DCCO: Towards Deformable Continuous Convolution Operators for Visual Tracking

Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan and Michael Felsberg

The self-archived postprint version of this conference article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-145373>

N.B.: When citing this work, cite the original publication.

The original publication is available at www.springerlink.com:

Johnander, J., Danelljan, M., Khan, F. S., Felsberg, M., (2017), DCCO: Towards Deformable Continuous Convolution Operators for Visual Tracking, *Computer Analysis of Images and Patterns*, , 55-67. https://doi.org/10.1007/978-3-319-64689-3_5

Original publication available at:

https://doi.org/10.1007/978-3-319-64689-3_5

Copyright: Springer Verlag (Germany)

<http://www.springerlink.com/?MUD=MP>

DCCO: Towards Deformable Continuous Convolution Operators for Visual Tracking

Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan, Michael Felsberg

Computer Vision Laboratory, Dept. of Electrical Engineering, Linköping University

Abstract. Discriminative Correlation Filter (DCF) based methods have shown competitive performance on tracking benchmarks in recent years. Generally, DCF based trackers learn a rigid appearance model of the target. However, this reliance on a single rigid appearance model is insufficient in situations where the target undergoes non-rigid transformations. In this paper, we propose a unified formulation for learning a deformable convolution filter. In our framework, the deformable filter is represented as a linear combination of sub-filters. Both the sub-filter coefficients and their relative locations are inferred jointly in our formulation. Experiments are performed on three challenging tracking benchmarks: OTB-2015, TempleColor and VOT2016. Our approach improves the baseline method, leading to performance comparable to state-of-the-art.

Keywords: Visual tracking

1 Introduction

Generic visual object tracking is the computer vision problem of estimating the trajectory of a target throughout an image sequence, given only the initial target location. Visual tracking is useful in numerous applications, including autonomous driving, smart surveillance systems and intelligent robotics. The problem is challenging due to large variations in appearance of the target and background, as well as challenging situations involving motion blur, target deformation, in- and out-of-plane rotations, and fast motion.

To tackle the problem of visual tracking, several paradigms exist in literature [13]. Among different paradigms, approaches based on the Discriminative Correlation Filters (DCF) based framework have achieved superior results, evident from recent the Visual Object Tracking (VOT) challenge results [14][13]. This improvement in performance, both in terms of precision and robustness, is largely attributed to the use of powerful multi-dimensional features such as HOG, Colornames, and deep features [5][20][10], as well as sophisticated learning models [8][9].

Despite the improvement in tracking performance, the aforementioned state-of-the-art DCF based approaches employ a single rigid model of the target. However, this reliance on a single rigid model is insufficient in situations involving rotations and deformable targets. In such complex situations, the rigid filters fail to capture information of the target parts that move relative to each other.

Fig. 1. Example tracking results of our deformable correlation filter approach on three challenging sequences. The circles mark sub-filter locations and the green box is the predicted target location. The red boxes show the baseline predictions. The sub-filter locations deform according to the appearance changes of the target in the presence of deformations.

This desired information can be retained by integrating deformability in the DCF filters. Several recent works aim at introducing part-based information into the DCF framework [18][16][19]. These approaches introduce an explicit component to integrate the part-based information in the learning. Different to these approaches, we investigate a deformable DCF model, which can be learned in unified fashion.

In many real-world situations, such as a running human or a rotating box, different regions of the target deform relative to each other. Ideally, such information should be integrated in the learning formulation by allowing the regions of the appearance model to deform accordingly. This flexibility in the tracking model reduces the need of highly invariant features, thereby increasing the discriminative power of the model. However, increasing the flexibility and complexity of the model introduces the risk of over-fitting and complex inference mechanisms, which degrades the robustness of the tracker. In this paper, we therefore advocate a unified formulation, where the deformable filter is learned by optimizing a single joint objective function. Additionally, this unified strategy enables the careful incorporation of regularization models to tackle the risk of over-fitting.

