CyF

This ACCV 2020 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2020 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

Adversarial Semi-Supervised Multi-Domain Tracking

Kourosh Meshgi¹ and Maryam Sadat Mirzaei¹

RIKEN Center for Advanced Intelligence Project (AIP), Tokyo, Japan {kourosh.meshgi, maryam.mirzaei}@riken.jp

Abstract. Neural networks for multi-domain learning empowers an effective combination of information from different domains by sharing and co-learning the parameters. In visual tracking, the emerging features in shared layers of a multi-domain tracker, trained on various sequences, are crucial for tracking in unseen videos. Yet, in a fully shared architecture, some of the emerging features are useful only in a specific domain, reducing the generalization of the learned feature representation. We propose a semi-supervised learning scheme to separate domain-invariant and domain-specific features using adversarial learning, to encourage mutual exclusion between them, and to leverage self-supervised learning for enhancing the shared features using the unlabeled reservoir. By employing these features and training dedicated layers for each sequence, we build a tracker that performs exceptionally on different types of videos. ...

1 Introduction

Multi-task learning (MTL) is a branch of supervised learning that strives to improve the generalization of the regression or classification task by leveraging the domain-specific information contained in the training signals of related tasks [1]. MTL has been investigated in various applications of machine learning, from natural language processing [2] and speech recognition [3] to computer vision [4]. The tasks can be defined as applying the same model on different data (aka multi-domain learning, MDL) [5, 6], or on various problems [7]. In NN-based MDL, different domains share a set of parameters, which means more training data for the shared feature space, faster training, and better generalization by averaging inherent noise of different domains [8].

MDNet [5] introduced this form of learning into a visual tracking problem, in which different video sequences are considered as different domains to learn from, and the task is defined as a foreground-background separation. This method, however, suffers from several setbacks: (i) this model captures domainindependent representation from different domains with a fully-shared (FS) network (Figure 1(a)). Such an architecture is unable to exclude the domain-specific features from the shared space, and either ignores them (underfit) or includes them in shared representation (overfit). In the former case, the training tries to compensate by non-optimally change the shared feature space. In contrast, in the latter, the learned feature in the shared space is merely useful for one or



Fig. 1. Different methods for obtaining shared representations using multi-domain learning (MDL). Panel (c) depicts our proposed architecture, Adversarial Semi-supervised Private-Shared (ASPS) MTL including adversarial, mutual exclusion, and self-supervised signals that promote the learning task.

a few domains, wasting the model capacity for an inferior representation learning [9]. Therefore, training on one domain inevitably hurts other domains and hinders the emergence of diverse shared features; *(ii)* a shallow network is selected to avoid vanishing gradients, enabling training with limited number of annotated videos, assuming that tracking is relatively easier than object classification and requires less number of layers, and *(iii)* the learned representation does not consider the patterns of the target motion.

To address these issues, we proposed a private-shared (PS) MDL architecture that separates the domain-specific and domain-invariant (shared) features to disentangle the training on different domains and allow for learning an effective, shared feature representation (Figure 1(b)). However, experiments show that PS architecture by itself is not enough to prevent domain-specific features from creeping into the share space [10]. Therefore, we proposed to have a virtual discriminator to predict which domain introduces the feature to the shared space. In a GAN-style optimization, we encourage the shared space to have only domain-independent features. We also introduced a regularization term to penalize redundancy in different feature spaces. To deal with the limited number of annotated video sequences, we employ self-supervised learning in which another virtual classifier using shared feature space is constructed to detect the playback direction (i.e., playing forward and backward) to enforce shared feature space to learn low-level features as well as semantics (Figure 1(c)). The use of both supervised and unsupervised videos improves the discriminative power of shared feature space, as shown with the experiments. Further, we employed ST-ResNets [11] to learn spatiotemporal features, addressing the vanishing gradients, and separating domain-specific motion patterns from the shared ones. Finally, the learned features are used in a custom-made tracking-by-detection method to asses its transferability and effectiveness. In general we (i) propose deep private-shared (PS) MTL architecture for tracking, (ii) propose adversarial training for PS-MTL, (iii) integrate semi-supervised and self-supervised training in a single MTL framework, (iv) enforce learning spatiotemporal features in multi-domain tracking, using self-supervision and a backbone network capable of doing that, and (v) propose SUS-based hard negative batch mining, and (vi)conducted extensive ablation, design choice, and performance experiments.

The problem at hand is closely similar to [12], which tries to perform domain adaptation, leverage multiple datasets with overlapping but distinct class sets, and tries to separate labeled and non-labeled data using multi-domain adversarial learning. However, we explicitly divide features into shared and specific groups, to obtain features that generalize well and push out the specific features that reveal the originating domain in a GAN-style setup.

2 Related Works

Deep Visual Tracking. Early studies in the use of deep learning in tracking utilized features from autoencoders [13, 14] and fully-connected layers of pre-trained (CNN-based) object detector [15]. Still, later the CNN layers themselves were used to serve as features that balance the abstraction level needed to localize the target [16], to provide a spatiotemporal relationship between the target and its background [17], to combine spatial and temporal aspects of tracking [18, 19], and to generate probability maps for the tracker [20]. Recently the tracking studies pay attention to other deep learning approaches including the use of R-CNN for tracking-by-segmentation [21], Siamese Networks for similarity measuring [22–25] and GANs to augment positive samples [26, 27].

Most of the current CNN-based techniques use architectures with 2D convolutions to achieve different invariances to the variations of the images. Meanwhile, the invariance to transformations in time is of paramount importance for video analysis [28]. Modeling temporal information in CNNs has been tackled by applying CNNs on optical flow images [29], reformulating R-CNNs to exploit the temporal context of the images [30] or by the use of separate information pathways for spatial and temporal pathways [31, 32]. Motion-based CNNs typically outperform CNN representations learned from images for tasks dealing with dynamic target, e.g., action recognition [28]. In these approaches, a CNN is applied on a 3-channel optical flow image [33], and different layers of such network provide different variances toward speed and the direction of the target's motion [34]. In visual tracking, deep motion features provide promising results [35], which requires the tracker to fuse the information from two different CNN networks, a temporal and a spatial [32]. However, due to the difficulties in fusing these two streams, some researchers use only the motion features [35]. Deep Multi-Task/Domain Learning. When used with deep learning, MTL

Deep Multi-Task/Domain Learning. When used with deep learning, M1L models tend to share the learned parameters across different tasks either by (i) hard parameter sharing [1] that the hidden layers are shared between all tasks, while task-specific ones are fine-tuned for each task, or by (ii) soft parameter sharing, where each task has its parameters, and the distance between the parameters of the model is then regularized to encourage the parameters to be similar using, e.g., ℓ_1 norm [36] or trace norm [37].

