



Pointly-Supervised Action Localization

Pascal Mettes¹ · Cees G. M. Snoek¹

Received: 29 May 2018 / Accepted: 31 August 2018 / Published online: 11 September 2018
© The Author(s) 2018

Abstract

This paper strives for spatio-temporal localization of human actions in videos. In the literature, the consensus is to achieve localization by training on bounding box annotations provided for each frame of each training video. As annotating boxes in video is expensive, cumbersome and error-prone, we propose to bypass box-supervision. Instead, we introduce action localization based on point-supervision. We start from unsupervised spatio-temporal proposals, which provide a set of candidate regions in videos. While normally used exclusively for inference, we show spatio-temporal proposals can also be leveraged during training when guided by a sparse set of point annotations. We introduce an overlap measure between points and spatio-temporal proposals and incorporate them all into a new objective of a multiple instance learning optimization. During inference, we introduce pseudo-points, visual cues from videos, that automatically guide the selection of spatio-temporal proposals. We outline five spatial and one temporal pseudo-point, as well as a measure to best leverage pseudo-points at test time. Experimental evaluation on three action localization datasets shows our pointly-supervised approach (1) is as effective as traditional box-supervision at a fraction of the annotation cost, (2) is robust to sparse and noisy point annotations, (3) benefits from pseudo-points during inference, and (4) outperforms recent weakly-supervised alternatives. This leads us to conclude that points provide a viable alternative to boxes for action localization.

Keywords Action localization · Point supervision · Spatio-temporal proposals

1 Introduction

This paper aims to recognize and localize actions such as *skiing*, *running*, and *getting out of a vehicle* in videos. Action recognition has been a vibrant topic in vision for several decades, resulting in approaches based on local spatio-temporal features (Dollár et al. 2005; Laptev 2005; Wang et al. 2009), dense trajectories (Jain et al. 2013; Wang et al. 2013) two-stream neural networks (Simonyan and Zisserman 2014; Feichtenhofer et al. 2016), 3D convolutions (Ji et al. 2013; Tran et al. 2015), and recurrent networks (Donahue et al. 2015; Li et al. 2018; Srivastava et al. 2015). We aim to not only recognize which actions occur in videos, but also discover when and where the actions are present.

Action localization in videos corresponds to finding tubes of consecutive bounding boxes in frames for each action. Initial work aimed at localizing actions by finding local discriminative parts and generating tubes through linking or sliding windows (Lan et al. 2011; Tian et al. 2013a; Wang et al. 2014). State-of-the-art localizers classify boxes per frame (or few frames) before linking them into tubes (Gkioxari and Malik 2015; Weinzaepfel et al. 2015; Hou et al. 2017; Kalogeton et al. 2017a). Regardless the approach, a requirement for all these works is the need for box-supervision per frame of each training video. As annotating boxes in videos is an expensive, cumbersome and error-prone endeavor, we prefer to perform action localization without the need for box supervision.

The first contribution of this paper is to localize actions in videos with the aid of point-supervision. For pointly-supervised action localization, we start from (unsupervised) spatio-temporal proposals. Spatio-temporal proposals reduce the search space of actions in videos to a few hundred to thousand tubes, where at least one tube matches well with the ground truth action location (Jain et al. 2014; van Gemert et al. 2015; Jain et al. 2017; Oneata et al. 2014). This is typ-

Communicated by M. Hebert.

✉ Pascal Mettes
P.S.M.Mettes@uva.nl

Cees G. M. Snoek
cgmsnoek@uva.nl

¹ Universiteit van Amsterdam, Amsterdam, the Netherlands

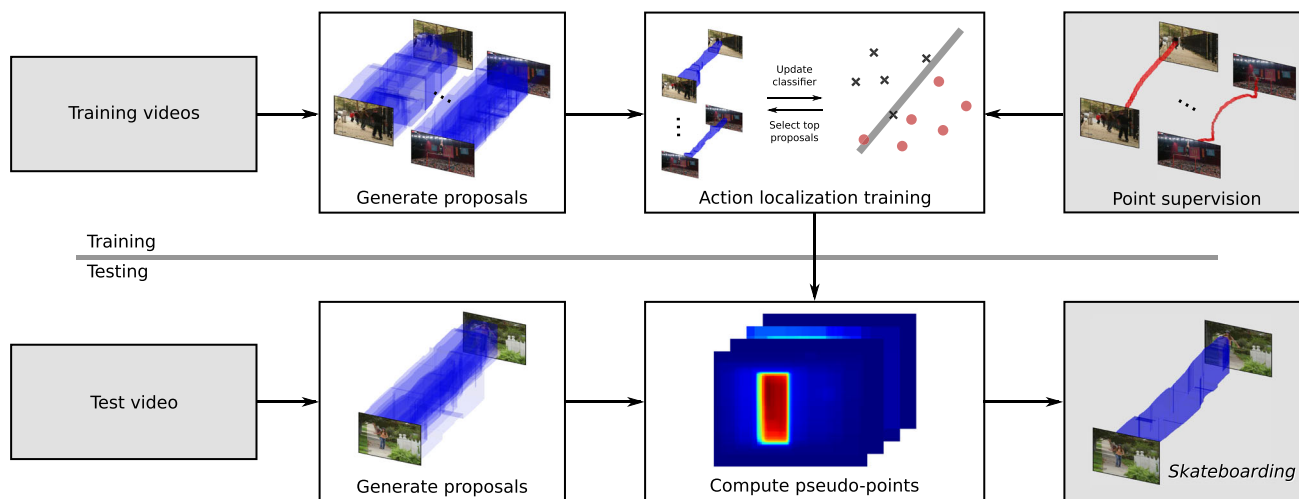


Fig. 1 Pointly-supervised action localization using spatio-temporal proposals and pseudo-points. During training, we start from point-supervision for each video. Our overlap measure computes the match between each proposal and the point annotations. We iteratively refine

ically achieved by clustering local representations such as supervoxels or dense trajectories. In the literature, the use of spatio-temporal proposals is restricted to the inference stage; training of the action localizer that select the best proposal still depends on box-supervision. While the spatio-temporal proposals may be unsupervised, they do not relax the need for box-supervision during the training stage of action localizers. We propose to bypass bounding box annotations by training action localizers on spatio-temporal proposals from *training* videos. We show that training on spatio-temporal proposals guided by point annotations, yields similar action localization performance to their box-supervised alternative at a fraction of the annotation time.

As our second contribution, we propose an overlap measure that matches the centers of spatio-temporal proposals with point annotations. To identify the best proposal to train on, we adopt a multiple instance learning perspective (Andrews et al. 2002), with the spatio-temporal proposals defining the instances and videos the bags. We employ the max-margin multiple instance learning formulation and extend it to incorporate information from the proposed overlap measure. This results in action localization using video labels and point annotations as the sole action supervision. Our first two contributions were previously presented in the conference version of this paper (Mettes et al. 2016).

For our third contribution we are inspired by Mettes et al. (2017), who propose to train action localizers with spatio-temporal proposals selected by automatic visual cues. Rather than employing the cues at training time, we prefer to exploit the cues during *inference* and call them pseudo-points. The pseudo-points are used as an unsupervised generalization of point-supervision during the testing stage. The pseudo-points cover cues from training point statistics, person detection

the proposal selection by extending the max-margin multiple instance learning formulation. During inference, we compute pseudo-points for all video frames and use them in conjunction with the learned action model to determine the top proposals per action over all test videos

(Yu and Yuan 2015), independent motion (Jain et al. 2014), spatio-temporal proposals (van Gemert et al. 2015), center bias (Tseng et al. 2009), and temporal information. To link the point-supervision in training videos to pseudo-points in test videos, we propose a function that both weights and selects pseudo-points based on how well they match with points annotated during training. We use the weighting function to determine which pseudo-points are most effective and how much they should contribute to the selection of spatio-temporal proposals in test videos. A complete overview of our proposed approach is shown in Fig. 1.

The rest of the paper is organized as follows. In Sect. 2, we describe related work. Section 3 details our algorithm for point-supervision during training. Section 4 presents pseudo-points and explains how to leverage them during inference. We detail our experimental setup on UCF Sports (Rodriguez et al. 2008), UCF-101 (Soomro et al. 2012) and Hollywood2Tubes (Mettes et al. 2016) in Sect. 5. Ablation studies, error diagnosis and comparisons are discussed in Sect. 6. We conclude our work in Sect. 7.

2 Related Work

2.1 Action Localization with Box-supervision

The problem of action localization is commonly tackled by supervision with video-level action class labels and frame-level box annotations during training. Initial approaches do so through figure-centric structures (Lan et al. 2011) and part-based models (Tian et al. 2013a; Wang et al. 2014). Inspired by the success of object proposals in images (Uijlings et al. 2013), several works have investigated spatio-temporal pro-

posals for action localization in videos. Such spatio-temporal proposals are typically generated by grouping supervoxels (Jain et al. 2017; Oneata et al. 2014; Soomro et al. 2015) or dense trajectories (van Gemert et al. 2015; Marian Puscas et al. 2015). Spatio-temporal proposals reduce the search space to a few hundred or thousand locations per video. In the literature, the use of spatio-temporal proposals is limited to the testing stage. Training is still performed on features derived from bounding box annotations per frame. In this paper, we extend the use of action proposals to the training stage. We show that proposals provide high quality training examples when leveraging our multiple instance learning variant, guided by point annotations, completely alleviating the need for box annotations.

Recently, a number of works have achieved success in action localization by separating spatial detection from temporal linking (Gkioxari and Malik 2015; Weinzaepfel et al. 2015). Such approaches have been further improved with better representations (Peng and Schmid 2016; Saha et al. 2016; Yang et al. 2017), joint linking algorithms (Singh et al. 2017), and by classifying a few consecutive frames before linking (Hou et al. 2017; Kalogeiton et al. 2017a; Saha et al. 2017). While effective, these approaches have an inherent requirement for box annotations to detect and regress the boxes in video frames. We focus on the use of unsupervised spatio-temporal proposals (Jain et al. 2014; van Gemert et al. 2015; Jain et al. 2017; Oneata et al. 2014), and we show how to utilize them during training to bypass the need for box-supervision.

