by comparing them to a number of exhaustively determined best boxes for various scenes.

The three main contributions of this paper are:

1. The introduction of the relative Intersection over Union accuracy (rIoU) measure, which allows an accurate measurement of object detector and tracker accuracies on densely segmented data.
2. The proposed evaluation removes the bias introduced by restricting the ground truth to boxes for densely segmented data (such as COCO 2014 Detection Challenge [13] or VOT2016 [11]).
3. A compact, easy-to-use, and efficient evaluation scheme for evaluating object trackers that allows a good interpretability of a trackers strengths and weaknesses.

The proposed measure and evaluation scheme is evaluated on a handful of state-of-the-art trackers for the DAVIS [16] and VOT2016 [11] datasets and made available to the community¹.

2 Related Work

In the object detection community, the most commonly used accuracy measure is the Intersection over Union (IoU), also called Pascal overlap or bounding box overlap [8]. It is commonly used as the standard requirement for a correct detection, when the IoU between the predicted detection and the ground truth is at least 0.5 [13].

In the tracking community, many different accuracy measures have been proposed, most of them center-based and overlap-based measures [11, 12, 15, 18, 22, 24]. To unify the evaluation of trackers, Čehovin *et al.* [20, 22] provide a highly detailed theoretical and experimental analysis of the most popular performance measures and show that many of the accuracy measures are highly correlated. Nevertheless, the appealing property of the IoU measure is that it accounts for both position and size of the prediction and ground truth simultaneously. This has led to the fact that, in recent years, it has been the most commonly used accuracy measure in the tracking community [11, 24]. For example, the VOT2016 [11] evaluation framework uses the IoU as the sole accuracy measure and identifies tracker failures when the IoU between the predicted detection and the ground truth is 0.0 [12].

Since bounding boxes are very crude approximations of objects [13] and cannot accurately capture an object's shape, location, or characteristics, numerous datasets with densely segmented ground truth have emerged. For example, the COCO 2014 dataset [13] includes more than 886,000 densely annotated instances of 80 categories of objects. Nevertheless, on the COCO detection challenge the segmentations are approximated by axis-aligned bounding boxes to simplify the

¹ <http://www.mvtec.com/company/research/>.

evaluation. As stated earlier, this introduces an unwanted bias in the evaluation. A further dataset with excellent pixel accurate segmentations is the DAVIS dataset [16], which was released in 2017. It consists of 50 short sequences of manually segmented objects which, although originally for video object segmentation, can also be used for the evaluation of object trackers. Furthermore, the segmentations used to generate the VOT2016 ground truths have very recently been released [23].

In our work, we enable the evaluation of object detection and tracking algorithms that are restricted to output boxes on densely segmented ground truth data. The proposed approach is easy to add to existing evaluations and improves the precision of the standard IoU accuracy measure.

3 Relative Intersection over Union (rIoU)

Using segmentations for evaluating the accuracy of detectors or trackers removes the bias a bounding-box abstraction induces. Nevertheless, the IoU of a box and an arbitrary segmentation generally does not range from 0 to 1, where the maximum value depends strongly on the objects' shape. For example, in Fig. 1(b) the best possible axis-aligned box only has an IoU of 0.66 with the segmentation.

To enable a more precise measurement of the accuracy, we introduce the relative Intersection over Union (rIoU) of a box \mathcal{B} and a dense segmentation \mathcal{S} as

$$\Phi_{rIoU}(\mathcal{S}, \mathcal{B}) = \frac{\Phi_{IoU}(\mathcal{S}, \mathcal{B})}{\Phi_{opt}(\mathcal{S})}, \quad (1)$$

where Φ_{IoU} is the Intersection over Union (IoU),

$$\Phi_{IoU}(\mathcal{S}, \mathcal{B}) = \frac{|\mathcal{S} \cap \mathcal{B}|}{|\mathcal{S} \cup \mathcal{B}|}, \quad (2)$$

and Φ_{opt} is the best possible IoU a box can achieve for the segmentation \mathcal{S} . In comparison to the usual IoU (Φ_{IoU}), the rIoU measure (Φ_{rIoU}) truly ranges from 0 to 1 for all possible segmentations. Furthermore, the measure makes it possible to interpret ground truth attributes such as scale change or occlusion, as is displayed later in Sect. 4.

The calculation of Φ_{opt} , required to obtain Φ_{rIoU} , is described in the following section.