In the hard parameter sharing architectures, shared parameters provide a global feature representation, while task-specific layers further process these features or provide a complementary set of features suitable for a specific task. Some MTL approaches are based on the intuition that learning easy tasks is the prereq-



Fig. 2. Some of features in MDNet are activated for a few sequences containing a particular motion pattern or appearance and and wastes the capacity of shared feature space. Domain-specific features emerged in shared space (left) motion features maximally activated by motion pattern in Joggling (right) appearance feature maximally activated by FaceOcc1 due to the book cover (visualized using [34] and [52]).

uisite for learning more complex ones [8], hence put tasks in hierarchies [38–40] or try to automatically group similar tasks to dynamically form shared layers [41]. When training a multi-task learner, training each task normally increases the task's accuracy (fine-tuning) and, at the same time, provides more information for the shared representation that affects the accuracy of the rest of the tasks (generalization). Balancing the fine-tuning-generalization trade-off has been the subject of several studies. Kendall et al. [42] adjusts tasks' relative weights in the loss function in proportion to the task uncertainty, [43] divides the feature space into task-specific and shared spaces and later employs adversarial learning to encourage shared feature space to contain more common and less taskspecific information [6], and [44] proposed orthogonality constraints to punish redundancy between shared and task layers. In line with this, learning through hints [45] trains a network to predict the most important features. To leverage correlations of different problems in MTL, the tasks could be adversarial [10], provide hints or attention for a main task [46, 47], use a shared memory [48], explicitly perform representation learning for a more complex task [49], facilitate training for a quickly-plateauing main task [50], or learn a shared unsupervised representation [51]. It can be helpful to learn the relations between tasks to enable efficient transfer learning between them [7].

Self-Supervised Learning. Self-supervised learning refers to a learning problem from unlabeled that is framed as a supervised learning problem by using proxy tasks on behalf of the main learning objective. Supervised learning algorithms are used to solve an alternate or pretext task, the result of which is a model or representation that can be used in the solution of the original (actual) modeling problem [53], e.g., by masking a known word in a sentence and trying to predict it [54]. Common examples of self-supervised learning in computer vision is to make colorful images grayscale and have a model to predict a color representation (colorization) [55], removing blocks of the image and have a to model predict the missing parts (inpainting) [56], or rotating an image by random multiples of 90° and predict if the image has the correct rotation [57]. Self-supervised learning extracts additional information from video sequences by exploiting the natural structure in them to use as labels for proxy learning problems. Such information may derive from low-level physics (e.g., gravity), high-level events (e.g., shooting the ball in a soccer match), and artificial cues (e.g., camera motion) [56, 58]. The video sequence order [59], video colorization [60], and video direction [58] have been explored as the pretext task to learn video representations. For this purpose, CNNs are more suitable to learn representations compared to autoencoders [61].

3 Proposed Tracker

This section describes the architecture of the proposed Adversarial Semi-supervised Multi-Domain (ASMD) tracker. We consider multi-domain (domain:=video) learning to share the acquired knowledge from one sequence to another. First, we explore the private-shared MDL framework that provides both domain-specific and shared features for the tracking-by-detection. The network architecture of private and shared modules to capture spatiotemporal features and their training is discussed next. To avoid redundancy and loss of generalization in the shared layers, we introduce adversarial regularization and orthogonality constraint into the MDL training (Fig. 3). Further, to boost the performance of the MDL, the shared representation is intermittently trained using unlabeled videos (Fig. 4).

3.1 Private-Shared Multi-Domain Learning

In FS-MDL (Figure 1(a)), all of the domains share the features extracted by a shared set of layers. Such an idea has been explored in MDNet tracker [5], which assume that all features extracted for one domain is useful for other domains. This model ignores that some features are domain-dependent and may not be useful for other domains. Such features are maximally activated for one or very few domains (i.e., video sequences) while having low activation for others. Figure 2 illustrates some examples of such domain-specific features. On the other hand, PS-MDL [6] divides the feature space of each domain into a domain-invariant part that is shared among all domains, and a domain-specific part. As Figure 1(b) depicts, each domain has its own set of features, that together with the shared features (i.e., concatenated with them), forms the domain features.

3.2 Network Architecture

We used a modified Spatiotemporal ResNet [11] for our private and shared networks. The first layer of motion stream is altered to have 2L = 10 filter channels to operate on the horizontal and vertical optical flow stack of 5 frames. To establish the correspondence between the object and its motion, a residual link connects two streams for each ResNet block. The network receives 224×224 input and reduces it to 7×7 via global average pooling (GAP) (Figure 3). The depth of the networks in MDNet [5] is kept low, however, with the use of ResNet blocks instead of ConvNet here, the network can be significantly deeper.

For each domain the shared features (i.e., concatenation of appearance and motion streams of shared layer) are concatenated with the domain-specific features, followed by two fully connected layers with 128 units and a binary classification fully-connected layer with softmax cross-entropy loss. The network is



Fig. 3. The architecture of the proposed Private-Shared Multi-Domain Network, which consists of one shared and K modified ST-ResNet domain-specific. For domain $k \in \{1..K\}$, the feature space consists of shared features and the corresponding domain-specific features, and the final fully connected layer calculated the classification score of the input patch from video k for a foreground-background classification.

trained to minimize the cross-entropy of the predicted label and the true label of all target candidates for all domains. The loss is computed as:

$$L_{dom} = \sum_{k=1}^{K} \alpha_k L_{track} \left(\hat{y}^{(k)}, y^{(k)} \right) \tag{1}$$

where α_k is the weight of domain k and $L_{track}\left(\hat{y}^{(k)}, y^{(k)}\right)$ is defined in eq(6).