2.2 Action Localization Without Box-supervision

Given the annotation burden for box-supervision in action localization, several works have investigated action localization from weaker supervision signals. Most works focus on localization from video labels only. Siva and Xiang (2011) employ spatio-temporal proposals and optimize for an action localization model through multiple instance learning (Andrews et al. 2002), where the videos are the bags and the proposals are the instances. We show that multiple instance learning yields suboptimal results for action localization; extending multiple instance learning with point-supervision alleviates this problem.

Chen and Corso (2015) also employ spatio-temporal proposals and video labels, but skip the multiple instance learning step. Instead, they train on the most dominant proposal per training video, without knowing whether the proposal fits the action location well. Recent work by Li et al. (2018) achieves action localization from video labels through attention. The action location is determined by a box around the center of attention in each frame, followed by a linking procedure. These approaches provide action localization without box annotations. However, using only the video

label restricts the localization performance. We show that point annotations have a direct impact on the performance at the expense of a small additional annotation cost, outperforming approaches using video labels only.

Several recent works have investigated action localization in a zero-shot setting, where no video training examples are provided for test actions. This is typically achieved through semantic word embeddings (Mikolov et al. 2013) between actions and objects as found in text corpora. Initial work by Jain et al. (2015) employed spatio-temporal proposals and assigned object classifier scores to each proposal. The object scores are combined with the word embedding scores given an action and the highest scoring proposal is selected for each test video. Mettes and Snoek (2017) perform zero-shot action localization by linking boxes that are scored based on a spatial-aware embedding between actors and objects. Kalogeiton et al. (2017b) perform zero-shot localization through joint localization of actions and objects. Soomro and Shah (2017) aim for unsupervised action localization through discriminative clustering on videos and spatio-temporal action proposal generation with 0-1 Knapsack. Such works are promising but do not perform on the level of (weakly) supervised alternatives, as detailed in our final experiment.

2.3 Speeding-up Box Annotations

Easing the annotation burden of bounding box annotations in videos has been investigated by Vondrick et al. (2013). They investigate different strategies to annotate boxes in videos, e.g., with expert annotators and tracking. Furthermore, several works have attempted faster ways to annotate boxes, e.g., through human verification (Russakovsky et al. 2015; Papadopoulos et al. 2016) or by clicking the extremes of objects (Papadopoulos et al. 2017). While such investigations and approaches provide faster alternatives to the costly ImageNet standard for box annotation (Su et al. 2012), annotating boxes remains a slow and manually expensive endeavor. In this work, we avoid box annotations and show that action localization can be done efficiently through simple point annotations.

Several recent works have investigated the merit of point annotations in other visual recognition challenges. Bearman et al. (2016) investigate point-supervision for semantic segmentation in images, which constitutes a fraction of the annotation cost compared to pixel-wise segmentation. In the video domain, Jain and Grauman (2016) investigate object segmentation based on point clicks. Similar in spirit to our work, Manen et al. (2017) show the spatio-temporal tracks from consecutive point annotations provide a rich supervision for multiple object tracking in videos. In this work, we investigate the potential of point-supervision for action localization in videos, showing we can reach comparable performance to full box-supervision approaches based on action proposals.

3 Point-supervision for Training

For point-supervised action localization with spatio-temporal proposals, we start from the hypothesis that the proposals themselves, normally used for testing, can substitute the ground truth box annotations during training without a significant loss in performance. Our main goal is to mine out of a set of action proposals the best one during training while minimizing the annotation effort. The first level of supervision constitutes the action class label for the whole video. Given such global labels, we follow the traditional approach of mining the best proposals through multiple instance learning, as introduced for object detection by Cinbis et al. (2017). In the context of action localization, each video is interpreted as a bag and the proposals in each video are interpreted as its instances. The goal of Multiple Instance Learning is to train a classifier that selects the top proposals and separates proposals from different actions.

Next to the global action class label we leverage easy to obtain annotations within each video: we simply point at the action. Point-supervision allows us to easily exclude those proposals that have no overlap with any annotated point. Nevertheless, there are still many proposals that intersect with at least one point, as points do not uniquely identify each proposal. Therefore, we introduce an overlap measure to associate proposals with points. We also extend the objective of multiple instance learning to include the proposed overlap measure for proposal mining.

3.1 Overlap Between Proposals and Points

Let us first introduce the following notation. For a video V of F_V frames, an action tube proposal $A = \{BB_i\}_{i=f}^m$ consists of connected bounding boxes through video frames (f, \dots, m) where $1 \leq f < m \leq F_V$. Let \overline{BB}_i denote the center of a bounding box i . The point supervision $C = \{(c_i^{(x)}, c_i^{(y)})\}^K$ is a set of $K \leq F_V$ sub-sampled video frames where each frame i has a single annotated point $(c_i^{(x)}, c_i^{(y)})$. We propose an overlap measure that provides a continuous bounded score based on the match between a proposal and the point annotations.

Our overlap measure, inspired by a mild center-bias in annotators (Tseng et al. 2009), consists of two terms. The first term $M(\cdot)$ states how close the center of a bounding box from a proposal is to an annotated point, relative to the bounding box size, within the same frame. This center-bias term normalizes the distance of a point annotation to the center of a bounding box by the distance between the center and closest edge of the bounding box. For point annotation $(c_i^{(x)}, c_i^{(y)})$ and for bounding box \overline{BB}_{K_i} in the same frame, the score is 1 if the box center \overline{BB}_{K_i} is the same as the point annotation. The score decreases linearly in value as

the distance between the point annotation and the box center grows and the score becomes 0 if the point annotation is not contained in BB_{K_i} :

$$M(A, C) = \frac{1}{K} \sum_{i=1}^K \max \left(0, 1 - \frac{\|(c_i^{(x)}, c_i^{(y)}) - \overline{BB}_{K_i}\|_2}{\max_{(u,v) \in e(BB_{K_i})} \|((u, v) - \overline{BB}_{K_i})\|_2} \right). \quad (1)$$

In Eq. 1, (u, v) denotes the center point of each of the four edges of box BB_{K_i} , given by the function $e(BB_{K_i})$.

The second term $S(\cdot)$ serves as a form of regularization on the overall size of a proposal. The regularization aims to alleviate the bias of the first term towards large proposals, since large proposals are more likely to contain points and the box centers of large proposals are by default closer to the center of the video frames. Since actions are more likely to be in the center of videos (Tseng et al. 2009), the first term $M(\cdot)$ tends to be biased to large proposals. The size regularization term $S(\cdot)$ addresses this bias by penalizing proposals with large bounding boxes $|BB_i| \in A$, compared to the size of a video frame $|F_i| \in V$:

$$S(A, V) = \left(\frac{\sum_{i=f}^m |BB_i|}{\sum_{i=1}^{|V|} |F_i|} \right)^2, \quad (2)$$

where $|b| = (b(xmax) - b(xmin)) \cdot (b(ymax) - b(ymin))$ denotes the size of box b . Using the center-bias term $M(\cdot)$ regularized by $S(\cdot)$, our overlap measure $O(\cdot)$ is defined as

$$O(A, C, V) = M(A, C) - S(A, V). \quad (3)$$

Recall that A are the proposals, C captures the point-supervision and V the video. Overlap measure $O(\cdot)$ provides an estimation of the quality of the proposals during training and we use the measure in an iterative proposal mining algorithm over all training videos in search for the best proposals. In Fig. 2, we provide three visual examples of spatio-temporal proposals ranked based on our overlap measure.




Fig. 2 **a** No overlap. **b** Small overlap. **c** High overlap. Illustration of overlap between proposals and points

3.2 Mining for Proposals with Points

To mine spatio-temporal proposals, we are given a training set of videos $\{A_i, \mathbf{x}_i, y_i, C_i\}_{i=1}^N$, where the collection of tubes for the $|A_i|$ proposals is denoted as $A_i = \{\{B B_j\}_{j=f}^m\}^{|A_i|}$. Variable $\mathbf{x}_i \in \mathbb{R}^{|A_i| \times D}$ is the D dimensional feature representation of each proposal in video i . Annotations consist of the action class label y_i and the points C_i .

The proposal mining combines the use of the overlap measure $O(\cdot)$ of Eq. 3 with a Multiple Instance Learning optimization. The optimization aims to train a classification model that can separate good and bad proposals for a given action. We start from a standard MIL-SVM (Andrews et al. 2002; Cinbis et al. 2017) and adapt its objective to include a mining score $P(\cdot)$ of each proposal, which incorporates our function $O(\cdot)$ as:

$$\begin{aligned} \min_{\mathbf{w}, b, \xi} \quad & \frac{1}{2} \|\mathbf{w}\|^2 + \lambda \sum_i \xi_i, \quad \text{s.t.} \\ \forall_i : y_i \cdot (\mathbf{w} \cdot \arg \max_{\mathbf{z} \in \mathbf{x}_i} P(\mathbf{z}|\mathbf{w}, b, A_i^{(z)}, C_i, V_i) + b) & \geq 1 - \xi_i, \\ \forall_i : \xi_i & \geq 0, \end{aligned} \tag{4}$$

where (\mathbf{w}, b) denote the classifier parameters, ξ_i denotes the slack variable and λ denotes the regularization parameter. Variable $\mathbf{z} \in \mathbf{x}_i$ denotes the representation of a single proposal in the set of all proposals \mathbf{x}_i for training video i . Variable $A_i^{(z)}$ denotes the tube corresponding to proposal representation \mathbf{z} . The proposal with the highest mining score per video is used to train the classifier.

Different from standard MIL-SVM, the proposals are not only conditioned on the classifier parameters, but also on the overlap scores from the point annotations. In other words, the standard maximum likelihood optimization of MIL is adapted to include point overlap scores that serve as a prior on the individual proposals. The objective of Eq. 4 is non-convex due to the joint minimization over the classifier parameters (\mathbf{w}, b) and the maximization over the mined proposals $P(\cdot)$. Therefore, we perform iterative block coordinate descent by alternating between clamping one and optimizing the other. Given a fixed selection of proposals, the optimization becomes a standard SVM optimization over the features of selected proposals (Cortes and Vapnik 1995). For fixed model parameters, the maximization over the proposals is determined by scoring proposals as:

$$P(\mathbf{z}|\mathbf{w}, b, A_i^{(z)}, C_i, V_i) \propto ((\mathbf{w}, \mathbf{z}) + b) + O(A_i^{(z)}, C_i, V_i). \tag{5}$$

In Eq. 5, the score of a proposal is the sum of two components, namely the score of the current model and the overlap with

the point annotations in the corresponding training video. The mining and classifier optimizations are alternated for a fixed amount of iterations. After the iterative optimization, a final SVM is trained on the best mined proposals. Identical to approaches using box-supervision, our model selects the best proposals from test videos, without requiring any box annotations during training.