3.1 Optimization

An oriented box \mathcal{B} can be parameterized with 5 parameters

$$b = (r_c, c_c, w, h, \phi), \quad (3)$$

where r_c and c_c denote the row and column of the center, w and h denote the width and height, and ϕ the orientation of the box with respect to the

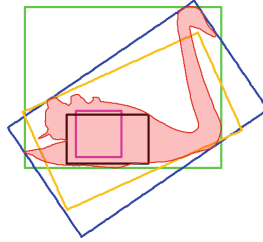


Fig. 2. blackswan from DAVIS [16]. The initial values of the optimization process of (4) are displayed. We use the axis-aligned bounding box (green), the oriented bounding box (blue), the inner square of the largest inner circle (magenta), the largest inner axis-aligned box (black) and the oriented box with the same second order moments as the segmentation (orange). (Color figure online)

column-axis. An axis-aligned box can equally be parameterized with the above parameters by fixing the orientation to 0° .

For a given segmentation \mathcal{S} , the box with the best possible IoU is

$$\Phi_{opt}(\mathcal{S}) = \max_b \Phi_{IoU}(\mathcal{S}, \mathcal{B}(b)) \quad s.t. \ b \in \mathbb{R}_{>0}^4 \times [0^\circ, 90^\circ]. \quad (4)$$

For a convex segmentation, the above problem can efficiently be optimized with the method of steepest descent. To handle arbitrary, possibly unconnected, segmentations, we optimize (4) with a multi-start gradient descent with a backtracking line search. The gradient is approximated numerically by the symmetric difference quotient. We use the diverse set of initial values for the optimization process displayed in Fig. 2. The largest axis-aligned inner box (black) and the inner box of the largest inner circle (magenta) are completely within the segmentation. Hence, in the optimization process, they will gradually grow and include background if it improves Φ_{IoU} . On the other hand, the bounding boxes (green and blue) include the complete segmentation and will gradually shrink in the optimization to include less of the segmentation. The oriented box with the same second order moments as the segmentation (orange) serves as an intermediate starting point [17]. Hence, only if the initial values converge to different optima do we need to expend more effort. In these cases, we randomly sample further initial values from the interval spanned by the obtained optima with an added perturbation. In our experiments we used 50 random samples. Although this may lead to many different optimizations, the approach is still very efficient. A single evaluation of $\Phi_{IoU}(\mathcal{S}, \mathcal{B})$ only requires an average of 0.04 ms for the segmentations within the DAVIS [16] dataset in HALCON² on an IntelCore i7-4810 CPU @2.8 GHz with 16 GB of RAM with Windows 7 (x64). As a consequence, the optimization of Φ_{opt} requires an average of 1.3 s for the DAVIS [16] and 0.7 s for the VOT2016 [11] segmentations.

The optimization of the IoU for axis-aligned rectangles bears some similarity to the 2D maximum subarray problem [1]. This might make an alternative

² MVTec Software GmbH, <https://www.mvtec.com/>.

algorithmic approach to the optimization possible. However, a straightforward adaptation of methods is difficult, since these methods rely on the additive nature of the maximum subarray problem. In contrast, the IoU is inherently non-linear due to the quotient in its definition.

3.2 Validation

To validate the optimization process, we exhaustively searched for the best boxes in a collection of exemplary frames from each of the 50 sequences in the DAVIS dataset [16]. The validation set consists of frames that were challenging for the optimization process. In a first step, we validated the optimization for axis-aligned boxes. The results in Fig. 3 indicate that the optimization is generally very close or identical to the exhaustively determined boxes.

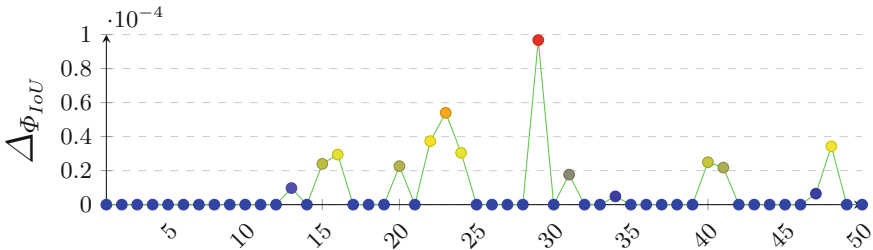


Fig. 3. The absolute difference $\Delta\phi_{IoU}$ of the exhaustively determined best axis-aligned box and the optimized axis-aligned box for a selected frame in each of the 50 DAVIS [16] sequences. Most boxes are identical, only a handful of boxes are marginally different (<0.0001).