3.3 Adversarial Representation Learning

PS-MDL dissects the representation for each domain into a domain-specific and shared feature spaces but does not guarantee that shared feature space does not include domain-specific features and vice versa. To encourage the shared feature space to include more common information between domains yet has no domain-specific ones, we proposed to use adversarial training. Adversarial training became very popular after the seminal work of Goodfellow et al. [62]. This approach sets up a min-max between networks in which a network G tries to fool the other (D) by generating better samples (i.e., closer to real data distribution), and the other excels in distinguishing the generated data (P_G) from the real ones (P_{data}) , by optimizing

$$\ell = \min_{G} \max_{D} \left(\mathbb{E}_{x \sim P_{data}} [\log D(x)] + \mathbb{E}_{z \sim p(z)} [\log(1 - D(G(z)))] \right)$$
(2)

We propose an adversarial loss function for our MDL-based tracking. The goal is to encourage the shared space to be free from domain-specific knowledge. Therefore, we introduce a discriminator D that predicts the domain k only from the features that emerged in shared representation. The discriminator in each iteration is created from the learned shared ST-ResNet followed by three fully connected layers, each having 512 units and ReLU activation, and a K unit softmax output layer. The multi-class adversarial loss [6] can be formulated as

$$L_{adv} = \min_{\theta_s} \left(\max_{\theta_D} \left(\sum_{k=1}^K \sum_{i=1}^{N_k} \delta_i^k \log \left[D_{\theta_D} \left(\mathbb{E}(\mathbf{x}^{(k)}) \right) \right] \right) \right)$$
(3)

in which θ_D and θ_S are the parameters of the discriminator D and shared layer respectively, δ_i^k indicates the domain of the current input i, and $\mathbb{E}(\mathbf{x})$ denotes shared representation of \mathbf{x} followed by 2 fully connected layers and just before the discriminator layer. The intuition here is that given a sample from domain k, the shared feature space tries to alter its representation to better discriminate the sample, while the discriminator guesses which domain is the sample from. If the learned feature helps the discriminator to find the domain, it includes some domain-specific knowledge [63], and the defined loss L_{adv} punishes this representation. It should be noted that L_{adv} ignores the label $y^{(k)}$ of the input $\mathbf{x}^{(k)}$, allowing for unsupervised training of the shared layer, which is essential to realize the semi-supervised training of the network .

A drawback of the PS-MDL is that the domain-specific and shared feature space may not be mutually exclusive, and redundant domain-invariant features may emerge in both [64]. To alleviate this, inspired by [44], we compare the activations of the private and shared networks, and penalize the high co-activations of the corresponding neurons that are likely to encode a similar latent feature,

$$L_{dif} = \sum_{k=1}^{K} \left\| (\mathbf{f}^s)^T \mathbf{f}^k \right\|_F^2 \tag{4}$$

in which \mathbf{f}^s and \mathbf{f}^k are shared and domain-specific feature vectors (i.e., last layer neuron activations) for a given input and $\|.\|_F^2$ denoted the Frobenius norm.

3.4 Self-Supervised Representation Learning

To enrich the share representation, we want to leverage the structure of the video data to use them as labels for closely related proxy problems that can be solved with the current shared architecture. We select the video direction (forward or backward) as a proxy learning task that is compatible with this architecture.

The order of the frames in a natural video contains spatiotemporal clues about the way objects move. The order of the frames can provide a trivial label for a proxy classification task: to classify if the movie as played in the natural direction or backward. During this task, the network learns spatiotemporal features, which is useful for other tasks such as video understanding, action recognition, and object tracking. Therefore, we select random clips from our dataset, randomly invert the order of some of them, and give them to our network to classify. This helps the features to emerge in the shared layers of our network.



Fig. 4. Self-supervised network using shared ST-ResNet layers, a global average pooling layer, and an output layer to classify "arrow of time" in the 10-frame input clip.

To pull out this task, we draw 10-frame samples from annotated sequences as well as numerous unlabeled videos in YouTubeBB [65] dataset, randomly reverse their order (to augment the data with negative samples), and feed them to the shared ST-ResNet. The extracted features are then aggregated similar to [58] using a GAP layer, and a binary cross-entropy output layer computes the final arrow-of-time loss L_{aot} for the forward-backward classification (Figure 4).

3.5 Semi-Supervised MDL Training

The final loss function of the tracker can be written as

$$L = L_{dom} + \lambda_1 L_{adv} + \lambda_2 L_{dif} + \lambda_3 L_{aot}$$
⁽⁵⁾

in which $\lambda_1, \lambda_2, \lambda_3$ are hyper-parameters. This network is periodically trained on labeled and unlabeled videos such that after every 100 supervised iterations, 1000 unsupervised iterations are conducted.

Supervised. For the supervised multi-domain learning, we collect 50 positive samples and 200 negative samples for each frame such that each positive sample has at least 70% overlap with the target ($IoU \ge 0.7$), and negative samples have $IoU \le 0.5$). For multi-domain representation learning, the network is trained for 1M iterations with learning rate 0.0001 for ST-ResNets and 0.001 for FC layers on GOT10K dataset [66] and its grayscale version. The networks are trained using backpropagation through gradient reversal layer [10] to assist the minimax optimization. For each iteration, five random 10-frame clip and another five with reversed frame-order are extracted for L_{aot} .

Unsupervised. We also conduct 10M iteration of unsupervised training on 240k videos of the YouTubeBB dataset. Again in each iteration, five random 10-frame clips and another five reversed one is extracted from the input video to provide the L_{aot} , and the learning rate for ST-ResNet is set to 0.00001.

4 Online Tracking

After learning the shared features, the private branches are discarded, and 3 new FC layers and a binary softmax output layer is used for every new sequence. During tracking (test time), no domain-specific network is used/trained. All of

FC layers (after shared features) are randomly initialized and trained on-the-fly. The goal of private networks is to capture features that are not captured by the shared network (by trying to overfit the input sequence), introduce them to the feature pool, and receive the features that are pushed out of the shared pool.

4.1 Tracking-by-Detection

For each test sequence, the final FC layers convert the domain-specific and shared feature for all n samples of each T_k frame of video sequence k into classification score $\hat{y} \in [0, 1]$, trained by minimizing the cross-entropy of the sample's label $y \in \{0, 1\}$ and the score, by applying SGD on the loss

$$L_{track}(\hat{y}^{(k)}, y^{(k)}) = -\frac{1}{nT_k} \sum_{t=1}^{T_k} \sum_{j=1}^n y_t^{j(k)} \log\left(\hat{y}_t^{j(k)}\right)$$
(6)

while the shared layers are kept frozen during tracking. This is important to avoid over-fitting to the video sequence [5] and for computational efficiency. After training on the first annotated frame of the sequence, n samples are extracted from each frame t following a Gaussian distribution around the last known target position. The samples are classified using the shared features and dedicated FC layers, and the sample with highest classification score is considered as the new target location. To find a tight bounding box around the target, the target candidates (i.e., score > 0.5) are used in bounding box regression similar to [67].