4 Pseudo-pointing for Inference

Inference is typically achieved through a maximum likelihood over all proposals in a test video. However, relying on a maximum likelihood estimate of the model is rather limited, as it only relies on the features of the proposals. We show that visual cues within the test videos help to guide the selection proposals during inference, similar to how point annotations provide guidance during training. We dub these automatic cues pseudo-points and investigate five of them. The pseudo-points rely on training point annotations, self-supervision, person detection, independent motion, and center bias. We show how to exploit and combine these pseudo-points to improve the action localization during inference. Lastly, we also provide two forms of regularization to further boost the localization results.

4.1 The Pseudo-points

In Fig. 3, we provide a visual overview of the visual cues for multiple video frames. Next, we outline each pseudo-point individually.

4.1.1 Training Point Statistics

The first pseudo-point focuses on the point annotations provided during training. Intuitively, actions do not occur at random locations in video frames. Recall that we are given N training videos, where y_i, C_i denote respectively the video label and point annotation of training video i . We exploit this observation by making a pseudo-point for an action class \mathcal{Y} as follows:

$$p_{\text{points}}(F, \mathcal{Y}) = \frac{1}{\sum_{i=1}^N \mathbb{I}[y_i = \mathcal{Y}]} \sum_{i=1}^N \mathbb{I}[y_i = \mathcal{Y}] \cdot C_i. \tag{6}$$

The above Equation states that for an action \mathcal{Y} , the pseudo-point in a test video is determined as the average point annotation location given the training videos of the same action. The reasoning behind this pseudo-point is that specific actions tend to re-occur in similar locations across videos. Note that the pseudo-point is independent of the frame F itself and only depends on the training point statistics.

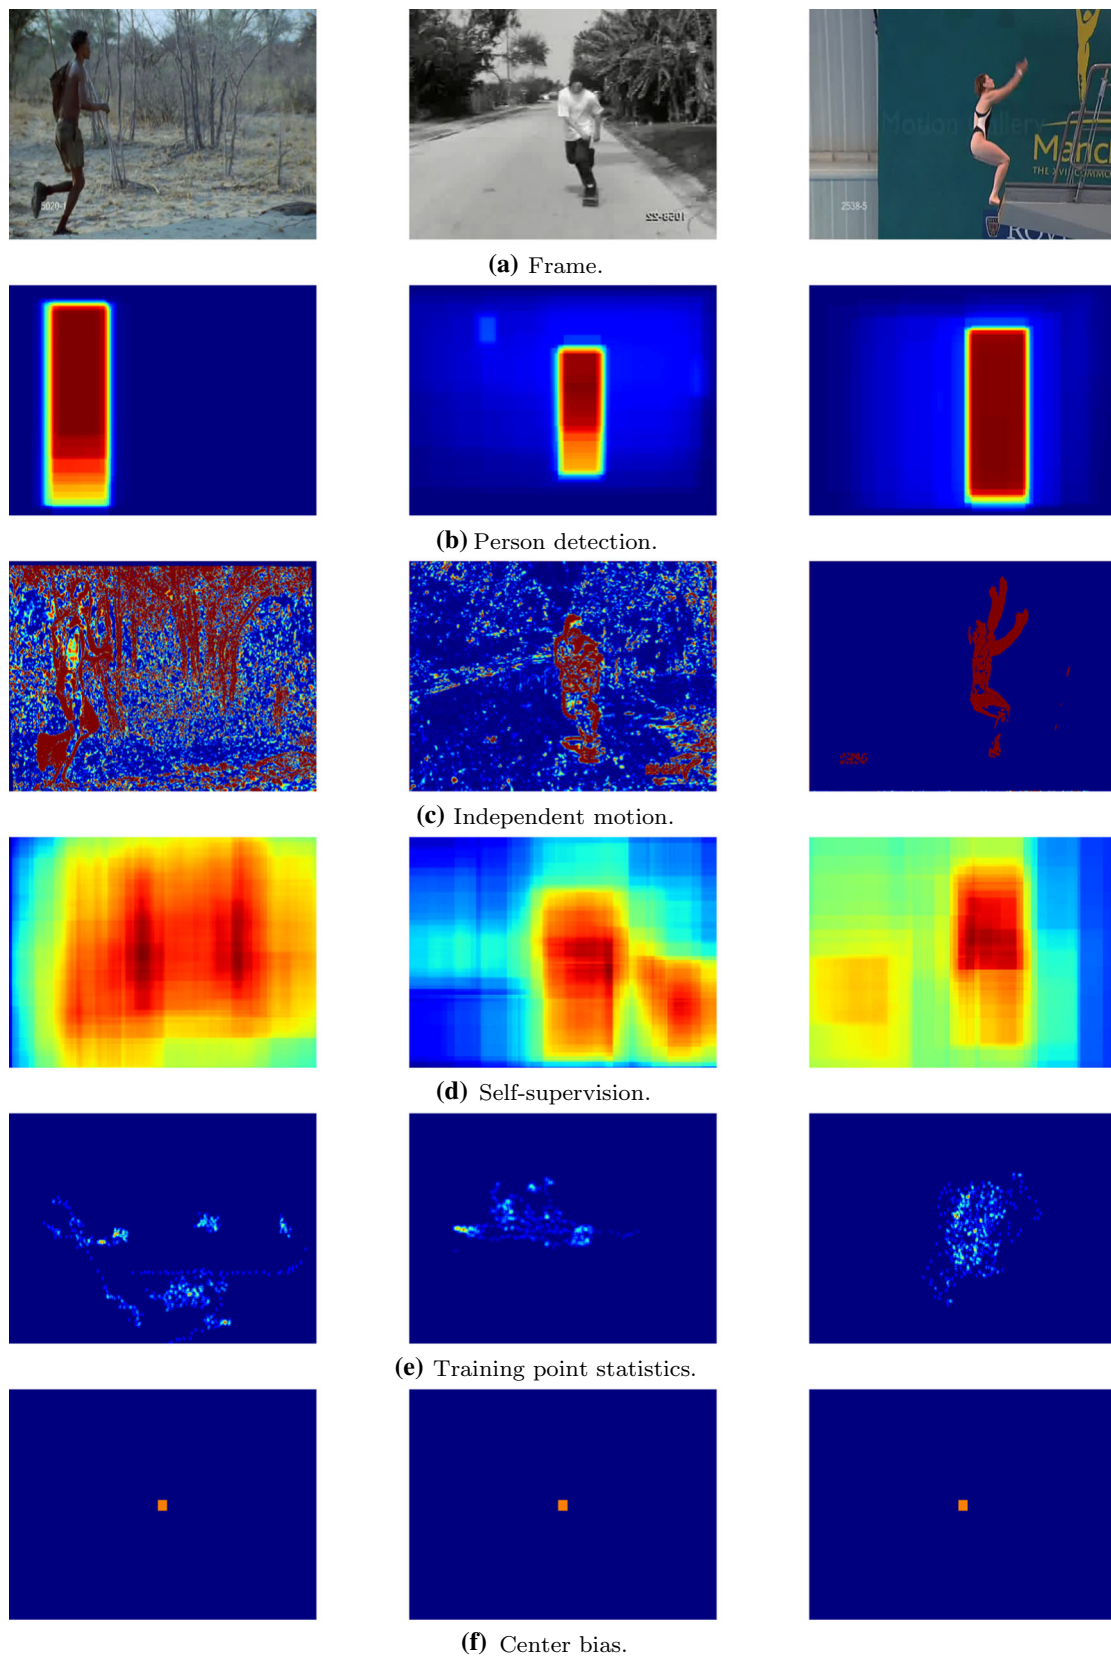


Fig. 3 Pseudo-points for an example frame (a) from three videos showing *running*, *skateboarding*, and *diving*. The pseudo-points derived from (b) person detection, (c) independent motion, and (d) self-supervision

focus on the primary action in the video. The pseudo-points derived from (d) training points and (e) center bias provide data-independent prior statistics to steer better proposal selection during inference

4.1.2 Self-Supervision from Proposals

The second pseudo-point we investigate does not require external information; the pseudo-point relies on the spatio-temporal proposals themselves. The main idea behind this pseudo-point is that the distribution over all proposals in a single frame provides information about its action uncertainty. It relies on the following assumption: the more the proposals are on the same spatial location, the higher the likelihood that the action occurs in that location. The pseudo-point can be seen as a form of self-supervision (Doersch et al. 2015; Fernando et al. 2017), since it provides an automatic annotation from proposals to guide the selection of the very same proposals.

More formally, for test video t , let A_t denote the spatio-temporal proposals. Furthermore, let F denote a video frame and let $C_{A_t}^*(u, v, F)$ denote the number of proposals that contain pixel (u, v) in F . We place an automatic pseudo-point at the center of mass over these pixel counts:

$$p_{\text{self}}(F) = \frac{1}{\sum_{u,v} C_{A_t}^*(x, y, F)} \sum_{(u,v) \in F} C_{A_t}^*(u, v, F) \cdot (u, v). \quad (7)$$

The function of Eq. 7 outputs a 2D coordinate in frame F , representing the center of mass over all pixels in F , with the mass of each pixel (u, v) given by $C_{A_t}^*(u, v, F)$. The 2D output coordinate will serve as the pseudo-point in frame F .

4.1.3 Person Detection

The third pseudo-point follows earlier work on action localization by incorporating knowledge about person detections (Siva and Xiang 2011; Yu and Yuan 2015). Actions are typically person-oriented, so the presence or absence of a person in a proposal provides valuable information. Here, we employ a Faster R-CNN network (Ren et al. 2015), pre-trained on MS-COCO (Lin et al. 2014), and use the person class for the detections. This results in roughly 50 box detections per frame after non-maximum suppression. We select

the box in each frame with the maximum confidence score as the automatic pseudo-point.