Contribution We propose a unified framework for learning a deformable convolution filter in a discriminative fashion. The deformable filter is represented as a linear combination of sub-filters. The deformable filter is learned by jointly optimizing the sub-filter coefficients *and* their relative locations. To avoid over-fitting, we propose to regularize the sub-filter locations with an affine deformation model. We further derive an efficient online optimization procedure to infer the parameters of the model. Experiments on three challenging tracking benchmarks suggest that our method improves the performance in challenging situations.

2 Related Work

In recent years, Discriminative Correlation Filters (DCF) based tracking methods have shown competitive performance in terms of accuracy and robustness on tracking benchmarks [13][22]. In particular, the success of DCF based methods is evident from the outcome of the Visual Object Tracking (VOT) 2014 and 2016 challenges [13] where the top-rank trackers employ variants of the DCF framework. In DCF framework, a correlation filter is learned from a set of training samples to discriminate between the target and background appearance. The training of the filter is performed in a sliding-window manner by exploiting the properties of circular correlation. The original DCF based tracking approach by Bolme et al. [3] was restricted to a single feature channel and was later extended to multi-channel feature maps [11][10][12]. Most recent advancement in DCF based tracking performance is attributed to including scale estimation [6][15], deep features [7][20], spatial regularization [8], and continuous convolution filters [9].

Several recent works have shown that integrating the part-based information improve the tracking performance. The work of [18] introduces a part-based approach where each part utilizes the kernelized correlation filter (KCF) tracker and argues that partial occlusions can effectively be handled by adaptive weighting of the parts. The work of [16] tracks several patches, each with a KCF, by fusing the information using a particle filter to estimate position, width and height. Lukezic et. al. [19] introduces a sophisticated model with several parts held together by a spring-like system by minimizing an energy function based on the part-filter responses.

Our approach: Different to aforementioned approaches, we propose a theoretical framework by designing a single deformable correlation filter. In our approach, the coefficients and locations of all sub-filters are learned jointly in a *unified* framework. Additionally, we integrate our deformable correlation filter in a recently introduced state-of-the-art DCF tracking framework [9].

3 Continuous Convolution Operators for Tracking

In this work, we propose a deformable correlation tracking formulation. As a starting point, we use the recent Continuous Convolution Operator Tracker (C-

COT) formulation [9] due to two main advantages compared to current template based correlation filter trackers. Firstly, the continuous reformulation of the learning problem benefits from a natural integration of multi-resolution deep features and continuous-domain score map predictions. Secondly, it provides an efficient optimization framework based on the Conjugate Gradient method. For efficiency, we also employ components of its descendant tracker *ECO* [4].

For a given target object in a video, the C-COT discriminatively learns a convolution filter f that acts as an instance-specific object detector. Different from previous approaches, the filter f is viewed as a continuous function represented by its Fourier series coefficients. The detection scores are computed by first extracting a D -dimensional feature map x from the local image region of interest. Typically, the sample x consists of HOG or multi-resolution deep convolutional features. We let $x_d[n_1, n_2]$ denote the value of the d -th feature channel at the spatial location (n_1, n_2) in the feature map. The continuous scores in the corresponding image region are determined by the convolution operation $S_f\{x\} = \sum_{d=1}^D f_d * J^d\{x^d\}$, where $J^d\{x^d\}$ is an interpolation operator mapping the samples from the discrete to the continuous domain.

The filter f is trained in a supervised fashion, given a set of sample feature maps $\{x^1, x^2, \dots, x^C\}$ and corresponding label score maps $\{y^1, y^2, \dots, y^C\}$, by minimizing the objective,

$$\epsilon(f) = \sum_{c=1}^C \alpha^c \|S_f\{x^c\} - y^c\|^2 + \sum_{d=1}^D \|w^d \cdot f_d\|^2. \quad (1)$$

The first term penalizes classification errors of each sample using the squared L^2 -norm. The sample c is weighted by the positive weight factor α^c , which is typically set using a learning rate parameter. The second term deploys a continuous spatial regularization function w^d , that penalizes high magnitude filter coefficients to alleviate the periodic boundary effects. Element-wise multiplication is denoted as $\cdot$. The label score function y^c is generally set to a Gaussian function with a narrow peak at the target center. Note that a sample feature map x^c contains both target appearance and the surrounding background. The filter is hence trained to predict high activation scores at the target center and low scores at the neighboring background. In practice, training and detection is performed directly in the Fourier domain, utilizing the FFT algorithm and the convolution properties of the Fourier series.