To update the tracker, a dual-memory [68] is employed in which the network is updated every Δ^s frames (short-memory) with all positive and selected negative samples between t and $t - \Delta^s$. In the long run, the network is likely to forget the initial samples (especially those obtained from user), thus the network is updated every Δ^l frames ($\Delta^l \gg \Delta^s$) with the all previous target estimations (if score > 0.5) and the most uncertain negative sample (score $\rightarrow 0.5$).

4.2 Stochastic Universal Negative Mining

In visual tracking, negative examples (samples) comes from (i) background, or (ii) occluders, distractors and patches that has small overlap with the target. To identify such negative examples, hard negative mining [69] explicitly predict examples' labels and selects misclassified examples (i.e., false positives). Hard minibatch mining [5] selects random subsets of negative examples and selects ones with the lowest classification score as the negative training examples without the need to the examples' labels.

To keep the focus on the hard negatives while keeping the diversity of the negative samples, we borrowed a stochastic universal sampling (SUS) technique [70] from genetic algorithms to select negative training examples via repeated random sampling based on the classification score of the example (Figure 5). This technique gives weaker members of the negative example pool (according to their score) a chance to be chosen to represent different parts of the background in the classifier.



Fig. 5. Stochastic universal sampling to select negative examples for retraining. Using a comb-like ruler, SUS starts from a random real number, and chooses next candidates from the rest of negative examples, preventing the examples with the highest scores to saturate the selection space. S is the sum of negative samples' classification scores.

4.3 Implementation Details

This framework contains several sets of parameters. The tracking parameters including the number of positive and negative samples during train and test time, the covariance of the sampling from each frame, and the parameters of the bounding box regression are kept similar to MDNet [5]. The update intervals are set to $\Delta^s = 15$ and $\Delta^l = 50$. On the MDL side, all domains are equally weighted $\alpha_k = 1/K$, and regularization hyper-parameters are tuned on GOT10K and YouTubeBB datasets via a grid search in $\lambda_i \in [0.01, 0.1]$ and kept fixed through test-time trackings/experiments. Final values are 0.08,0.01,0.6 for $\lambda_{1..3}$.

The FC layers of tracker during the test are trained on the first (annotated) frame for 30 iterations with the learning rate of 0.0001 for hidden layers and 0.001 for the softmax layer. To update the classifier, the learning rate is increased to 0.0003 and 0.003 for hidden and output FC layers for ten iterations, with mini-batches composed of 32 positive and 96 SUS-sampled negatives out of 1024 available negative examples. The tracker is implemented in Matlab using MatConvNet [72] and ST-ResNet Github Repos, and runs around 5.18 fps on Intel Core i7 @ 4.00GHz with Nvidia V100 GPU.

5 Experiment

In a series of experiments using OTB50 dataset [73], the effect of the MDL architecture and the network used in it, and the proposed negative minibatch mining are investigated. An ablation study is conducted to clarify the contributions of each component in the tracker. Finally, the performance of the tracker on recent challenging datasets such as OTB-100 [74], LaSOT [75], UAV123 [76], TrackingNet [77], and VOT-2018 [78] is benchmarked against the state-of-the-art trackers. For this comparison, we have used success and precision plots, where their area under curve provides a robust metric for comparing tracker performances [73]. We also compare all the trackers by the success rate at the

conventional thresholds of 0.50 $(IoU > \frac{1}{2})$ [73]. For the state-of-the-art comparisons, the results are averaged for five different runs using the original semisupervised training (*cf* Section 3.5). The tracker is compared with latest trackers (DiMP [79], ATOM [80], STResCF [81], CBCW [82]), those with bounding box regression (SiamRPN [24], SiamRPN++ [25]), those performing adversarial learning (VITAL [27], SINT++ [26]), and those using MTL (MDNet [5], RT-MDNet [83]) as well as other popular trackers.

Multi-Domain Learning Model. Moving from FS-MDL in MDNet [5] to PS-MDL enables the model to separate the domain-invariant features from domain-specific ones. Adversarial and mutual-exclusion regularizations encourage this effect. The proposed adversarial training cannot operate on the fully-shared architecture of MDNet. Additionally, MDNet's backbone (VGG-M) is unable to learn spatiotemporal features, thus ineffective for our self-supervised learning.

Table 1 presents five variations for the MDL architecture (i) vanilla FS (different from MDNet), that uses only a fully shared ST-ResNet, (ii) PS, that uses a task-specific ST-ResNet for each domain on top of the shared ST-ResNet, (iii) PS+dif, that adds L_{dif} to the PS variation, (iv) APS-, that adds L_{adv} to the PS variation, and (v) APS that uses both L_{adv} and L_{dif} for the network. All variations are **only** trained with supervised data for 200k iterations, and regularization weights are tuned by 5-fold cross-validation for each case. Using PS-MDL without regularizations is not always beneficial, but using both in APS significantly improves the learning performance.

Table 1. Comparison of MDL models with 200k iterations of *supervised* training on GOT10K [66] and tested OTB-50. We also retrained MDNet on this dataset for the same number of iterations (\star) . The first, second and third best methods are highlighted.

	\mathbf{FS}	MDNet^{\star}	\mathbf{PS}	PS+dif	APS-	APS
Average Success	0.62	0.67	0.59	0.41	0.71	0.73
Average Precision	0.73	0.76	0.70	0.55	0.80	0.81
$IoU > \frac{1}{2}$	0.67	0.71	0.61	0.45	0.74	0.76

Backbone Network. We tested our tracker with different backbone architecture such as MDNet-style VGG-M, Two Stream ConvNets with convolutional fusion [84], TwoStream ResNets with late fusion, Spatiotemporal ConvNet [85] and SpatioTemporal ResNets [11]. Table 2 shows that similar to the action recognition domain [11], Spatiotemporal ResNets yields the best results for our intended tasks using to residual connections that enables learning the difference of sequences [54]. Notice the big jump between single-stream and two-stream architectures that is partly because the arrow-of-time self-supervised training is mostly beneficial to train the motion stream.

Negative Minibatch Selection. To select representative negative samples from the pool, and to avoid model drift by insufficiently sampling from critical distractors and occluder, we proposed stochastic universal negative mining (SUNM). Table 3 shows the effectiveness of this method compared to a random

Table 2. Comparison of backbone networks with 200k iterations of supervised training on OTB-50 (1S: single stream, 2S: two-stream, ST: spatiotemporal).

	1S Conv [5]	2S Conv [84]	2S Res	ST Conv [85]	ST Res [11]
Avg Succ	0.64	0.69	0.67	0.63	0.73
Avg Prec	0.71	0.77	0.77	0.71	0.81
$IoU > \frac{1}{2}$	0.66	0.70	0.69	0.66	0.76

Table 3. Comparison of negative minibatch selection for tracking on OTB-50.