4.1.4 Independent Motion

The independent motion of a pixel (u, v) in frame F provides information as to where foreground actions are occurring. More precisely, independent motion states the deviation from the global motion in a frame (Jain et al. 2017). Let $C_I^*(u, v, F) \in [0, 1]$ denote the inverse of the residual in the global motion estimation at pixel (u, v) in frame F . The higher $C_I^*(u, v, F)$, the less likely it is the pixel contributes to the global motion. Akin to the second pseudo-point, we place an automatic pseudo-point at the center of mass, now over the independent motion estimates:

$$p_{\text{motion}}(F) = \frac{1}{\sum_{u,v} C_I^*(u, v, F)} \sum_{u,v} C_I^*(u, v, F) \cdot (u, v). \quad (8)$$

Equation 8 outputs a 2D coordinate, but now using the independent motion as mass for each pixel in F .

4.1.5 Direct Center Bias

Lastly, we again focus on an observation made during training; actions and annotators have a bias towards the center of the video (Tseng et al. 2009). We exploit this bias directly in our fifth pseudo-point by simply placing a point on the center of each frame:

$$p_{\text{center}}(F) = (F_W/2, F_H/2), \quad (9)$$

where F_W and F_H denote the width and height of frame F respectively.

Figure 4 provides the spatio-temporal evolution and focus area of the pseudo-points for four example videos.




Fig. 4 Pseudo-point visualization on four example videos for training points, center bias, self-supervision, independent motion, and person detection (depicted as points for visualization). In general, the pseudo-

points are present around the action or even follow the action. When actions are not in the frame however, as shown in the right example, pseudo-points may place automatic annotations in phantom positions

4.2 Exploiting Pseudo-points

4.2.1 Rescoring Test Proposals

Given a test video t , the standard approach in action localization with spatio-temporal proposal for finding the best proposal is done from a set of proposals $\{A_t i\}_{i=1}^{|A_t|}$, given a model (\mathbf{w}, b) is given as:

$$t_{\star} = \arg \max_{i=1, \dots, |A_t|} ((\mathbf{w}, \mathbf{z}_i) + b). \quad (10)$$

We exploit the pseudo-points to adjust this likelihood estimate:

$$t_{\star} = \arg \max_{i=1, \dots, |A_t|} ((\mathbf{w}, \mathbf{z}_i) + b) + \lambda_P \cdot O(A_t i, P, t_V), \quad (11)$$

where P denotes the pseudo-point of interest and t_V denotes the test video itself.

The above equation is similar in spirit to Eq. 5, but now automatic cues are employed, rather than manual point annotations. Adjusting the proposal selection using pseudo-points during testing can be seen as a form of regularization. Rather than a single-point maximum likelihood given a trained model, we add continuous restrictions on the proposals based on their match with automatic pseudo-points, which aid the selection towards proposals with a high overlap to the ground truth action location.

4.2.2 Weighting and Selecting Pseudo-points

Intuitively, not all pseudo-points are equally effective, stating the need for the weights in Eq. 11. However, setting proper values for λ_P can not be done directly through standard (cross-)validation, as this requires box-supervision. To overcome both problems, we provide a score function to estimate the quality of each pseudo-point. This score will be used to both determine which pseudo-point is most favourable to select and directly serve as weighting value in Eq. 11.

The score function for the person detection pseudo-point (the only pseudo-point that outputs boxes), is identical to the overlap function in Eq. 1. This entails that if the center of the top person detection in each frame of a training video matches with the point annotations of the same video, a high score is achieved. We compute the average match over all training videos as the weight (λ_P) for person detection. For the other pseudo-points, we are only given points. In these cases, we use the distance to the nearest image border to normalize the distance between the manual point annotation and the automatic pseudo-point annotation. The overall score function is computed in identical fashion as for the person detection.

By matching automatic pseudo-points in training videos with the manual point annotations, we arrive at an automatic quality measure for pseudo-points, which can be used to weight and select pseudo-points.

4.3 Temporal Pseudo-points

Besides knowing where specific actions occur spatially over a complete dataset, the temporal extent of actions is also helpful for proposal selection. Here, we provide a temporal pseudo-point, again relying on training point statistics. For action \mathcal{Y} , we retrieve its temporal extent by comparing the temporal span of point annotations to the temporal extent of the videos in which the actions occur. The fraction of the annotation span relative to the video length is computed for each action instance and averaged over a complete dataset. Let $\mathcal{F}_{\mathcal{Y}}$ denote the average temporal length of action \mathcal{Y} and let $\mathcal{F}_{t_{jk}}$ denote the temporal length of proposal k in test video j . Then we compute this temporal pseudo-point as:

$$s_{\mathcal{Y}}^{\star} = \arg \max_{k \in \{1, \dots, |\mathbf{x}|\}} \left(((\mathbf{w}_j, \mathbf{x}_k) + b_j) - \lambda_T \cdot \frac{|\mathcal{F}_{\mathcal{Y}} - \mathcal{F}_{t_{jk}}|}{\mathcal{F}_{\mathcal{Y}}} \right). \quad (12)$$

In Eq. 12, the match between the temporal pseudo-point of an action and the temporal extent of a proposal also acts as a regularization. The better the match, the lower the penalty in the likelihood, resulting in a better selection of proposals.

With pseudo-points, we are able to guide the selection of the top proposal per action per video for action localization, akin to how point-supervision is used during training. Having defined the complete point-supervised regime for training and inference we are now ready for the experiments.

5 Experimental Setup

5.1 Datasets

5.1.1 UCF Sports

The UCF Sports dataset consists of 150 videos from 10 sport action categories, such as *Skateboarding*, *Horse riding*, and *Walking* (Rodriguez et al. 2008). We employ the train/test split suggested in (Lan et al. 2011). Example frames from the dataset are shown in Fig. 5a.

5.1.2 UCF-101

The UCF-101 dataset consists of 13,320 videos from 101 action categories, such as *Skiing*, *Basketball dunk*, and *Surfing* (Soomro et al. 2012). For a subset of 3,204 videos



Fig. 5 Example frames of the UCF Sports, UCF-101, and Hollywood2Tubes dataset with box annotations. Different from the UCF Sports and UCF-101 datasets, Hollywood2Tubes provides new chal-

lenges for spatio-temporal action localization, due to large occlusions, small inter-class difference, and large action size variation. **a** UCF sports. **b** UCF-101. **c** Hollywood2Tubes

and 24 categories, spatio-temporal annotations are provided (Soomro et al. 2012). We will use this subset in the experiments and use the first train/test split suggested in (Soomro et al. 2012). In Fig. 5b, we show dataset example frames.

5.1.3 Hollywood2Tubes

The Hollywood2Tubes dataset consists of 1,707 videos from 12 action categories (Mettes et al. 2016), such as *getting out of a car*, *sitting down*, and *eating*. The dataset is derived from the Hollywood2 dataset (Marszałek et al. 2009), with point annotations for training and box annotations for evaluation. Different from current action localization datasets, Hollywood2Tubes is multi-label and actions can be multi-shot, i.e., can span over multiple non-continuous shot, adding new challenges for action localization. We show an example frame with box annotations for each of the 12 actions in the dataset in Fig. 5c. Annotations are available at <http://tinyurl.com/hollywood2tubes>.

5.2 Implementation Details

5.2.1 Proposals

Our proposal mining algorithm is agnostic to the underlying spatio-temporal proposal algorithm. Through this work, we employ the unsupervised APT proposals (van Gemert et al. 2015), since the algorithm provides high action recall, is fast to execute, and the code is publicly available. For each proposal, we extract Improved Dense Trajectories and compute HOG, HOF, Traj, and MBH features (Wang and Schmid 2013). The combined features are concatenated and aggregated into a fixed-size representation using Fisher Vectors (Sánchez et al. 2013). We construct a codebook of 128 clusters, resulting in a 54,656-dimensional representation per proposal. The same proposals and representations are also used in (van Gemert et al. 2015; Mettes et al. 2016, 2017) allowing for a fair comparison.

5.2.2 Training

The proposal mining is performed for 5 iterations; more iterations have little effect on performance. Following the suggestions of Cinbis et al. (2017), the training videos are randomly split into 3 splits to train and select the proposals. For training a classification for one action, 100 proposals of each video are randomly sampled from the other actions as negatives. The regularization parameter λ in the max-margin optimization is fixed to 10 throughout the experiments.

5.2.3 Evaluation

For an action, we select the top scoring proposal for each test video given the trained model. To evaluate the action localization performance, we compute the Intersection-over-Union (IoU) between proposal p_1 and the box annotations of the corresponding test example p_2 as: $iou(p_1, p_2) = \frac{1}{|\Gamma|} \sum_{f \in \Gamma} IoU_{p_1, p_2}(f)$, where Γ is the set of frames where at least one of (p_1, p_2) is present. The function IoU states the box overlap within a specified frame. For IoU threshold τ , a top selected proposal is deemed a positive detection if $iou(p_1, p_2) \geq \tau$. After aggregating the top tubes from all videos, we compute either the Average Precision score or AUC using the proposal scores and positive/negative detection labels.

6 Results

6.1 Action Localization with Point-supervision

6.1.1 Setup

In the first experiment, we evaluate our main notion of localizing actions using point-supervision. We perform this

evaluation on UCF Sports and UCF-101. We compare our approach to the following three baselines:

- *box-supervision* This baseline follows the training protocol of van Gemert et al. (2015), where for each action, a classifier is trained using the features from ground truth boxes. Additionally, spatio-temporal proposals with an overlap higher than 0.6 and lower than 0.1 are added as positives and negatives, respectively.
- *best proposal* This baseline trains an action localizer using the spatio-temporal proposal with the highest overlap to the ground truth box tube.
- *video label (ours)* This baseline employs MIL optimization with a uniform prior, i.e., only video labels are used as annotations. This baseline is inspired by (Cinbis et al. 2017), but performed on action proposals in videos instead of object proposals in images.

Unless stated otherwise, we employ the centers of the original box annotations on UCF Sports and UCF-101 as the point annotations throughout our experiments.

