



HAL
open science

Towards Weakly-Supervised Action Localization

Philippe Weinzaepfel, Xavier Martin, Cordelia Schmid

► **To cite this version:**

Philippe Weinzaepfel, Xavier Martin, Cordelia Schmid. Towards Weakly-Supervised Action Localization. 2016. hal-01317558v1

HAL Id: hal-01317558

<https://inria.hal.science/hal-01317558v1>

Preprint submitted on 18 May 2016 (v1), last revised 24 May 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards Weakly-Supervised Action Localization

Philippe Weinzaepfel, Xavier Martin, and Cordelia Schmid

Inria

Abstract. This paper presents a novel approach for weakly-supervised action localization, *i.e.*, that does not require per-frame spatial annotations for training. We first introduce an effective method for extracting human tubes by combining a state-of-the-art human detector with a tracking-by-detection approach. Our tube extraction leverages the large amount of annotated humans available today and outperforms the state of the art by an order of magnitude: with less than 5 tubes per video, we obtain a recall of 95% on the UCF-Sports and J-HMDB datasets. Given these human tubes, we perform weakly-supervised selection based on multi-fold Multiple Instance Learning (MIL) with improved dense trajectories and achieve excellent results. We obtain a mAP of 84% on UCF-Sports, 54% on J-HMDB and 45% on UCF-101, which outperforms the state of the art for weakly-supervised action localization and is close to the performance of the best fully-supervised approaches.

The second contribution of this paper is a new realistic dataset for action localization, named **DALY** (Daily Action Localization in YouTube). It contains high quality temporal and spatial annotations for 10 actions in 31 hours of videos (3.3M frames), which is an order of magnitude larger than standard action localization datasets. On the DALY dataset, our tubes have a spatial recall of 82%, but the detection task is extremely challenging, we obtain 10.8% mAP.

1 Introduction

Action classification has been widely studied over the past decade and state-of-the-art methods [1–4] now achieve excellent performance. However, to analyze video content in more detail, we need to localize actions in space and time. A potential application is action detection in the context of self-driving cars, where it is necessary to know where the people are and in which direction they are moving. Other examples are the analysis of sport videos, where we would like to annotate the motion of individual players, and the auto-annotation of movies, where action localization is necessary to describe the content.

Localizing actions in videos is a challenging task which has received increasing attention over the past few years. Recently, significant progress was achieved, see for example [5–7]. Nevertheless, these methods require a large amount of supervision. For instance, per-frame bounding box annotations are used for training class-specific detectors [5, 6]. Wang *et al.* [7] additionally require pose annotations, as they represent actions as a sequence of skeleton models. Several works

have suggested to generate action proposals before classifying them [8, 9]. However, supervision is still required for learning to classify these hundreds or thousands of proposals. Consequently, all these approaches are restricted to relatively small datasets because of the supervision cost and can not be generalized easily to more classes.

In this paper, we propose a weakly-supervised action localization method, *i.e.*, that does not require any spatial annotation for training. The first step of our approach consists in extracting human tubes from videos. Using human tubes for action recognition is not a novel idea [10–12]. However, we show for the first time that extracting highly robust human tubes is possible by leveraging a recent state-of-the-art object detection approach [13], a large annotated dataset of humans in a variety of poses [14] and state-of-the-art tracking-by-detection [15, 16]. We show that a small number of human tubes per video is sufficient to obtain more than 95% recall on challenging action localization datasets. Our approach outperforms existing video proposal methods by an order of magnitude, as shown in Section 4.3. As a second step, we describe human tubes with improved dense trajectories [1] and use multi-fold Multiple Instance Learning (MIL) [17] to select human tubes containing the action. This requires only labels at the video level and manages to select proposals corresponding to the action category accurately. During testing, we extract human tubes and we obtain a mAP using classifiers learned from the tubes selected in the training videos. Figure 1 illustrates our weakly-supervised action localization method. Our approach significantly outperforms the state of the art for weakly-supervised action localization and is only slightly worse than the best fully-supervised methods, with a mAP of 84% on UCF-Sports, 54% on J-HMDB and 45% on UCF-101.

We also introduce a new action localization dataset, named **DALY** for Daily Action Localization in YouTube. It overcomes the limitations of existing datasets which mostly consist of videos trimmed to the action, have limited action types, *i.e.*, often sports only, and contain in many cases only one human per video. The DALY dataset consists of more than 31 hours of videos from YouTube with high quality spatial and temporal annotations for 10 realistic daily actions. The task is to localize relatively short actions in long untrimmed videos. Furthermore, it includes videos with multiple humans performing actions at the same time. On the DALY dataset, our tubes have a spatial recall of 82%, but the detection task is extremely challenging, and we obtain 10.8% mAP.

This paper is organized as follows. We first review related work in Section 2. Next, Section 3 presents the datasets used for evaluation and introduces our DALY dataset. We then describe our approach for extracting human tubes in a video (Section 4) and our weakly-supervised learning method using multi-fold MIL (Section 5). Finally, experimental results are presented in Section 6.