	RAND	GRDY	HMM [5]	SUNM (ours)
Average Success	0.72	0.61	0.77	0.80
Average Precision	0.81	0.69	0.89	0.88
$IoU > \frac{1}{2}$	0.75	0.67	0.87	0.93

Table 4. Comparison of the transferability of learned representations with different versions of the proposed framework as well as the pre-training on ILSRVC-2012 action recognition task for leave-one-out evaluation. For each video sequence of OTB50, 49 is used for semi-supervised training, and the remaining video is used for tracking, and the average of the success and precision of the tracker is reported in the table.

	PRE	LOO-FS	LOO-PS	LOO-APS	LOO-ASPS
Avg Succ	0.41	0.70	0.67	0.72	0.78
Avg Prec	0.52	0.76	0.72	0.80	0.85
$IoU > \frac{1}{2}$	0.46	0.79	0.70	0.79	0.87

selection of negative samples (RAND), hard minibatch mining (HMM) [5], and a greedy selection of the most uncertain negative examples (GRDY).

Shared Knowledge Transfer. To test the transferability of the learned representation, we design a leave-on-out experiment on OTB50 dataset. in which out of 50 unique sequences in this dataset, we train the MDL on 49 of them for training, and the remaining one for testing. We use only 100k supervised and 10M unsupervised iterations, and the shared layer is frozen during tracking (LOO-ASPS). This process is repeated for three ablated versions of the proposed model, (i) with fully-shared architecture (LOO-FS), (ii) with the private-shared model with no regularization (LOO-PS), and (iii) without unsupervised iterations (LOO-APS), all trained with 100k supervised iterations. We also set the parameters of shared layers with the ILSRVC-2012 pre-trained network (PRE). Table 4 shows that even FS-MDL in the LOO setting outperforms the tracker made with pre-training for the action recognition task. PS-MDL with regularization but without unsupervised iterations outperform FS-MDL, however, when unsupervised training is used, such distinction becomes more apparent. In summary, it is shown that the full treatment yields better results due to multi-task learning (PRE<LOO-FS), adversarial training (LOO-PS<LOO-FS<LOO-SPS), and self-supervised training (LOO-APS<LOO-ASPS).

Tracking Challenges Analysis. Table 5 presents the performance of the proposed tracker under challenging subsets of OTB50. The results reveal that using proposed PS-MTL architecture as well as the semi-supervised training scheme almost improved every aspect of the tracker compared to the baseline (MDNet).

Attribute	TLD	STRK	TGPR	MEEM	MUSTer	STAPLE	CMT	SRDCF	CCOT	MDNet	Ours
	[86]	[87]	[88]	[89]	[90]	[91]	[92]	[93]	[94]	[5]	
Illumination Variation	0.48	0.53	0.54	0.62	0.73	0.68	0.73	0.70	0.75	0.76	0.80
Deformation	0.38	0.51	0.61	0.62	0.69	0.70	0.69	0.67	0.69	0.73	0.74
Occlusion	0.46	0.50	0.51	0.61	0.69	0.69	0.69	0.70	0.76	0.75	0.81
Scale Variation	0.49	0.51	0.50	0.58	0.71	0.68	0.72	0.71	0.76	0.78	0.82
In-plane Rotation	0.50	0.54	0.56	0.58	0.69	0.69	0.74	0.70	0.72	0.75	0.78
Out-of-plane Rotation	0.48	0.53	0.54	0.62	0.70	0.67	0.73	0.69	0.74	0.76	0.80
Out-of-View (Shear)	0.54	0.52	0.44	0.68	0.73	0.62	0.71	0.66	0.79	0.79	0.84
Low Resolution	0.36	0.33	0.38	0.43	0.50	0.47	0.55	0.58	0.70	0.72	0.71
Background Clutter	0.39	0.52	0.57	0.67	0.72	0.67	0.69	0.70	0.70	0.76	0.77
Fast Motion	0.45	0.52	0.46	0.65	0.65	0.56	0.70	0.63	0.72	0.73	0.78
Motion Blur	0.41	0.47	0.44	0.63	0.65	0.61	0.65	0.69	0.72	0.72	0.78
Avg. Succ	0.49	0.55	0.56	0.62	0.72	0.69	0.72	0.70	0.75	0.76	0.80
Avg. Prec	0.60	0.66	0.68	0.74	0.82	0.76	0.83	0.78	0.84	0.85	0.88
IoU > 0.5	0.59	0.64	0.66	0.75	0.86	0.82	0.83	0.83	0.90	0.93	0.93

Table 5. Quantitative evaluation of trackers under different visual tracking challenges of OTB50 [73] using AUC of success plot and their overall precision.

Table 6. Quantitative evaluation on OTB100 using success rate and precision.

	dSRDCF	CCOT	BACF	dSTRCF	STResCF	CBCW	ECO	SiamRPN	SiamRPN++	SINT++	VITAL	MDNet	RT-MDNet	ATOM	DiMP	Ours
	[95]	[94]	[96]	[97]	[81]	[82]	[<mark>98</mark>]	[24]	[25]	[26]	[27]	[5]	[83]	[80]	[79]	
Avg. Succ	0.69	0.68	0.62	0.68	0.59	0.61	0.69	0.63	0.69	0.57	0.68	0.67	0.65	0.66	0.68	0.73
Avg. Prec↑	0.81	0.85	0.82	-	0.83	0.81	0.91	0.85	0.91	0.76	0.91	0.90	0.88	-	0.89	0.91
$IoU > \frac{1}{2}$	0.78	0.88	0.77	0.77	0.76	0.76	-	0.80	0.83	0.78	0.75	0.80	0.79	0.86	0.87	0.88

Table 7. Evaluation on VOT2018 by expected avg overlap, robustness and accuracy.

	STURCK	MEEM	STAPLE	SRDCF	CCOT	SiamFC	ECO	SiamRPN	SiamRPN++	ATOM	DiMP	Ours
	[87]	[89]	[91]	[93]	[94]	[23]	[98]	[24]	[25]	[<mark>80</mark>]	[79]	
EAO↑	0.097	0.192	0.169	0.119	0.267	0.188	0.280	0.383	0.414	0.401	0.440	0.427
$\mathrm{Acc}\uparrow$	0.418	0.463	0.530	0.490	0.494	0.503	0.484	0.586	0.600	0.590	0.597	0.604
$\mathrm{Rob}{\downarrow}$	1.297	0.534	0.688	0.974	0.318	0.585	0.276	0.276	0.234	0.204	0.153	0.169

This improvement is more significant for occlusions, out-of-view and motion challenges due to the usage of spatiotemporal features in the tracker.