6.1.2 Results

The results on UCF Sports and UCF-101 are shown in Fig. 6a (mean average precision) and Fig. 6b (AUC). We first observe that traditional box-supervision yield identical results to using the best possible spatio-temporal proposal. This result validates our starting hypothesis that spatio-temporal proposals provide viable training examples. Second, we observe that across both datasets and all overlap thresholds, point-supervision performs similar to both the box-supervision and best proposal approaches. This result highlights the effectiveness of point-supervision for action localization. With point-supervised action localization we no longer require

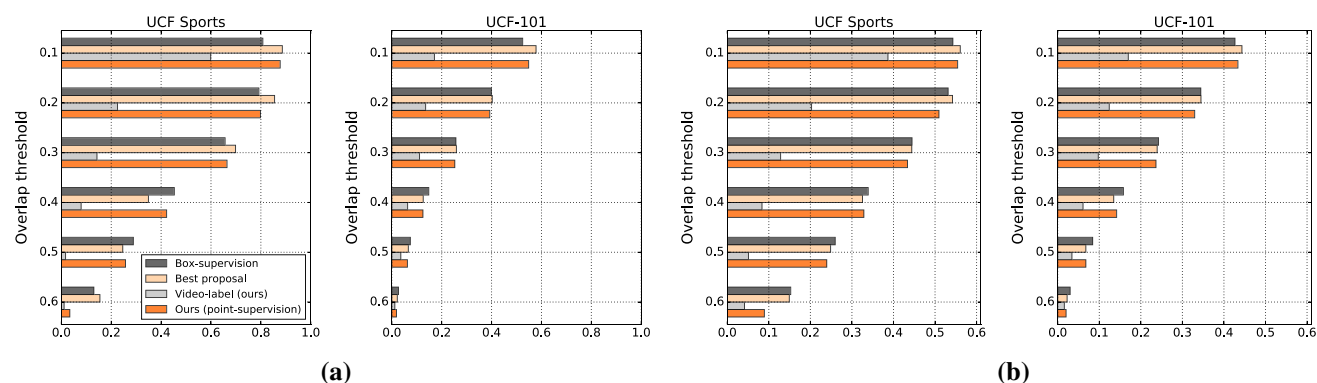


Fig. 6 **a** Mean average precision. **b** AUC. Localization results on UCF Sports and UCF-101 using box-supervision, point-supervision, and video-labels only. Across both datasets and all overlap thresholds, point-supervision is as effective as box-supervision, while they both

outperform video-label supervision. We conclude that spatial annotations are vital and that points provide sufficient support for effective localization

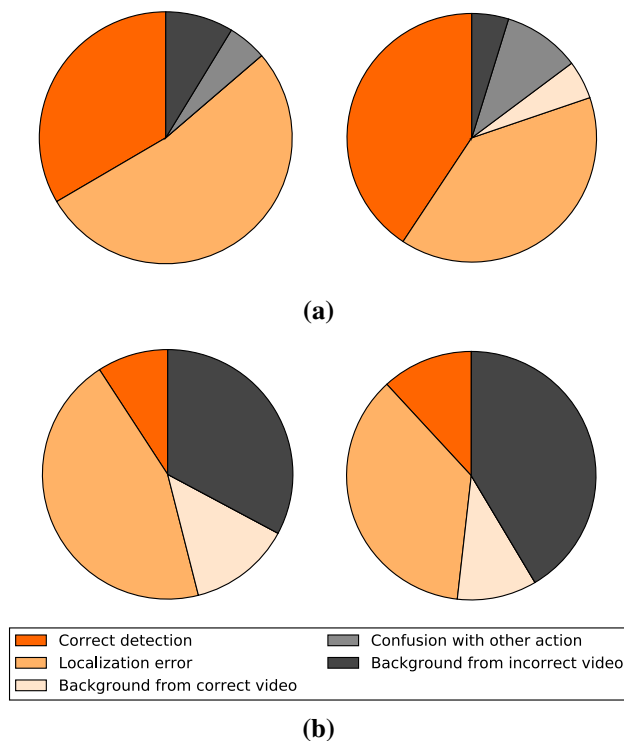


Fig. 7 **a** UCF Sports. **b** UCF-101. Action localization error diagnosis on UCF Sports and UCF-101 when using point-supervision (left) and box-supervision (right). On both datasets, we observe that averaged over all actions, approaches using point- and box-supervision yield similar error type distributions, explaining their similar localization performance

expensive box annotations. As results using video-labels only are limited compared to points, we conclude that points provide vital information about the spatial location of actions.

6.1.3 Error Analysis

To gain insight into why point supervision is effective for action localization, we perform an error diagnosis and corresponding qualitative analysis. We perform the diagnosis on the approaches using box-supervision and point-supervision. Akin to error diagnosis for object detection (Hoiem et al. 2012), we quantify the types of errors made by each localization approach. We take the top R detections for each action, where R is equal to the number of ground truth instances in the test set. We categorize each detection into five classes relevant for action localization: (1) correct detection, (2) localization error, (3) confusion with other action, (4) background from video containing the action, and (5) background from video not containing the action. The categorization definition is provided in “Appendix A”.

The error diagnosis, averaged over all actions, is shown in Fig. 7 for UCF Sports and UCF-101. We observe that overall, the types of errors made by both approaches are similar. The predominant error type is localization error, which

means that proposals from positive videos with a low match to the ground truth are the main errors. Proposals from background proposals of both positive and negative videos are hardly ranked high. Overall, using boxes and points result in similar errors, which matches with their similarity in localization performance. A common limitation is the quality of the spatio-temporal proposals themselves; only few proposals have a high overlap with the ground truth, making the localization a needle in the haystack problem regardless of the model. On UCF-101, a large part of the errors also comes from confusion with the background from other videos. This is because the UCF-101 dataset can have more than one instance of the same action in each video. If such additional instances are missed, non-distinct regions of negative videos are automatically ranked higher.

6.1.4 Qualitative Analysis

In Fig. 8, we provide qualitative results on UCF Sports and UCF-101. The results show where point- and box-supervision yield similar and dissimilar action tubes. For simple actions like *walking* and *soccer juggling*, both approaches yield (near-)identical results. For actions such as *skateboarding* and *walking with a dog*, we observe that point-supervision tends to focus on the invariant object (here: skateboard, horse, and dog), since these are distinctive elements for the action. This is because the spatial extent of actions is no longer known with points, which means that the extent is learned from examples, rather than from manual annotations. We also note that limitations in the model can result in different results, as shown in the leftmost example of Fig. 8b.

6.1.5 Conclusions

From the localization results, error diagnosis, and qualitative analysis, we make the following conclusions: (i) point-supervision yield results comparable to full box-supervision for action localization, (ii) averaged over all actions, the approaches using box and point annotations have approximately similar error type distributions, and (iii) models learned with point-supervision learn the spatial extent of actions discriminatively from examples.

6.2 Influence of Spatio-temporal Proposal Quality

In the second experiment, we evaluate the influence of the spatio-temporal proposals upon which our approach is built. Spatio-temporal proposals optimize recall, i.e., for a video, *at least* one proposal should have a high overlap to the ground truth action localization. An inconvenient side-effect from this requirement is that each video outputs many proposals that have a low overlap, making the selection of the best proposal during testing a needle in the haystack problem.

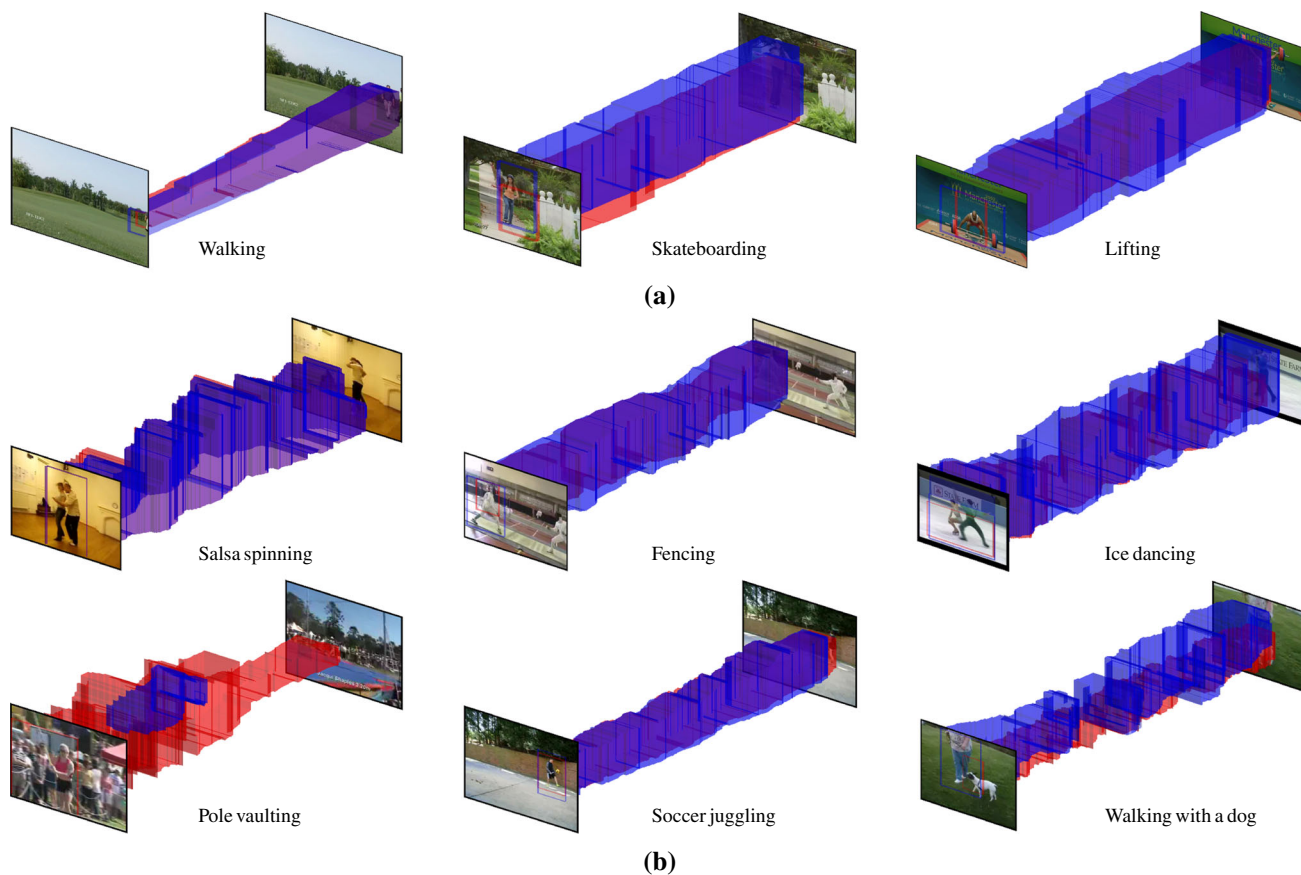


Fig. 8 Qualitative results on (a) UCF Sports and (b) UCF-101 of selected proposals using point- (red) and box-supervision (blue). For simple actions with static backgrounds, such as *walking*, *salsa spinning* and *soccer juggling*, both approaches converge to similar locations. For actions with a more dynamic background and interacting objects, point-

supervision might lead to a selection of different proposal locations. Examples include *fencing* and *walking with a dog*. We conclude that point- and box-supervision typically leads to similar results, although point-supervision tends to focus more on the most discriminative element of actions (Color figure online)

This problem was observed in the error diagnosis of the first experiment.