As related methods, the C-COT method works in two main steps. (i) When a new sample is received, the target position and scale are estimated, i.e. $S_f\{x\}$ is calculated using the estimated filter f for different scales using a scale pyramid. The new target state is then estimated as the position and scale that maximizes the detection score. (ii) To update the model, a sample (x^c, y^c) is first added to the training set, where x^c is extracted in the estimated target scale. The filter is then refined by minimizing the objective (1). This is done by using conjugate gradient to solve the arising normal equations. We refer to [9] for further details. To enhance the efficiency of the tracker, we further deploy the factorized convolution approach and update strategy recently proposed in [4].

4 Method

Here, we introduce a deformable correlation filter tracking model. A classic DCF contains an assumption that the target is rigid and will not rotate. The filter can handle violations to this assumption if a significant part of the target still fulfills it, or by using features with sufficient invariance. Examples of such model violations are sequences showing humans running or a change of perspective. By dividing the filter into sub-filters which can move relative to each other, they can fit more accurately onto a smaller part of the target. A standard DCF may choose to discard or weigh down information about a moving part whereas our approach allows one sub-filter to focus on this information explicitly, and move with that part. By writing the filter as a linear combination of sub-filters we can optimize a joint loss over all the sub-filter coefficients and the sub-filter positions jointly.

4.1 Deformable Correlation Filter

We construct a deformable convolution filter as a linear combination of trainable sub-filters. The filter becomes deformable by allowing the relative locations of the filters to change along to the target transformations. Formally, we denote the sub-filter with f^m and let $p^{c,m} = (p_1^{c,m}, p_2^{c,m})$ be its relative location in the frame c . The filter f at frame c is obtained as a linear combination of the shifted sub-filters,

$$f(t_1, t_2) = \sum_{m=1}^M f^m(t_1 - p_1^{c,m}, t_2 - p_2^{c,m}). \quad (2)$$

We jointly learn both the sub-filter coefficients f^m and their locations $p^{c,m}$ by minimizing a joint loss.,

$$\epsilon(f, p) = \epsilon_1(f, p) + \epsilon_2(f) + \epsilon_3(p), \quad (3)$$

where each term is described below.

Classification Error The loss for the discrepancy between the desired response and the filter response for sample x^c is

$$\epsilon_1(f, p) = \sum_{c=1}^C \alpha^c \|S_f\{x^c\} - y^c\|^2, \quad (4)$$

where α^c is the weight for sample c . From the translation invariance of the convolution operation and the definition (2), the classification scores can be computed as,

$$S_f\{x^c\}(t_1, t_2) = \sum_{m=1}^M S_{f^m}\{x^c\}(t_1 - p_1^{c,m}, t_2 - p_2^{c,m}). \quad (5)$$

The score operator $S_{f^m}\{x^c\}$ is defined as described in section 3.

Spatial Regularization A spatial regularization of the filters enforces low filter coefficients close to the edges,

$$\epsilon_2(f) = \sum_{m=1}^M \sum_{d=1}^D \|w^{m,d} \cdot f_d^m\|^2, \quad (6)$$

where $w^{m,d}$ is the continuous spatial regularization function for filter m . We assume different spatial regularization functions for the different sub-filters as it may be desirable for the sub-filters to track regions of different size. In our experiments, by using two different spatial regularizations where one is much tighter, we let one sub-filter track the whole target while the others track smaller patches. Please note that $\epsilon_2(f)$ does not depend on the sub-filter positions.