Evaluation on OTB-100. We benchmarked the tracker against the competing trackers on OTB100 [74] as presented in Table 6. In comparison with trackers who perform BB regression, those using adversarial learning, and the state-of-the-art in the tracking, our proposed tracker shows beter performance.

Evaluation on VOT-2018. Table 7 shows the comparison of our method with the competing algorithms on 60 challenging videos of VOT-2018 [78].

Evaluation on LaSOT [75]. In phase I, our tracker is trained on the GOT10K dataset for 1M supervised iterations and on YouTubeBB [65] for 10M unsupervised iterations. As showed in table 8, our proposed algorithm obtained the best results compared to state-of-the-art trackers when tested on all 1400 video sequences of LaSOT. Protocol II brings an even more interesting challenge, that limits the training data to the given 1120 videos, and test the trained trackers on the remaining 280 sequences of the dataset. Using the given training videos, we conduct 1M supervised and 10M unsupervised training. The shared layer obtained a significant improvement compared to the MDNet and the state-of-the-art, which can be attributed to the abundance of labeled data, and the close distributions of training and test data, essential for unsupervised training.

Table 8. Evaluation on LaSOT with protocol I (testing on all videos) and protocol II (training on given videos and testing on the rest). We get better results with dataset's own videos as training due to large training set and matching domain.

	STAPLE	SRDCF	SiamFC	SINT	ECO	BACF	SiamRPN++	VITAL	MDNet	ATOM	DiMP	Ours
	[91]	[93]	[23]	[22]	[98]	[96]	[25]	[27]	[5]	[<mark>80</mark>]	[79]	
(I) $Acc\uparrow$	0.266	0.271	0.358	0.339	0.340	0.277	0.496	0.412	0.413	0.515	0.569	0.554
(I) $\operatorname{Rob}\uparrow$	0.231	0.227	0.341	0.229	0.298	0.239	-	0.372	0.374	-	-	0.487
$(II)Acc\uparrow$	0.243	0.245	0.336	0.314	0.324	0.259	-	0.390	0.397	-	-	0.499
(II)Rob↑	0.239	0.219	0.339	0.295	0.301	0.239	-	0.360	0.373	-	-	0.495

Table 9. Evaluation on UAV123 by success rate and precision.

	STRUC	K MEEM	STAPLE	E SRDCF	MUSTer	ECO	CCOT	SiamRPN	SiamRPN++	MDNet	RT-MDNet	ATOM	DiMP	Ours
	[87]	[89]	[91]	[93]	[<mark>90</mark>]	[98]	[94]	[24]	[25]	[5]	[83]	[80]	[79]	
$\operatorname{Succ}\uparrow$	0.387	0.398	0.453	0.473	0.517	0.522	0.513	0.527	0.613	0.528	0.528	0.643	0.654	0.655
$\operatorname{Prec}\uparrow$	0.578	0.627	0.666	0.676	0.391	0.591	-	0.748	0.807	-	0.772	0.856	0.858	0.793

Table 10. Evaluation on TrackingNet by precision, normalized precision, and success.

	ECO	SiamFC	CFNet	MDNet	DaSiam-RPN	ATOM	SiamRPN++	DiMP	Ours
	[98]	[23]		[5]	[99]	[80]	[25]	[79]	
Precision	0.492	0.533	0.533	0.565	0.591	0.648	0.694	0.687	0.687
Norm. Prec.	0.618	0.666	0.654	0.705	0.733	0.771	0.800	0.801	0.802
Success (AUC)	0.554	0.571	0.578	0.606	0.638	0.703	0.733	0.740	0.741

Evaluation on UAV123. The result in table 9 suggest that the shared feature space is good enough to surpass the state-of-the-art trackers for tracking in this dataset, but not good enough to obtain a high success score. In this regard, we conducted an additional experiment by 5-fold cross validating the network trained for 100K supervised iterations and 1M unsupervised iteration for each fold (we ensured that each fold has exactly 4 of the 20 longer videos of the dataset, called UAV20L subset). The average success score for this task on UAV123 reaches 0.655 that is almost 7.8% improvement compared to the generally trained version of the proposed tracker (on GOT10K and YouTubeBB).

Evaluation on TrackingNet. Table 10 shows our results on the TrackingNet test set [77] (511 videosh) that is comparable to the SOTA approaches.

6 Conclusion

We proposed a semi-supervised private-shared MDL to learn the domain-invariant and domain-specific features, push the domain-specific features out of the shared feature space using an adversarial regularization, and use the direction of video playing as a proxy learning task to further train the shared representation using unannotated videos. We then proposed a tracker that classifies the target from the background for each video that uses the learned feature representation. The results of the demonstrate the superior performance of the proposed tracker compared to the state-of-the-art. Our next step is to tailor our adversarial learning and global search for long-term tracking as proposed in [100, 101].