Here, we investigate the influence of the high ratio of proposals with a low overlap during training and testing. During both training and testing, we add the oracle ground truth tube to the proposals. We furthermore add a parameter ϵ , which controls the fraction of proposals with an overlap below 0.5. We train several models with varying values for ϵ . We evaluate this oracle experiment on UCF Sports.

6.2.1 Results

The localization performance for several values of ϵ is shown in Fig. 9. The baseline is the result achieved in the first experiment. From the Figure, it is evident that removing low quality proposals positively affects the localization performance. However, a large portion of low quality proposals need to be removed to achieve better results. This is because of the large amount of low quality proposals. On UCF Sports, only 7% of the proposals have an overlap of at least 0.5 (!). This means that when 50% of the low quality proposals are




Fig. 9 Influence of spatio-temporal proposal quality on UCF-Sports. The baseline corresponds to the result from the first experiment. For the others, the ground truth location is added as one of the proposals during testing. Where, ϵ states the fraction of low quality (overlap ≤ 0.5) proposals that are removed. Action localization performance increases when large amounts of low quality proposals are removed. We conclude that better quality action proposals will have a positive impact on pointily-supervised action localization

Table 1 Action localization performance on UCF-101 as a function of the annotation stride for point-supervision, compared to box-supervision

	Box supervision	Point supervision						
		Annotation stride						
		1	2	5	10	20	50	100
mAP@0.2	0.399	0.393	0.404	0.389	0.384	0.395	0.379	0.371
mAP@0.5	0.074	0.063	0.060	0.068	0.064	0.061	0.064	0.053
Annotation speed-up	1.0	9.8	19.3	46.0	85.0	147.6	264.6	359.6

The annotation-speedup is relative to a box annotation in each frame. Fewer point annotations result in large annotation-time speed-ups, while the performance is hardly affected
 Bold values indicate the importance of the obtained speed-up

removed, the ratio of low to high quality proposals is still 6 to 1. When removing 50% of the low quality proposals, the result increases from 0.23 to 0.32. This further increases to 0.49 when removing 95% of the low quality proposals. When only using the ground truth tube and high overlapping proposals (i.e., $\epsilon = 1.0$), we achieve a performance of 0.90, indicating the large gap between current performance and the upper bound given the current set of features. We conclude from this experiment that for point-supervised action localization with spatio-temporal proposals, a limiting factor is the quality of the proposals themselves. With better action proposals, point-supervision can achieve even better results.

6.3 Sparse Point Annotations

In the third experiment, we evaluate: (i) how much faster is point-supervision compared to box-supervision and (ii) how many point annotations are sufficient for effective localization. Intuitively, point-supervision is not required for every frame, since the amount of change between consecutive frames is small. We evaluate the influence of the annotation stride and we also estimate how much faster the annotation process becomes compared to dense box-supervision. We perform this experiment on UCF-101, since the videos in this dataset are the longest for action localization, allowing for a wide range of annotation strides.

6.3.1 Annotation Times

To obtain an estimate of the annotation times for box- and point-supervision, we have re-annotated several hundreds of videos while keeping track of the annotation times. We found that annotating a video with an action label takes roughly 5 s. Furthermore, annotating a box in a frame takes roughly 15 s. This estimate is in between the estimate for image annotation of Su et al. (2012) (roughly 30 s) and the estimate of Russakovsky et al. (2015) (10 to 12 s). Annotating a point takes roughly 1.5 s, making points ten times faster than boxes to annotate. This estimate is in line with point annotations in images (0.9–2.4 s (Bearman et al. 2016)).

6.3.2 Results

In Table 1, we provide the localization performance for two overlap thresholds and the annotation speed-up for point supervision at seven annotation strides. The Table shows that when annotating fewer frames, performance is retained. Only when annotating fewer than 5% of the frames (i.e., for a stride larger than 20), the performance drops marginally. This result shows that our approach is robust to sparse annotation, a point at every frame is not required. The bottom row of the Table shows the corresponding speed-up in annotation time compared to box-supervision. An almost 50-fold speed-up can be achieved while maintaining comparable localization performance. A 300–500 fold speed-up can be attained with a marginal drop in localization performance. We conclude that point-supervision is robust to sparse annotations, opening up the possibility for further reductions in annotation cost for action localization.

6.4 Noisy Point Annotations

Human annotators, while center biased (Tseng et al. 2009), do not always precisely pinpoint center locations while annotating (Bearman et al. 2016). In the fourth experiment, we evaluate how robust the action localization performance is with respect to noise in the point-supervision. We start from the original point annotations and add zero-mean Gaussian noise with varying levels of isotropic variance. This experiment is performed on the UCF-101 dataset.

6.4.1 Results

The localization performance for six levels of annotation noise is shown in Fig. 10. The performance for $\sigma = 0$ corresponds to the performance of point-supervision in the first experiment. We observe that across all overlap thresholds, the localization performance is unaffected for noise variations up to a σ of 5. For $\sigma = 10$, the results are only affected for thresholds of 0.3 and 0.4, highlighting the robustness of our approach to annotation noise. For large noise variations ($\sigma = 50$ or 100), we observe a modest drop in performance




Fig. 10 Localization performance on UCF-101 for various levels of noise in the point annotations. Up to a noise deviation of 10 pixels can be handled robustly. For large deviations (50 pixels an up), performance drops for lower overlap thresholds. Point-supervision can accommodate human error in the point annotations up to 10 pixels

for the overlaps thresholds 0.1 to 0.4. We conclude that points do not need to be annotated precisely at the center of actions. Annotating points in the vicinity of the action is sufficient for action localization.

6.5 Exploiting Pseudo-points

In the fifth experiment, we investigate the effect of each of the pseudo-points on the action localization performance during inference. We perform this experiment on both UCF Sports and UCF-101.

6.5.1 Pseudo-point weights

To utilize the pseudo-points effectively during inference and to know a priori which pseudo-point is most effective, we compute the weight per pseudo-point as outlined in Sect. 4.2.2. This has resulted in the following values:

1. Person detection: $\lambda_P = 0.76$.
2. Independent motion: $\lambda_P = 0.57$.
3. Center bias: $\lambda_P = 0.48$.
4. Self-supervision: $\lambda_P = 0.32$.
5. Training points: $\lambda_P = 0.25$.

The weights computed based on the match with point-supervision in training videos provide the degree to which each pseudo-point should contribute to the selection of spatio-temporal proposals in test videos and they also provide a measure to select the best pseudo-point.

6.5.2 Results

In Fig. 11, we show the localization performance for overlap thresholds of 0.2 and 0.5. On UCF Sports, we observe for an overlap of 0.2, the performance improves for training points, self-supervision, and person detection. For center bias and independent motion, there is a minimal drop in performance. For an overlap of 0.5, the results diverge more clearly. Independent motion (+6.7%), self-supervision (+6.8%), and person detection (+20.0%) benefit directly from inclusion. This does not hold for center bias and training points. While pseudo-points can have a positive impact on the performance, it is not effective for all types of pseudo-points. Discovering which pseudo-points are most effective is a necessity. On UCF-101, we observe similar trends. Person detection and self-supervision yield increased localization performance, while the data independent center bias and train points have a negative effect. Video-specific visual cues, such as persons and motion, are effective; generic statistics less so.

We observe that the order of the weights correlates with the localization performance. This indicates the effectiveness of the proposed weighting function, as it provides insight into the quality of the pseudo-points without having to evaluate their performance at test time. Person detection is the most effective pseudo-point. The center bias and training points score lower, which is also visible in their localization performance. Only the self-supervision scores low, while it has a positive effect on the localization. We conclude that the proposed pseudo-point weighting is a reliable way to determine the effectiveness of pseudo-points with point-supervision.

On both datasets, we also investigate the effect of temporal pseudo-points on the localization performance. On UCF Sports, we observe an increase in performance for both overlap thresholds (+7.3% at 0.2, +2.4% at 0.5). On UCF-101, we also observe a positive effect, albeit with smaller improvements (+0.8% at 0.2, +0.6% at 0.5). We conclude that regularizing spatio-temporal proposals using the temporal extent of actions, which is provided by point-supervision, aids action localization in videos.

Based on the weights of the pseudo-annotations, we recommend to use person detection during inference. We will use this setup for the state-of-the-art comparison.