Regularization of Sub-filter Positions To regularize the sub-filter positions, we add a deformable model that incorporates prior information of typical target deformations. In this work, we use a simple yet effective model, namely that the current sub-filter positions are related to their initial positions by a linear mapping. The resulting regularization term is thus given by,

$$\epsilon_3(p) = \lambda_p \sum_{m=1}^M \|p^{c,m} - Rp^{1,m}\|^2. \quad (7)$$

Here, $p^{c,m}$ is the position of sub-filter m in frame c , and $R \in \mathbb{R}^{2 \times 2}$ is a transformation matrix. In our experiments we use a full linear transform, which is optimized jointly during the learning. λ_p is a parameter determining the regularization impact. This part of the loss does not depend on the sub-filter coefficients.

4.2 Fourier Domain Formulation

The optimization is performed in the Fourier domain using Parseval's formula. This results in a finite representation of the continuous filters using truncated Fourier series.

Let $\hat{\cdot}$ denote the Fourier coefficients for any given, sufficiently nice function. By linearity of the Fourier transform

$$\widehat{S_f\{x^c\}}[k_1, k_2] = \sum_{m=1}^M \beta[k_1, k_2] \widehat{S_{f^m}\{x^c\}}[k_1, k_2] \quad (8)$$

where

$$\beta[k_1, k_2] = e^{-i2\pi p_1^{c,m} k_1 / T_1} e^{-i2\pi p_2^{c,m} k_2 / T_2} \quad (9)$$

and

$$\widehat{S_{f^m}\{x^c\}}[k_1, k_2] = \left(\sum_{d=1}^D \hat{f}_d^m[k_1, k_2] \widehat{J^d\{x^c\}}[k_1, k_2] \right). \quad (10)$$

Given C samples, we optimize the filter in the C-COT framework. The objective 3 is minimized by using Parseval's formula. We get the corresponding objective

$$\epsilon(f, p) = \sum_{c=1}^C \alpha^c \|\widehat{S_f\{x^c\}} - \hat{y}^c\|^2 + \sum_{m=1}^M \sum_{d=1}^D \|\hat{w}^{m,d} * \hat{f}_d^m\|^2 + \lambda_p \sum_{m=1}^M \|p^{c,m} - Rp^{1,m}\|^2 \quad (11)$$

which will be minimized by an alternate optimization strategy where we iteratively update the sub-filter coefficients and positions.

4.3 Updating the Filter Coefficients

The Fourier coefficients are truncated such that for feature dimension d only the K^d first coefficients are used (resulting in $2K^d + 1$ coefficients in total for that dimension). Also define $K = \max_d K^d$. To minimize the functional we rewrite it as a least squares problem which can be solved via its normal equations. The normal equations are then solved using conjugate gradient. Let $\cdot^H$ be the conjugate transpose. We define a block matrix with $C \times MD$ blocks

$$A = \begin{pmatrix} A^1 \\ \vdots \\ A^C \end{pmatrix}, \quad A^c = (A^{c,1} \dots A^{c,M}), \quad A^{c,m} = (A^{c,m,1} \dots A^{c,m,D}) \quad (12)$$

where $A^{c,m,d}$ is a diagonal matrix of size $K \cdot K \times K^d \cdot K^d$

$$A^{c,m,d} = \text{diag} \begin{pmatrix} \beta[-K^d, -K^d] \widehat{J^d\{x^c\}}[-K^d, -K^d] \\ \vdots \\ \beta[-K^d, K^d] \widehat{J^d\{x^c\}}[-K^d, K^d] \\ \vdots \\ \beta[K^d, K^d] \widehat{J^d\{x^c\}}[K^d, K^d] \end{pmatrix}. \quad (13)$$

Further define

$$\hat{\mathbf{f}} = \begin{pmatrix} \hat{\mathbf{f}}^1 \\ \vdots \\ \hat{\mathbf{f}}^M \end{pmatrix}, \quad \hat{\mathbf{f}}^m = \begin{pmatrix} \hat{\mathbf{f}}_1^m \\ \vdots \\ \hat{\mathbf{f}}_D^m \end{pmatrix}, \quad \hat{\mathbf{f}}_d^m = \begin{pmatrix} f_d^m[-K^d, -K^d] \\ \vdots \\ f_d^m[-K^d, K^d] \\ \vdots \\ f_d^m[K^d, K^d] \end{pmatrix} \quad (14)$$

and

$$\hat{\mathbf{y}} = \begin{pmatrix} \hat{\mathbf{y}}^1 \\ \vdots \\ \hat{\mathbf{y}}^C \end{pmatrix}. \quad (15)$$