References

- 1. Caruana, R.: Multitask learning: A knowledge-based source of inductive bias. (1993) 1, 3
- Collobert, R., Weston, J.: A unified architecture for natural language processing: Deep neural networks with multitask learning. In: ICML'08, ACM (2008) 160–167
- Deng, L., Hinton, G., Kingsbury, B.: New types of deep neural network learning for speech recognition and related applications: An overview. In: ICASSP'13, IEEE (2013) 8599–8603 1
- 4. Girshick, R.: Fast R-CNN. In: ICCV'15. (2015) 1440–1448 1
- 5. Nam, H., Han, B.: Learning multi-domain convolutional neural networks for visual tracking. (In: CVPR'16) 1, 5, 9, 10, 11, 12, 13, 14
- Liu, P., Qiu, X., Huang, X.: Adversarial multi-task learning for text classification. In: ACL'17. (2017) 1–10 1, 4, 5, 7
- Roshan Zamir, A., Sax, A., Shen, W., Guibas, L.J., Malik, J., Savarese, S.: Taskonomy: Disentangling task transfer learning. In: CVPR'18. (2018) 3712–3722 1, 4
- Ruder, S.: An overview of multi-task learning in deep neural networks. arXiv preprint arXiv:1706.05098 (2017) 1, 4
- Bengio, Y., Courville, A., Vincent, P.: Representation learning: A review and new perspectives. PAMI (2013) 2
- Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: ICML'15. (2015) 1180–1189 2, 4, 8
- Feichtenhofer, C., Pinz, A., Wildes, R.: Spatiotemporal residual networks for video action recognition. In: NIPS'16. (2016) 3468–3476 2, 5, 11, 12
- Sebag, A.S., Heinrich, L., Schoenauer, M., Sebag, M., Wu, L., Altschuler, S.: Multi-domain adversarial learning. In: ICLR'2019. (2019) 3
- Wang, N., Yeung, D.Y.: Learning a deep compact image representation for visual tracking. In: NIPS. (2013) 809–817
- Zhou, X., Xie, L., Zhang, P., Zhang, Y.: An ensemble of deep neural networks for object tracking. In: Image Processing (ICIP), 2014 IEEE International Conference on, IEEE (2014) 843–847 3
- Fan, J., Xu, W., Wu, Y., Gong, Y.: Human tracking using convolutional neural networks. IEEE Transactions on Neural Networks 21 (2010) 1610–1623 3
- Ma, C., Huang, J.B., Yang, X., Yang, M.H.: Hierarchical convolutional features for visual tracking. In: ICCV'15. (2015) 3074–3082 3
- 17. Zhang, K., Liu, Q., Wu, Y., Yang, M.: Robust visual tracking via convolutional networks without training. IEEE TIP **25** (2016) 1779–1792 **3**
- Zhu, Z., Huang, G., Zou, W., Du, D., Huang, C.: Uct: learning unified convolutional networks for real-time visual tracking. In: ICCVw. (2017) 1973–1982
 3
- 19. Chen, K., Tao, W.: Once for all: a two-flow convolutional neural network for visual tracking. IEEE CSVT (2018) 1–1 $\stackrel{3}{-}$
- Wang, N., Li, S., Gupta, A., Yeung, D.Y.: Transferring rich feature hierarchies for robust visual tracking. arXiv (2015) 3
- Drayer, B., Brox, T.: Object detection, tracking, and motion segmentation for object-level video segmentation. arXiv preprint arXiv:1608.03066 (2016) 3
- Tao, R., Gavves, E., Smeulders, A.W.: Siamese instance search for tracking. In: CVPR. (2016) 1420–1429 3, 14

- 16 K. Meshgi & M.S. Mirzaei
- Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fullyconvolutional siamese networks for object tracking. In: ECCV, Springer (2016) 850–865 3, 13, 14
- 24. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: CVPR'18. (2018) 3, 11, 13, 14
- Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., Yan, J.: Siamrpn++: Evolution of siamese visual tracking with very deep networks. In: CVPR'19. (2019) 3, 11, 13, 14
- Wang, X., Li, C., Luo, B., Tang, J.: Sint++: Robust visual tracking via adversarial positive instance generation. In: CVPR'18. (2018) 3, 11, 13
- 27. Song, Y., Ma, C., Wu, X., Gong, L., Bao, L., Zuo, W., Shen, C., Rynson, L., Yang, M.H.: Vital: Visual tracking via adversarial learning. In: CVPR. (2018) 3, 11, 13, 14
- Varol, G., Laptev, I., Schmid, C.: Long-term temporal convolutions for action recognition. PAMI 40 (2018) 1510–1517 3
- 29. Gkioxari, G., Malik, J.: Finding action tubes. In: CVPR. (2015) 759-768 3
- Chao, Y.W., Vijayanarasimhan, S., Seybold, B., Ross, D.A., Deng, J., Sukthankar, R.: Rethinking the faster R-CNN architecture for temporal action localization. In: CVPR'18. (2018) 1130–1139 3
- Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS. (2014) 568–576
- Zhu, Z., Wu, W., Zou, W., Yan, J.: End-to-end flow correlation tracking with spatial-temporal attention. CVPR 42 (2017) 20 3
- 33. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: Flownet: Learning optical flow with convolutional networks. In: ICCV'15. (2015) 2758–2766 3
- 34. Feichtenhofer, C., Pinz, A., Wildes, R.P., Zisserman, A.: What have we learned from deep representations for action recognition? connections **19** (2018) 29 3, 4
- Gladh, S., Danelljan, M., Khan, F.S., Felsberg, M.: Deep motion features for visual tracking. In: ICPR, IEEE (2016) 1243–1248 3
- Duong, L., Cohn, T., Bird, S., Cook, P.: Low resource dependency parsing: Crosslingual parameter sharing in a neural network parser. In: ACL-IJCNLP'15. (2015) 845–850 3
- Yang, Y., Hospedales, T.M.: Trace norm regularised deep multi-task learning. ICLR'2017 (2017) 3
- Søgaard, A., Goldberg, Y.: Deep multi-task learning with low level tasks supervised at lower layers. In: ACL'16. (2016) 231–235 4
- Hashimoto, K., Tsuruoka, Y., Socher, R., et al.: A joint many-task model: Growing a neural network for multiple nlp tasks. In: EMNLP'17. (2017) 1923–1933
- 40. Sanh, V., Wolf, T., Ruder, S.: A hierarchical multi-task approach for learning embeddings from semantic tasks. In: AAAI'19. Volume 33. (2019) 6949–6956 4
- Liu, S., Pan, S.J., Ho, Q.: Distributed multi-task relationship learning. In: ACM SIGKDD'17, ACM (2017) 937–946
- Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: CVPR'18. (2018) 7482–7491
- Liu, P., Qiu, X., Huang, X.: Recurrent neural network for text classification with multi-task learning. arXiv preprint arXiv:1605.05101 (2016) 4
- 44. Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., Erhan, D.: Domain separation networks. In: NIPS'16. (2016) 343–351 4, 7