6.5.3 Qualitative Results

To gain insight into which types of videos benefit from pseudo-points, we provide a qualitative analysis on UCF Sports. In Fig. 12a, we show three test videos where the localization improved due to the effect of pseudo-points. We show the effect for the independent motion, self-supervision, and person detection respectively. In all three cases, the inclusion of pseudo-points resulted in a better fit on the action by enlarging its scope. For self-supervision and independent

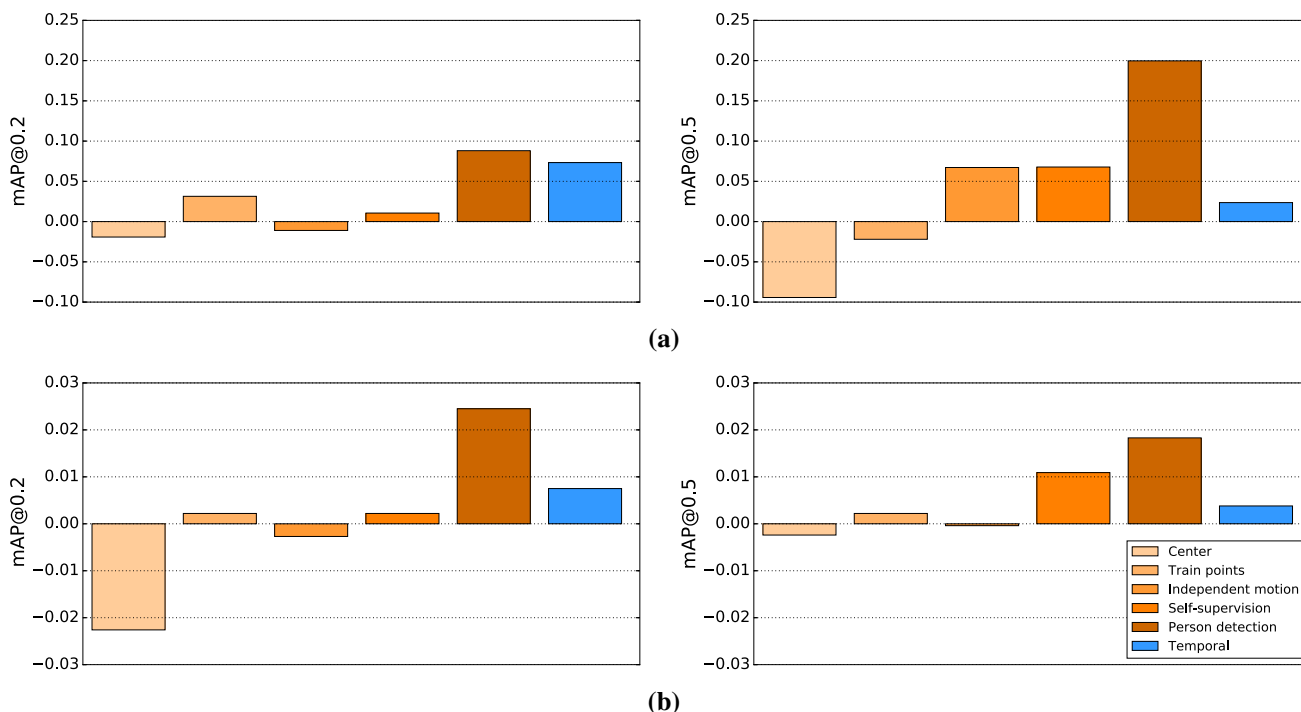


Fig. 11 **a** UCF Sports. **b** UCF-101. The effect of pseudo-points during inference for action localization on UCF Sports and UCF-101 at an overlap of 0.2 (left column) and 0.5 (right column). All results are provided relative to the performance without pseudo-points. Data-dependent pseudo-points such as person detection, self-supervision, and independent motion have a positive effect on the localization perfor-

mance. Data-independent pseudo-points such as center bias and training points are not effective for action localization. Incorporating the temporal extent of actions as a pseudo-point can further boost performance. We conclude that pseudo-points, when chosen correctly, aid action localization performance

motion, the wider motion evidence resulted in a better fitting localization. For person detection, the evidence of the whole person had a positive effect. These examples show the potential of pseudo-points to guide the selection of spatio-temporal proposals during inference.

In Fig. 12b, we show three test videos where the inclusion of pseudo-points resulted in a worse localization. We show this effect for the less successful center bias and training points, as well as for the most successful pseudo-point person detection. For center bias, this resulted in a shift from precise fit on the action (red) to a large generic location (blue). This is because the center bias is data independent and might undo correct localizations. This also holds for the training point statistics, which are identical for each test video. Lastly, the person detection can yield diverging localizations when many people are present in the scene. We conclude that motion-based and person-based pseudo-points can aid action localization, while data independent pseudo-points are less suited.

6.6 Comparison to Others

In our final experiment, we compare pointly-supervised action localization with alternatives using either box-

supervision, or weaker forms of supervision. We perform this experiment on UCF Sports, UCF-101, and Hollywood2Tubes. To compare with as many methods as possible, we evaluate with AUC on UCF Sports and with mAP on UCF-101 and Hollywood2Tubes. We evaluate action localization for the standard overlap threshold of 0.2.

6.6.1 Results

We present results on all three datasets In Table 2. On UCF Sports, we observe that our approach outperforms the state-of-the-art in weakly-supervised action localization, as well as the point-supervision in our previous work (Mettes et al. 2016) which lacks the pseudo-points during inference.

Naturally, our pointly-supervised approach also outperforms the state-of-the-art in zero-shot and unsupervised action localization, emphasizing the effectiveness of points as supervision. Lastly, we perform competitive or even better than the state-of-the-art using box-supervision.¹

On UCF-101, we also outperform all existing weakly-supervised alternatives. Our approach reaches an mAP of 0.418, compared to 0.369 of Li et al. (2018) and 0.351 of

¹ Note that we only use the top-1 proposal per action per video, as more proposals per video skews AUC performance (Weinzaepfel et al. 2015).




Fig. 12 Qualitative analysis of the effect of pseudo-points for action localization on UCF Sports. Red indicates the localization without pseudo-points, blue with pseudo-points. **(a)** The top row show three examples where pseudo-points improve the localization. For the self-supervision and independent motion examples, motion information helped to widen the scope of the action. For person detection, information from the whole person enhances the scope of the action. **(b)** The second row shows failure cases. For data independent pseudo-points

such as center bias and training point statistics, the action can deviate from its true location. Person detection can furthermore be problematic when many people are present in the scene. The qualitative analysis shows that pseudo-points can alleviate problems in point-supervised action localization. To aid the performance, data-dependent pseudo-points are informative, while data-independent pseudo-points appear less effective (Color figure online)

Mettes et al. (2017), the state-of-the-art in weakly-supervised action localization. In comparison to box-supervision, we outperform the approach of van Gemert et al. (2015), which employs identical spatio-temporal proposals and representations. On UCF-101, the state-of-the-art approaches in action localization from box-supervision perform better (Kalogiton et al. 2017a; Yang et al. 2017). These approaches score and link 1 to 10 consecutive boxes into tubes, rather than opting for spatio-temporal proposals. Based on our second experiment, we posit that better spatio-temporal proposals can narrow this gap in performance.

Lastly, we provide results with our approach on Hollywood2Tubes. We first observe that overall, the performance on this dataset is lower than on UCF Sports and UCF-101 in terms of mAP scores. This indicates the challenging nature of the dataset. The combination of temporally untrimmed videos, multi-shot actions, and actions of complex semantic nature make for a difficult action localization. Our approach provides a new state-of-the-art result in this dataset with an mAP of 0.178, compared to 0.143 of (Mettes et al. 2016) and 0.172 of (Mettes et al. 2017).

7 Conclusions

This paper introduces point-supervision for action localization in videos. We start from spatio-temporal proposals, normally used during inference to determine the action location. We propose to bypass the need for box-supervision by learning directly from spatio-temporal proposals in training videos, guided by point-supervision. Experimental evaluation on three action localization datasets shows that our approach yields similar results to box-supervision. Moreover, our approach can handle sparse and noisy point annotations, resulting in a 20 to 150 times speed-up for action supervision. To help guide the selection of spatio-temporal proposals during inference, we propose pseudo-points, automatic visual cues in videos that hallucinate points in test videos. When weighted and selected properly with our quality measure, pseudo-points can have a positive impact on the action localization performance. We conclude that points provide a fast and viable alternative to boxes for spatio-temporal action localization.

Table 2 Comparative evaluation of point-supervised action localization to the state-of-the-art using box-supervision as well as weakly-supervised alternatives

	Proposals	Supervision	UCF sports(AUC)	UCF-101(mAP)	Hollywood2Tubes(mAP)
Lan et al. (2011)	×	Box	0.380	–	–
Tian et al. (2013b)	×	Box	0.420	–	–
Wang et al. (2014)	×	Box	0.470	–	–
Jain et al. (2014)	✓	Box	0.520	–	–
van Gemert et al. (2015)	✓	Box	0.546	0.345	–
Soomro et al. (2015)	✓	Box	0.550	–	–
Gkioxari and Malik (2015)	×	Box	0.559	–	–
Weinzaepfel et al. (2015)	×	Box	0.559	0.468	–
Jain et al. (2017)	✓	Box	0.570	0.475	–
Hou et al. (2017)	×	Box	0.580	0.471	–
Saha et al. (2017)	×	Box	–	0.631	–
Yang et al. (2017)	×	Box	–	0.735	–
Kalogeiton et al. (2017a)	×	Box	–	0.772	–
Jain et al. (2015)	✓	Zero-shot	0.232	–	–
Mettes and Snoek (2017)	×	Zero-shot	0.393	–	–
Soomro and Shah (2017)	×	Unsupervised	0.450	–	–
Sharma et al. (2015) from (Li et al. 2018)	×	Video-label	–	0.055	–
Cinbis et al. (2017) from (Mettes et al. 2016)	✓	Video-label	0.278	0.136	0.009
Chen and Corso (2015)	✓	Video-label	0.530	–	–
Li et al. (2018)	×	Video-label	–	0.369	–
Mettes et al. (2017)	✓	Video-label	0.556	0.374	0.172
Mettes et al. (2016)	✓	Point	0.545	0.348	0.143
This paper	✓	Point	0.598	0.418	0.178

All results are shown for an overlap threshold of 0.2. On all datasets our approach compares favorably to all weakly-supervised localization approaches, indicating the effectiveness of point-supervision. On UCF Sports, we perform comparable or better to approaches that require box-supervision. On UCF-101, we outperform the approach based on box-supervision with the same proposals and features (van Gemert et al. 2015), but we are outperformed by approaches that score and link individual boxes (Kalogeiton et al. 2017a; Saha et al. 2017; Yang et al. 2017). We expect that higher quality spatio-temporal proposals, can narrow this gap (see Fig. 9). On Hollywood2Tubes, which only provides point annotations for training, we set a new state-of-the-art

Bold values are used to correctly identify which methods perform best per dataset

Acknowledgements This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number D17PC00343. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix A: Types of Localization Errors

For the error diagnosis, we consider five types of detections, parameterized by an overlap threshold τ . The first detection type is a correct detection (detection from positive video with an overlap of at least τ). The second type is a localization error (detection from positive video with an overlap less than τ , but greater than 0.1). The third type is confusion with another action (detection from negative video with an overlap of at least 0.1). The fourth type is background detection from own action (detection from positive video with an overlap less than 0.1). The fifth and final type is background detection from another action (detection from negative video with an overlap less than 0.1). These five types cover all possible types of detections. These types are similar to (Hoiem et al. 2012). Different from (Hoiem et al. 2012), we do not split actions into similar and dissimilar (since no such subdivision

exists). Instead, we split background detections into detection from own and other actions.