For the oriented boxes, one of the restrictions we can make is that the area must at least be as large as the smallest inner box of the segmentation and may not be larger than the bounding oriented box. Nevertheless, even with further heuristics, the number of candidates to test is in the number of billions for the sequences in the DAVIS dataset. Given a pixel-precise discretization for r_c, c_c, w, h and a 0.5° discretization of ϕ , it was impossible to find boxes with a better IoU than the optimized oriented boxes in the validation set. This is mostly due to the fact that the sub-pixel precision of the parameterization (especially in the angle ϕ) is of paramount importance for the IoU of oriented boxes.

4 Theoretical Trackers

The concept of theoretical trackers was first introduced by Čehovin *et al.* [22] as an “*excellent interpretation guide in the graphical representation of results*”. In their paper, they use perfectly robust or accurate theoretical trackers to create bounds for the comparison of the performance of different trackers. In our case, we use the boxes with an optimal IoU to create upper bounds for the accuracy

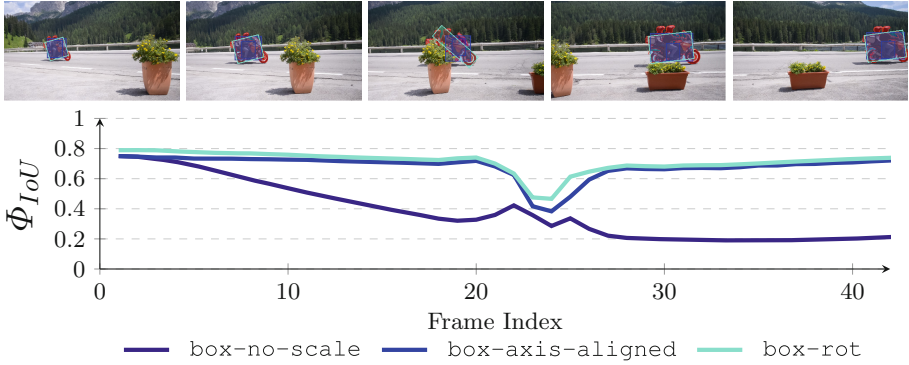


Fig. 4. motorbike from DAVIS [16]. The increasing gap between the `box-no-scale` and the other two theoretical trackers indicates a scale change of the motorbike. The drop in all three theoretical trackers around frame 25 indicates that the object is being occluded. The best possible IoU is never above 0.80 for the complete sequence.

of trackers that underlie the box-world assumption. We introduce three theoretical trackers that are obtained by optimizing (4) for a complete sequence. Given the segmentation \mathcal{S} , the first tracker returns the best possible axis-aligned box (`box-axis-aligned`), the second tracker returns the optimal oriented box (`box-rot`) and the third tracker returns the optimal axis-aligned box with a fixed scale (`box-no-scale`). The scale is initialized in the first frame with the scale of the box determined by `box-axis-aligned`.

The theoretical tracker can be used to normalize a tracker’s IoU for a complete sequence, which enables a fair interpretation of a tracker’s accuracy and removes the bias from the box-world assumption. Furthermore, the three different theoretical trackers make it possible to interpret a tracking scene without the need of by-frame labels. As is displayed in Fig. 4, the difference between

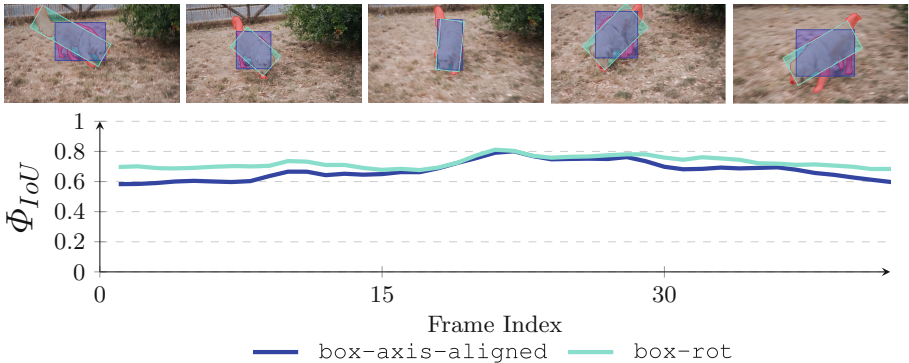


Fig. 5. dog from DAVIS [16]. The gaps between the `box-axis-aligned` and `box-rot` tracker indicate a rotation of the otherwise relatively compact segmentation of the dog. The best possible IoU is never above 0.80 for the complete sequence.