- Abu-Mostafa, Y.S.: Learning from hints in neural networks. Journal of Complexity 6 (1990) 192–198 4
- 46. Yu, J., Jiang, J.: Learning sentence embeddings with auxiliary tasks for crossdomain sentiment classification. In: EMNLP'16. (2016) 236–246 4
- 47. Caruana, R.: Multitask learning. Machine learning 28 (1997) 41-75 4
- Liu, P., Qiu, X., Huang, X.: Deep multi-task learning with shared memory for text classification. In: EMNLP'2016. (2016) 4
- Rei, M.: Semi-supervised multitask learning for sequence labeling. In: ACL'17. (2017) 2121–2130 4
- Bingel, J., Søgaard, A.: Identifying beneficial task relations for multi-task learning in deep neural networks. In: ACL'15. (2017) 164–169
- Doersch, C., Zisserman, A.: Multi-task self-supervised visual learning. In: ICCV'17. (2017) 2051–2060 4
- Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: ECCV, Springer (2014) 818–833 4
- Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. arXiv preprint arXiv:1901.09005 (2019) 4
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018) 4, 11
- Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: ECCV'2016, Springer (2016) 4
- Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: ICCV'15. (2015) 1422–1430 4, 5
- 57. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. arXiv preprint arXiv:1803.07728 (2018) 4
- Wei, D., Lim, J.J., Zisserman, A., Freeman, W.T.: Learning and using the arrow of time. In: CVPR'18. (2018) 8052–8060 5, 8
- Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: ECCV'16, Springer (2016) 527–544
- Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos. In: ECCV'18. (2018) 391–408 5
- Li, H., Li, Y., Porikli, F.: Deeptrack: Learning discriminative feature representations online for robust visual tracking. IEEE TIP 25 (2016) 1834–1848 5
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS'14. (2014) 2672–2680 6
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. (Machine learning) 7
- Jia, Y., Salzmann, M., Darrell, T.: Factorized latent spaces with structured sparsity. In: NIPS'10. (2010) 982–990 7
- Real, E., Shlens, J., Mazzocchi, S., Pan, X., Vanhoucke, V.: Youtubeboundingboxes: A large high-precision human-annotated data set for object detection in video. In: CVPR'17. (2017) 5296–5305 8, 13
- Huang, L., Zhao, X., Huang, K.: Got-10k: A large high-diversity benchmark for generic object tracking in the wild. PAMI (2019) 8, 11
- Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR'14. (2014) 580– 587 9
- Meshgi, K., Oba, S., Ishii, S.: Efficient diverse ensemble for discriminative cotracking: Supplementary material. CVPR (2018) 9

- 18 K. Meshgi & M.S. Mirzaei
- Sung, K.K., Poggio, T.: Example-based learning for view-based human face detection. PAMI (1998) 9
- Baker, J.E.: Reducing bias and inefficiency in the selection algorithm. In: Proceedings of the second international conference on genetic algorithms. (1987) 9
- Lewis, D.D., Catlett, J.: Heterogeneous uncertainty sampling for supervised learning. In: ICML'94. (1994) 148–156 10
- Vedaldi, A., Lenc, K.: Matconvnet: Convolutional neural networks for matlab. In: Proceedings of the 23rd ACM international conference on Multimedia, ACM (2015) 689–692 10
- Wu, Y., Lim, J., Yang, M.H.: Online object tracking: A benchmark. In: CVPR'13, IEEE (2013) 2411–2418 10, 11, 13
- 74. Wu, Y., Lim, J., Yang, M.H.: Object tracking benchmark. PAMI 37 (2015) 1834–1848 11, 13
- Fan, H., et al.: LaSOT: A high-quality benchmark for large-scale single object tracking. CVPR'19 (2019) 11, 13
- Mueller, M., et al.: A benchmark and simulator for uav tracking. In: ECCV'16. (2016) 11
- 77. Muller, M., Bibi, A., Giancola, S., Alsubaihi, S., Ghanem, B.: Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In: ECCV'2018. (2018) 11, 14
- Kristan, M., et al.: The sixth visual object tracking vot2018 challenge results. In: ECCV'18. (2018) 11, 13
- Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: ICCV'2019. (2019) 11, 13, 14
- Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: Atom: Accurate tracking by overlap maximization. In: CVPR'19. (2019) 11, 13, 14
- Zhu, Z., et al.: STResNet_cf tracker: The deep spatiotemporal features learning for correlation filter based robust visual object tracking. IEEE Access (2019) 11, 13
- Zhou, Y., et al.: Efficient correlation tracking via center-biased spatial regularization. (TIP) 11, 13
- Jung, I., Son, J., Baek, M., Han, B.: Real-time mdnet. In: ECCV. (2018) 11, 13, 14
- Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional two-stream network fusion for video action recognition. In: CVPR'16. (2016) 11, 12
- Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: CVPR'14. (2014) 11, 12
- Kalal, Z., Mikolajczyk, K., Matas, J.: Tracking-learning-detection. PAMI 34 (2012) 1409–1422 13
- Hare, S., Saffari, A., Torr, P.H.: Struck: Structured output tracking with kernels. In: ICCV'11. (2011) 13, 14
- Gao, J., Ling, H., Hu, W., Xing, J.: Transfer learning based visual tracking with gaussian processes regression. In: ECCV'14. Springer (2014) 188–203 13
- Zhang, J., Ma, S., Sclaroff, S.: Meem: Robust tracking via multiple experts using entropy minimization. (In: ECCV'14) 13, 14
- Hong, Z., Chen, Z., Wang, C., Mei, X., Prokhorov, D., Tao, D.: Multi-store tracker (muster): a cognitive psychology inspired approach to object tracking. (In: CVPR'15) 13, 14

- Bertinetto, L., Valmadre, J., Golodetz, S., Miksik, O., Torr, P.H.: Staple: Complementary learners for real-time tracking. In: CVPR'16. (2016) 1401–1409 13, 14
- Meshgi, K., Oba, S., Ishii, S.: Active discriminative tracking using collective memory. (In: MVA'17) 13
- Danelljan, M., Hager, G., Shahbaz Khan, F., Felsberg, M.: Learning spatially regularized correlation filters for visual tracking. In: ICCV'15. (2015) 4310–4318 13, 14
- Danelljan, M., Robinson, A., Khan, F.S., Felsberg, M.: Beyond correlation filters: Learning continuous convolution operators for visual tracking. (In: ECCV'16) 13, 14
- Danelljan, M., Hager, G., Shahbaz Khan, F., Felsberg, M.: Convolutional features for correlation filter based visual tracking. In: ICCVw. (2015) 58–66 13
- Kiani Galoogahi, H., Fagg, A., Lucey, S.: Learning background-aware correlation filters for visual tracking. (In: ICCV'17) 13, 14
- Li, F., et al.: Learning spatial-temporal regularized correlation filters for visual tracking. In: CVPR'18. (2018) 13
- Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: Eco: Efficient convolution operators for tracking. In: CVPR. (2017) 13, 14
- 99. Zhu, Z., Wang, Q., Li, B., Wu, W., Yan, J., Hu, W.: Distractor-aware siamese networks for visual object tracking. In: ECCV'2018. (2018) 14
- 100. Wang, X., Yang, R., Sun, T., Luo, B.: Learning target-aware attention for robust tracking with conditional adversarial network. In: BMVC. (2019) 131 15
- Huang, L., Zhao, X., Huang, K.: Globaltrack: A simple and strong baseline for long-term tracking. AAAI (2020) 15