References

- Andrews, S., Tsochantaridis, & Hofmann, T. (2002). Support vector machines for multiple-instance learning. In *Advances in neural information processing systems*.
- Bearman, A., Russakovsky, O., Ferrari, V., & Fei-Fei, L. (2016). Whats the point: Semantic segmentation with point supervision. In *European conference on computer vision*.
- Chen, W., & Corso, J. J. (2015). Action detection by implicit intentional motion clustering. In *International conference on computer vision*.
- Cinbis, R. G., Verbeek, J., & Schmid, C. (2017). Weakly supervised object localization with multi-fold multiple instance learning. *IEEE transactions on pattern analysis and machine intelligence*, 39(1), 189–203.
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine learning*, 20(3), 273–297.
- Doersch, C., Gupta, A., & Efros, A. A. (2015). Unsupervised visual representation learning by context prediction. In *International conference on computer vision*.
- Dollár, P., Rabaud, V., Cottrell, G., & Belongie, S. (2005). Behavior recognition via sparse spatio-temporal features. In *Visual surveillance and performance evaluation of tracking and surveillance workshops, IEEE*, (pp. 65–72).
- Donahue, J., Anne Hendricks, L., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., & Darrell, T. (2015). Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Feichtenhofer, C., Pinz, A., & Zisserman, A. (2016). Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Fernando, B., Bilen, H., Gavves, E., & Gould, S. (2017). Self-supervised video representation learning with odd-one-out networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Gkioxari, G., & Malik, J. (2015). Finding action tubes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Hoiem, D., Chodpathumwan, Y., & Dai, Q. (2012). Diagnosing error in object detectors. In *European conference on computer vision*.
- Hou, R., Chen, C., & Shah, M. (2017). Tube convolutional neural network (t-cnn) for action detection in videos. In *International conference on computer vision*.
- Jain, M., Jegou, H., & Bouthemy, P. (2013). Better exploiting motion for better action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Jain, M., Van Gemert, J., Jégou, H., Bouthemy, P., & Snoek, C. G. M. (2014). Action localization with tubelets from motion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Jain, M., van Gemert, J. C., Mensink, T., & Snoek, C. G. M. (2015). Objects2action: Classifying and localizing actions without any video example. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Jain, M., van Gemert, J., Jégou, H., Bouthemy, P., & Snoek, C. G. M. (2017). Tubelets: Unsupervised action proposals from spatiotemporal super-voxels. *International Journal of Computer Vision*, 124(3), 287–311.
- Jain, S. D., & Grauman, K. (2016). Click carving: Segmenting objects in video with point clicks. arXiv preprint [arXiv:1607.01115](https://arxiv.org/abs/1607.01115).
- Ji, S., Xu, W., Yang, M., & Yu, K. (2013). 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1), 221–231.
- Kalogeiton, V., Weinzaepfel, P., Ferrari, V., & Schmid, C. (2017a). Action tubelet detector for spatio-temporal action localization. In *International conference on computer vision*.
- Kalogeiton, V., Weinzaepfel, P., Ferrari, V., & Schmid, C. (2017b). Joint learning of object and action detectors. In *International conference on computer vision*.
- Lan, T., Wang, Y., & Mori, G. (2011). Discriminative figure-centric models for joint action localization and recognition. In *International conference on computer vision*.
- Laptev, I. (2005). On space-time interest points. *International Journal of Computer Vision*, 64(2–3), 107–123.
- Li, Z., Gavriluk, K., Gavves, E., Jain, M., & Snoek, C. G. M. (2018). Videolstm convolves, attends and flows for action recognition. *Computer Vision and Image Understanding*, 166, 41–50.
- Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *European conference on computer vision*.
- Manen, S., Gygli, M., Dai, D., & Van Gool, L. (2017). Pathtrack: Fast trajectory annotation with path supervision. In *International conference on computer vision*.
- Marian Puscas, M., Sangineto, E., Culibrk, D., & Sebe, N. (2015). Unsupervised tube extraction using transductive learning and dense trajectories. In *International conference on computer vision*.
- Marszałek, M., Laptev, I., & Schmid, C. (2009). Actions in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Mettes, P., & Snoek, C. G. M. (2017). Spatial-aware object embeddings for zero-shot localization and classification of actions. In *International conference on computer vision*.
- Mettes, P., van Gemert, J. C., & Snoek, C. G. M. (2016). Spot on: Action localization from pointly-supervised proposals. In *European conference on computer vision*.
- Mettes, P., Snoek, C. G. M., & Chang, S. F. (2017). Localizing actions from video labels and pseudo-annotations. In *British machine vision conference*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Workshops at the twenty-sixth AAAI conference on artificial intelligence*.
- Oneata, D., Revaud, J., Verbeek, J., & Schmid, C. (2014). Spatio-temporal object detection proposals. In *European conference on computer vision*.
- Papadopoulos, D. P., Uijlings, J. R., Keller, F., & Ferrari, V. (2016). We don't need no bounding-boxes: Training object class detectors using only human verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Papadopoulos, D. P., Uijlings, J. R., Keller, F., & Ferrari, V. (2017). Extreme clicking for efficient object annotation. In *International conference on computer vision*.
- Peng, X., & Schmid, C. (2016). Multi-region two-stream r-cnn for action detection. In *European conference on computer vision*.
- Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*.
- Rodriguez, M. D., Ahmed, J., & Shah, M. (2008). Action MACH: a spatio-temporal maximum average correlation height filter for action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Russakovsky, O., Li, L. J., & Fei-Fei, L. (2015). Best of both worlds: human-machine collaboration for object annotation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.

- Saha, S., Singh, G., Sapienza, M., Torr, P., & Cuzzolin, F. (2016). Deep learning for detecting multiple space-time action tubes in videos. In British machine vision conference.
- Saha, S., Singh, G., & Cuzzolin, F. (2017). Amtnet: Action-micro-tube regression by end-to-end trainable deep architecture. In International conference on computer vision.
- Sánchez, J., Perronnin, F., Mensink, T., & Verbeek, J. (2013). Image classification with the fisher vector: Theory and practice. *International Journal of Computer Vision*, 105(3), 222–245.
- Sharma, S., Kiros, R., & Salakhutdinov, R. (2015). Action recognition using visual attention. In Advances in neural information processing systems workshop.
- Simonyan, K., & Zisserman, A. (2014). Two-stream convolutional networks for action recognition in videos. In Advances in neural information processing systems.
- Singh, G., Saha, S., Sapienza, M., Torr, P., & Cuzzolin, F. (2017). Online real-time multiple spatiotemporal action localisation and prediction. In International conference on computer vision.
- Siva, P., Xiang, T. (2011). Weakly supervised action detection. In British machine vision conference.
- Soomro, K., & Shah, M. (2017). Unsupervised action discovery and localization in videos. In International conference on computer vision.
- Soomro, K., Zamir, A. R., & Shah, M. (2012). Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint [arXiv:1212.0402](https://arxiv.org/abs/1212.0402).
- Soomro, K., Idrees, H., & Shah, M. (2015). Action localization in videos through context walk. In International conference on computer vision.
- Srivastava, N., Mansimov, E., & Salakhudinov, R. (2015). Unsupervised learning of video representations using lstms. In International conference on machine learning.
- Su, H., Deng, J., & Fei-Fei, L. (2012). Crowdsourcing annotations for visual object detection. In AAAI workshops.
- Tian, Y., Sukthankar, R., & Shah, M. (2013a). Spatiotemporal deformable part models for action detection. In Proceedings of the IEEE conference on computer vision and pattern recognition.
- Tian, Y., Sukthankar, R., & Shah, M. (2013b). Spatiotemporal deformable part models for action detection. In Proceedings of the IEEE conference on computer vision and pattern recognition.
- Tran, D., Bourdev, L., Fergus, R., Torresani, L., & Paluri, M. (2015). Learning spatiotemporal features with 3d convolutional networks. In International conference on computer vision.
- Tseng, P. H., Carmi, R., Cameron, I. G., Munoz, D. P., & Itti, L. (2009). Quantifying center bias of observers in free viewing of dynamic natural scenes. *Journal of Vision*, 9(7), 4.
- Uijlings, J. R. R., Van De Sande, K. E. A., Gevers, T., & Smeulders, A. W. M. (2013). Selective search for object recognition. *International Journal of Computer Vision*, 104(2), 154–171.
- van Gemert, J. C., Jain, M., Gati, E., & Snoek, C. G. M. (2015). Apt: Action localization proposals from dense trajectories. In British machine vision conference.
- Vondrick, C., Patterson, D., & Ramanan, D. (2013). Efficiently scaling up crowdsourced video annotation. *International Journal of Computer Vision*, 101(1), 184–204.
- Wang, H., & Schmid, C. (2013). Action recognition with improved trajectories. In International conference on computer vision.
- Wang, H., Ullah, M. M., Klaser, A., Lapedv, I., & Schmid, C. (2009). Evaluation of local spatio-temporal features for action recognition. In British machine vision conference.
- Wang, H., Kläser, A., Schmid, C., & Liu, C. L. (2013). Dense trajectories and motion boundary descriptors for action recognition. *International Journal of Computer Vision*, 103(1), 60–79.
- Wang, L., Qiao, Y., & Tang, X. (2014). Video action detection with relational dynamic-poselets. In European conference on computer vision.
- Weinzaepfel, P., Harchaoui, Z., & Schmid, C. (2015). Learning to track for spatio-temporal action localization. In International conference on computer vision.
- Yang, Z., Gao, J., & Nevatia, R. (2017). Spatio-temporal action detection with cascade proposal and location anticipation. In British machine vision conference.
- Yu, G., & Yuan, J. (2015). Fast action proposals for human action detection and search. In Proceedings of the IEEE conference on computer vision and pattern recognition.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.