the `box-no-scale`, `box-axis-aligned`, and `box-rot` trackers indicates that the object is undergoing a scale change. Furthermore, the decreasing IoUs of all theoretical trackers indicate that the object is either being occluded or deforming to a shape that can be approximated less well by a box. For compact objects, the difference of the `box-rot` tracker and the `box-axis-aligned` tracker indicates a rotation or change of perspective, as displayed in Fig. 5.

5 Experiments

We evaluate the accuracy of a handful of state-of-the-art trackers on the DAVIS [16] and VOT2016 [11] datasets with the new rIoU measure. We initialize the trackers with the best possible axis-aligned box for the given segmentation. Since we are primarily interested in the accuracy and not in the trackers robustness, we do not reinitialize the trackers when they move off target. Please note that the accuracy of the robustness measure is also improved when using segmentations; The failure cases (hence $\Phi_{IoU} = 0$) are identified earlier since Φ_{IoU} is zero when the tracker has no overlap with the segmentation and not with a bounding box abstraction of the object (which may contain a large amount of background, see, e.g., Fig. 1).

We restrict our evaluation to the handful of (open source) state-of-the-art trackers displayed in Table 1. A thorough evaluation and comparison of all top ranking trackers is beyond the scope of this paper. The evaluation framework is made available and constructed such that it is easy to add new trackers from MATLAB³, Python⁴ or HALCON.

We include the Kernelized Correlation Filter (KCF) [9] tracker since it was a top ranked tracker in the VOT2014 challenge even though it assumes the scale of the object to stay constant. The Discriminative Scale Space Tracker (DSST) [6] tracker is essentially an extension of KCF that can handle scale changes and outperformed the KCF by a small margin in the VOT2014 challenge. As further axis-aligned trackers, we include ANT [21], L1APG [2], and the best performing tracker from the VOT2016 challenge, the continuous convolution filters (CCOT) from Danelljan *et al.* [7]. We include the LGT [19] as one of the few open source trackers that estimates the object position as an oriented box.

In Table 1, we compare the average IoU with the average rIoU for the DAVIS and the VOT2016 datasets. Please note that we normalize each tracker with the IoU of the theoretical tracker that has the same abilities. Hence, the KCF tracker is normalized with the `box-no-scale` tracker, the LGT tracker with `box-rot`, and the others with `box-axis-aligned`. By these means, it is possible to observe how well each tracker is doing with respect to its abilities. For the DAVIS dataset, the KCF, ANT, L1APG, and LGT trackers all have the same absolute IoU, but when normalized by Φ_{opt} , differences are visible. Hence, it is evident that the KCF is performing very well, given the fact that it does not estimate the scale. On the other hand, the LGT tracker, which has three more

³ The MathWorks, Inc., <https://www.mathworks.com/>.

⁴ Python Software Foundation, <https://www.python.org/>.

Table 1. Comparison of different tracking approaches and their average absolute (Φ_{IoU}) and relative IoU (Φ_{rIoU}) for the DAVIS [16] and the VOT2016 [11] segmentations

	DAVIS		VOT2016	
	Φ_{IoU}	Φ_{rIoU}	Φ_{IoU}	Φ_{rIoU}
<i>Axis-aligned boxes (fixed scale)</i>				
KCF [9]	0.40	0.78	0.23	0.45
<i>Axis-aligned boxes</i>				
DSST [6]	0.43	0.67	0.24	0.32
CCOT [7]	0.47	0.73	0.41	0.56
ANT [21]	0.40	0.64	0.26	0.37
LIAPG [2]	0.40	0.63	0.18	0.25
<i>Oriented boxes</i>				
LGT [19]	0.40	0.60	0.25	0.34