Lastly, let Γ denote a diagonal matrix containing the learning rate α^c , of size $CK \times CK$; and W denote a Toeplitz matrix corresponding to summation of the convolutions with $w^{m,d}$. Using these definitions the objective becomes

$$\epsilon(f, p) = \sum_{c=1}^C \alpha^c \|A^c \hat{\mathbf{f}} - \hat{\mathbf{y}}^c\|^2 + \|W \hat{\mathbf{f}}\|^2 + \epsilon_3(p). \quad (16)$$

We discard $\epsilon_3(p)$ while minimizing the objective over f , as it will be addressed in the next step. The objective is then minimized by solving

$$(A^H \Gamma A + W^H W) \hat{\mathbf{f}} = A^H \hat{\mathbf{y}} \quad (17)$$

using the method of conjugate gradient.

4.4 Displacement Estimation of the Sub-Filters

The sub-filters are moved by minimizing the objective with respect to the sub-filter positions. This problem is not convex, and we resort to gradient descent utilizing Barzilai-Borwein's method [1]. The perk of their method is that the steplength is adaptive. The gradient is found as

$$\frac{d}{dp^{c,m}} \epsilon(f) = \frac{d}{dp^{c,m}} \epsilon_1(f) + \frac{d}{dp^{c,m}} \epsilon_3(p) \quad (18)$$

where

$$\frac{d}{dp^{c,m}} \epsilon_1(f) = 2(\widehat{S_f\{x^c\}} - \hat{y}^c) e^{-i2\pi p_1^{c,m} k_1/T_1} e^{-i2\pi p_2^{c,m} k_2/T_2} \widehat{S_{f^m}\{x^c\}} \begin{pmatrix} -i2\pi k_1/T \\ -i2\pi k_2/T \end{pmatrix} \quad (19)$$

and

$$\frac{d}{dp^{c,m}} \epsilon_3(p) = 2\lambda_p(p^{c,m} - Rp^{1,m}). \quad (20)$$

Note that $\epsilon_2(f)$ does not depend on the sub-filter positions, and hence the derivative with respect to the sub-filter positions is zero. In our experiments we let R be either the identity matrix, or an affine transform. The translation part of the affine transform is handled during the target position estimation described in section 3. Hence the affine transform can be considered equivalent to a linear transform. The linear transform is estimated in each step of gradient descent using a closed form expression. This is done by rewriting the problem as an over-determined linear system of equations and solve it via its normal equations.

5 Experiment and Results

We validate our approach by performing comprehensive experiments on three tracking benchmarks: OTB-2015 [22], TempleColor [17] and VOT2016 [13].

Table 1. Baseline comparison on the OTB-2015 dataset with the two different regularizations of the sub-filter positions. The affine transform provides the best results.

	Baseline, no deformability	Affine	Identity
Mean OP	83.2	83.9	<i>83.4</i>
Mean AUC	68.4	69	<i>68.5</i>

Table 2. Baseline comparison on the OTB-2015 dataset when using different set of features for the sub-filters.

	Baseline	Shallow + CN	Shallow	Shallow + Deep	Deep	CN
Mean OP	83.2	83.6	83.5	83.6	<i>83.9</i>	83.9
Mean AUC	68.4	69	<i>68.9</i>	<i>68.9</i>	69	68.8

5.1 Implementation Details

In our experiments we employ two types of features: Color Names, and “Deep Features” extracted from the Convolutional Neural Network (CNN). We use the network VGG-m and extract features from the layers Conv-1 and Conv-5. We use different number of sub-filters depending on the target size. We employ a “root-filter” which is a subfilter that is always centered around the target and utilizes both shallow features and deep features from a CNN. The locations of the sub-filters are continuously updated and has a strong regularization to enforce locality. We test different feature sets for these sub-filters. The sub-filters are initialized in the first frame where they are placed in a grid. We use $\lambda_P = 3 \cdot 10^{-6}$ on VOT2016 and TempleColor datasets, and use $\lambda_P = 3 \cdot 10^{-4}$ on the OTB-2015 dataset. We use the same set of parameters for all videos in each dataset.