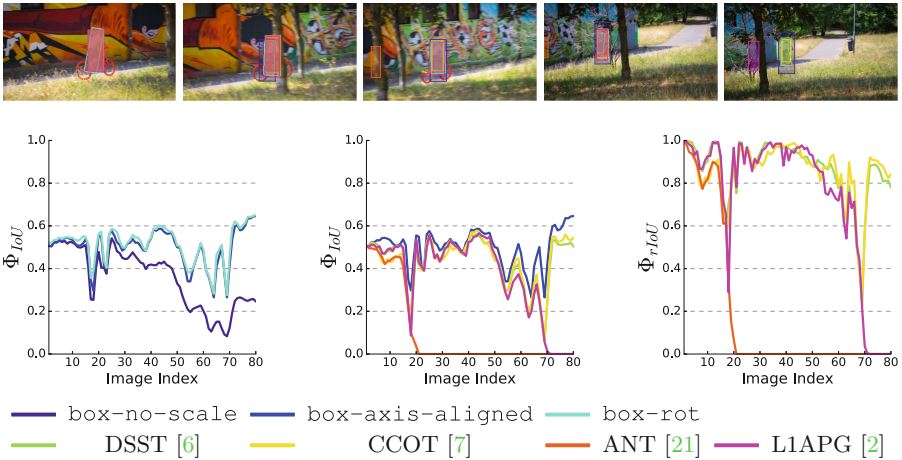


Fig. 6. *bmx-trees* from DAVIS [16]. On the left, differences between *box-no-scale* and *box-axis-aligned* indicate that the object is changing scale and is occluded at frame 18 and around frames 60–70. In the middle plot, we compare the IoU of the axis-aligned box trackers and *box-axis-aligned*. The corresponding rIoU plot is shown on the right. It becomes evident that the ANT tracker fails when the object is occluded for the first time and the LIAPG tracker at the second occlusion. The rIoU shows that DSST and CCOT perform well, while the IoU would imply they are weak.

degrees of freedom, is relatively weak. A more detailed example analysis of the *bmx-trees* sequence from DAVIS [16] is displayed in Fig. 6. Please note, the significantly higher difference between the IoU and the rIoU for KCF compared to the other trackers is due to the different normalization factors used in the rIoU measure. The optimal IoU value for a box with fixed size is usually considerably lower than for a general axis-aligned box.

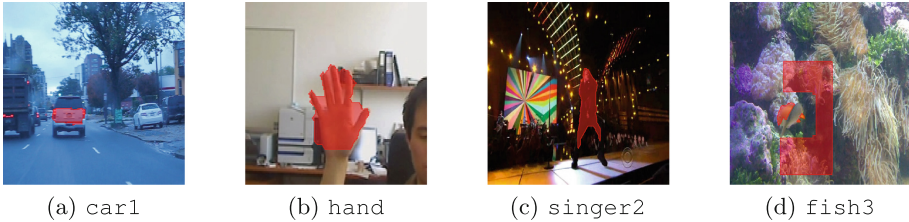


Fig. 7. Examples from VOT2016 [11] where the segmentations are degenerated. Sometimes due to motion blur (e.g., (a) and (b)) or a weak contrast of the object and its background (c). (Color figure online)

For the VOT2016 dataset, the overall accuracies are significantly worse than for DAVIS. On the one hand, this is due to the longer, more difficult sequences, and, on the other hand, due to the less accurate and noisier segmentations (see Fig. 7). Nevertheless, the rIoU allows a more reliable comparison of different trackers. For example, ANT, LGT and DSST have almost equal average IoU value, while ANT clearly outperforms LGT and DSST with respect to rIoU. Again, we can see that the KCF tracker is quite strong regarding the fact that it cannot estimate the scale.

6 Conclusion

In this paper, we have proposed a new accuracy measure that closes the gap between densely segmented ground truths and box detectors and trackers. We have presented an efficient optimization scheme to obtain the best possible detection boxes for arbitrary segmentations that are required for the new measure. The optimization was validated on a diverse set of segmentations from the DAVIS dataset [16]. The new accuracy measure can be used to generate three very expressive theoretical trackers, which can be used to obtain meaningful accuracies and help to interpret scenes without requiring by-frame labels. We have evaluated state-of-the-art trackers with the new accuracy measure on all segmentations within the DAVIS [16] and VOT2016 [11] datasets to display its advantages. The complete code and evaluation system will be made available to the community to encourage its use and make it easy to reproduce our results.