5.2 Baseline Comparison

We perform baseline comparisons on the OTB-2015 dataset with 100 videos. We compare different features for the sub-filters, and different regularization for their positions. We evaluate the tracking performance in terms of mean overlap precision (OP) and area-under-the-curve (AUC). The overlap precision (OP) is calculated as the fraction of frames in the video where the intersection-over-union (IoU) overlap with the ground truth exceeds a threshold of 0.5 (PASCAL criterion). The area-under-the-curve (AUC) is calculated from the success plot where the mean OP is plotted over the range of IoU thresholds over all videos.

Table 1 shows the results of the baseline and proposed approach with the sub-filter positions regularized either with an affine transform, or the identity transform (Sec. 4.4). The proposed approach based on an affine transform provides improved tracking performance. This shows that regularization of the sub-filter positions is important and using an affine transform is superior compared to an identity transform. Table 2 shows the baseline comparison when using different set of features. The deep features provide improved performance. However, performance comparable to deep features is also achieved by using colornames.

Fig. 2. Success plots on the OTB-2015 (left) and TempleColor (right) datasets, compared to state-of-the-art. The AUC score of each tracker is shown in the legend. We show slight performance increases on both datasets.

Fig. 3. Attribute-based comparison on the OTB-2015 dataset. Success plots are shown for six attributes. Our approach achieves improved performance compared to existing trackers in these scenarios.

5.3 State-of-the-art Comparison

OTB-2015 Figure 2 (on the left) shows the success plot for the OTB-2015 dataset which consists of 100 videos. The area-under-the-curve (AUC) score for each tracker is represented in the legend. Among existing approaches, the C-COT tracker [9] achieves an AUC score of 68.2%. It is worth to mention that the recently introduced ECO tracker [4] achieves the best results with an AUC score of 70.0%. However, the ECO tracker also employs HOG features together

Table 3. State-of-the-art in terms of expected area overlap (EAO), robustness (failure rate), and accuracy on the VOT2016 dataset. The proposed approach show a slight decrease in EAO but a slight improvement to failure rate.

	SRBT [13]	EBT [23]	DDC [13]	Staple [2]	MLDF [13]	SSAT [13]	TCNN [21]	C-COT [9]	ECO [4]	Proposed Our
EAO	0.290	0.291	0.293	0.295	0.311	0.321	0.325	0.331	0.374	<i>0.368</i>
Fail. rt.	1.25	0.90	1.23	1.35	0.83	1.04	0.96	0.85	<i>0.72</i>	0.70
Acc.	0.50	0.44	0.53	<i>0.54</i>	0.48	0.57	<i>0.54</i>	0.52	<i>0.54</i>	<i>0.54</i>

with colormnames (CN) and deep features. Instead, our deformable convolution filter approach achieves competitive performance without using HOG features, with an AUC score of 69.0%. Figure 3 shows the attribute based comparison on the OTB-2015 dataset. All videos in the OTB-2015 dataset are annotated with 11 different attributes. Our approach provides the best results on 7 attributes.

5.4 TempleColor

Figure 2 (on the right) shows the success plot for the TempleColor dataset consisting of 128 videos. The SRDCF tracker [8] and its deep features variant (Deep-SRDCF) [7] achieve AUC scores of 51.6% and 54.3% respectively. The C-COT tracker yields an AUC score of 58.1%. Our approach improves the performance by 1.4% compared to the C-COT tracker.

5.5 VOT2016

The VOT2016 which consists of 60 videos compiled from a set of more than 300 videos. On the VOT2016 dataset, the tracking performance is evaluated both in terms of accuracy (average overlap during successful tracking) and robustness (failure rate). The overall tracking performance is calculated using Expected Average Overlap (EAO) which takes into account both accuracy and robustness. For more details, we refer to [14]. Table 3 shows the comparison on the VOT2016 dataset. We present the results in terms of EAO, failure rate, and accuracy. Our approach provides competitive performance in terms of accuracy and provides the best results in terms of robustness, with a failure rate of 0.70.

6 Conclusions

We proposed a unified formulation to learn a deformable convolution filter. We represented our deformable filter as a linear combination of sub-filters. Both the coefficients and locations of all sub-filters are learned jointly in our framework. Experiments are performed on three challenging tracking datasets: OTB-2015, TempleColor and VOT2016. Our results clearly suggest that the proposed deformable convolution filter provides improved results compared to the baseline, leading to competitive performance compared to state-of-the-art trackers.

References

1. Barzilai, J., Borwein, J.M.: Two-point step size gradient methods. *IMA journal of numerical analysis* 8(1), 141–148 (1988)
2. Bertinetto, L., Valmadre, J., Golodetz, S., Miksik, O., Torr, P.H.S.: Staple: Complementary learners for real-time tracking. In: *CVPR* (2016)
3. Bolme, D.S., Beveridge, J.R., Draper, B.A., Lui, Y.M.: Visual object tracking using adaptive correlation filters. In: *CVPR* (2010)
4. Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M.: Eco: Efficient convolution operators for tracking. In: *CVPR* (2017)
5. Danelljan, M., Häger, G., Khan, F., Felsberg, M.: Adaptive decontamination of the training set: A unified formulation for discriminative visual tracking. In: *CVPR* (2016)
6. Danelljan, M., Häger, G., Shahbaz Khan, F., Felsberg, M.: Accurate scale estimation for robust visual tracking. In: *BMVC* (2014)
7. Danelljan, M., Häger, G., Shahbaz Khan, F., Felsberg, M.: Convolutional features for correlation filter based visual tracking. In: *ICCV Workshop* (2015)
8. Danelljan, M., Häger, G., Shahbaz Khan, F., Felsberg, M.: Learning spatially regularized correlation filters for visual tracking. In: *ICCV* (2015)
9. Danelljan, M., Robinson, A., 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., Sim, T., Lucey, S.: Multi-channel correlation filters. In: *ICCV* (2013)
12. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. *TPAMI* 37(3), 583–596 (2015)
13. Kristan, M., Leonardis, A., Matas, J., Felsberg, M., Pflugfelder, R., Čehovin, L., Vojír, T. and Häger, G., et al. The visual object tracking vot2016 challenge results. In: *ECCV workshop* (2016)
14. Kristan, M., Matas, J., Leonardis, A., Felsberg, M., Čehovin, L., Fernández, G., Vojír, T., Nebehay, G., Pflugfelder, R., Häger, G.: The visual object tracking vot2015 challenge results. In: *ICCV workshop* (2015)
15. Li, Y., Zhu, J.: A scale adaptive kernel correlation filter tracker with feature integration. In: *ECCV Workshop* (2014)
16. Li, Y., Zhu, J., Hoi, S.C.: Reliable patch trackers: Robust visual tracking by exploiting reliable patches. In: *CVPR* (2015)
17. Liang, P., Blasch, E., Ling, H.: Encoding color information for visual tracking: Algorithms and benchmark. *TIP* 24(12), 5630–5644 (2015)
18. Liu, T., Wang, G., Yang, Q.: Real-time part-based visual tracking via adaptive correlation filters. In: *CVPR* (2015)
19. Lukežič, A., Čehovin, L., Kristan, M.: Deformable parts correlation filters for robust visual tracking. *arXiv preprint arXiv:1605.03720* (2016)
20. Ma, C., Huang, J.B., Yang, X., Yang, M.H.: Hierarchical convolutional features for visual tracking. In: *ICCV* (2015)
21. Nam, H., Baek, M., Han, B.: Modeling and propagating cnns in a tree structure for visual tracking. *CoRR abs/1608.07242* (2016)
22. Wu, Y., Lim, J., Yang, M.H.: Object tracking benchmark. *TPAMI* 37(9), 1834–1848 (2015)
23. Zhu, G., Porikli, F., Li, H.: Tracking randomly moving objects on edge box proposals. *arXiv preprint arXiv:1507.08085* (2015)