References

1. An, S., Peursum, P., Liu, W., Venkatesh, S.: Efficient algorithms for subwindow search in object detection and localization. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 264–271, June 2009
2. Bao, C., Yi, W., Ling, H., Ji, H.: Real time robust L1 tracker using accelerated proximal gradient approach. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 1830–1837 (2012)

3. Böttger, T., Eisenhofer, C.: Efficiently tracking extremal regions in multichannel images. In: International Conference on Pattern Recognition Systems (ICPRS) (2017)
4. Böttger, T., Ulrich, M., Steger, C.: Subpixel-precise tracking of rigid objects in real-time. In: Sharma, P., Bianchi, F.M. (eds.) SCIA 2017. LNCS, vol. 10269, pp. 54–65. Springer, Cham (2017). doi:[10.1007/978-3-319-59126-1_5](https://doi.org/10.1007/978-3-319-59126-1_5)
5. Caelles, S., Maninis, K.-K., Pont-Tuset, J., Leal-Taixé, L., Cremers, D., Van Gool, L.: One-shot video object segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (2017)
6. Danelljan, M., Häger, G., Khan, F.S., Felsberg, M.: Accurate scale estimation for robust visual tracking. In: British Machine Vision Conference (2014)
7. Danelljan, M., Robinson, A., Shahbaz Khan, F., Felsberg, M.: Beyond correlation filters: learning continuous convolution operators for visual tracking. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9909, pp. 472–488. Springer, Cham (2016). doi:[10.1007/978-3-319-46454-1_29](https://doi.org/10.1007/978-3-319-46454-1_29)
8. Everingham, M., Ali Eslami, S.M., Van Gool, L.J., Williams, C.K.I., Winn, J.M., Zisserman, A.: The pascal visual object classes challenge: a retrospective. *Int. J. Comput. Vis.* **111**(1), 98–136 (2015)
9. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(3), 583–596 (2015)
10. Juránek, R., Herout, A., Dubská, M., Zemčík, P.: Real-time pose estimation piggybacked on object detection. In: IEEE International Conference on Computer Vision, pp. 2381–2389 (2015)
11. Kristan, M., et al.: The visual object tracking VOT2016 challenge results. In: Hua, G., Jégou, H. (eds.) ECCV 2016. LNCS, vol. 9914, pp. 777–823. Springer, Cham (2016). doi:[10.1007/978-3-319-48881-3_54](https://doi.org/10.1007/978-3-319-48881-3_54)
12. Kristan, M., Matas, J., Leonardis, A., Vojtír, T., Pflugfelder, R.P., Fernández, G., Nebel, G., Porikli, F., Čehovin, L.: A novel performance evaluation methodology for single-target trackers. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(11), 2137–2155 (2016)
13. Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). doi:[10.1007/978-3-319-10602-1_48](https://doi.org/10.1007/978-3-319-10602-1_48)
14. Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K.: MOT16: a benchmark for multi-object tracking. [arXiv:1603.00831](https://arxiv.org/abs/1603.00831) [cs], March 2016
15. Nawaz, T., Cavallaro, A.: A protocol for evaluating video trackers under real-world conditions. *IEEE Trans. Image Process.* **22**(4), 1354–1361 (2013)
16. Perazzi, F., Jordi Pont-Tuset, B., McWilliams, L.J., Gool, V., Gross, M.H., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 724–732 (2016)
17. Rosin, P.L.: Measuring rectangularity. *Mach. Vis. Appl.* **11**(4), 191–196 (1999)
18. Smeulders, A.W.M., Chu, D.M., Cucchiara, R., Calderara, S., Dehghan, A., Shah, M.: Visual tracking: an experimental survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(7), 1442–1468 (2014)
19. Čehovin, L., Kristan, M., Leonardis, A.: Robust visual tracking using an adaptive coupled-layer visual model. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(4), 941–953 (2013)

20. Čehovin, L., Kristan, M., Leonardis, A.: Is my new tracker really better than yours? In: IEEE Winter Conference on Applications of Computer Vision, pp. 540–547 (2014)
21. Čehovin, L., Leonardis, A., Kristan, M.: Robust visual tracking using template anchors. In: IEEE Winter Conference on Applications of Computer Vision, pp. 1–8 (2016)
22. Čehovin, L., Leonardis, A., Kristan, M.: Visual object tracking performance measures revisited. *IEEE Trans. Image Process.* **25**(3), 1261–1274 (2016)
23. Vojir, T., Matas, J.: Pixel-wise object segmentations for the VOT 2016 dataset. Research report CTU-CMP-2017-01, Center for Machine Perception, Czech Technical University, Prague, Czech Republic, January 2017
24. Yi, W., Lim, J., Yang, M.-H.: Object tracking benchmark. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(9), 1834–1848 (2015)