2 Related work

Initial attempts for temporal and spatio-temporal action localization are based on a sliding-window scheme and focus on improving the search complexity [11,

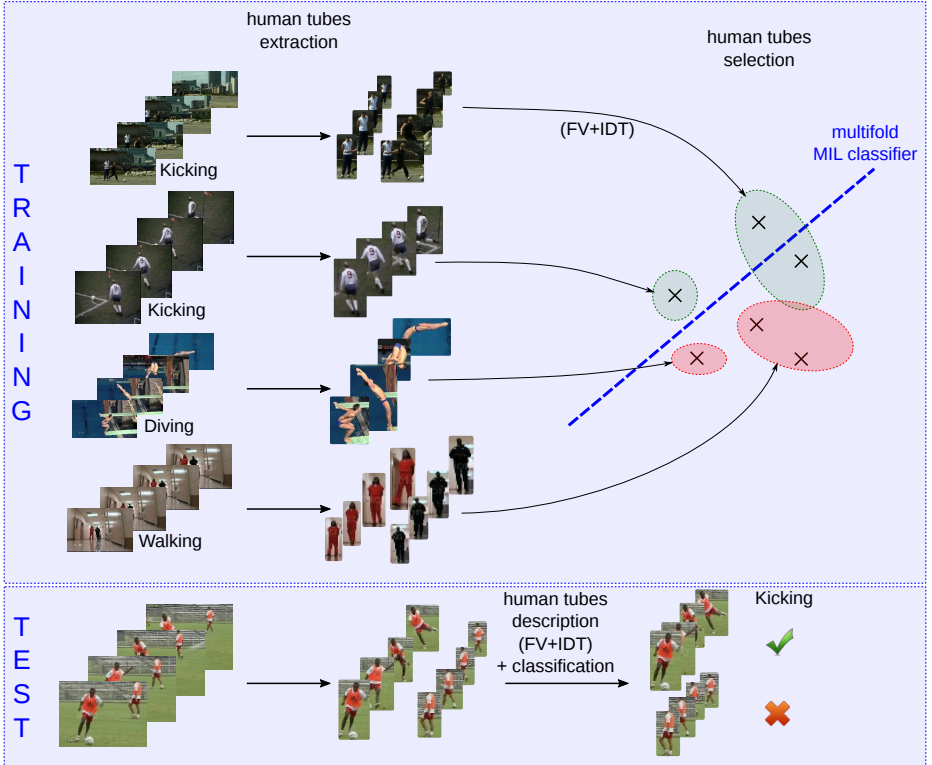


Fig. 1. Illustration of our weakly-supervised action localization approach. For training, we extract human tubes in each video and learn a classifier with weak supervision using Multiple Instance Learning for selecting tubes containing the action. At test time, human tubes are extracted and scored using the learned classifiers.

18–20]. Other approaches rely on figure-centric models. For instance, Lan *et al.* [21] consider the human position as a latent variable and infer it jointly with the action label. Kläser *et al.* [10] use a human detector and build human tubes using KLT features tracks. The human tracks are then classified with HOG-3D descriptors [22]. Our approach is also based on human tubes but is significantly more robust to huge variations in pose and appearance thanks to a human-specific tracking-by-detection approach [15, 16] as well as recent advances in detectors [13] and datasets [14]. In addition, our approach is weakly-supervised, *i.e.*, does not require bounding box annotation for labeling the training samples.

Several recent methods for action localization are based on action tubes to reduce the search complexity. Jain *et al.* [23] construct action tubes by hierarchically merging supervoxels and rely on dense trajectory features for tube classification. Similarly, van Gemert *et al.* [9] cluster trajectories and use the resulting tubes for action detection. In [12], proposals are based on an action-ness measure [24], which requires localized training samples. In parallel, several works [8, 25] have attempted to further improve the quality of tubes. While all these methods generate thousands of proposals and require ground-truth to annotate the proposals used for training, we obtain only few human tube proposals, thus allowing us to apply Multiple Instance Learning (MIL) effectively.

Recently, CNNs for human action localization have emerged [5, 6]. These approaches rely on appearance and motion CNNs for classifying region proposals in individual frames. Tracks are obtained by combining class-specific detections with either temporal linking based on proximity [5] or by class-specific tracking-by-detection [6]. Our approach also relies on tracking-by-detection, but performs generic human detections and tracking. Thus the complexity is divided by the number of classes and the tracks can be used for weakly-supervised learning. Furthermore, our tracker is significantly faster, as it uses the region-pooling layer of faster R-CNN [13] for feature computation.

These later approaches require bounding box annotations in every frame from training videos. Such annotation effort is not realistic for large-scale datasets. Weakly-supervised action localization is thus necessary, but has received little attention so far. The case of temporal localization was studied in [26–28]. Bojanowski *et al.* [26] assume an ordered list of the actions in each video as input. Duchenne *et al.* [27] use a discriminative clustering on short video segments to identify the localization in the training set and learn a classifier. Hoai *et al.* [28] extend a Multiple Instance SVM to time series while allowing discontinuities in the positive samples. For the case of spatio-temporal localization, Siva and Xiang [29] define cuboids of different time lengths around detected humans, describe them with STIPs [30] and then use Multiple Instance Learning (MIL). Their method can thus only be used for static humans whereas ours can be generalized to more complex and realistic motions. Mosabbeh *et al.* [31] use a subspace segmentation clustering approach applied on groups of trajectories in order to segment videos into parts. Low-rank matrix completion then estimates the contribution of each cluster to the different labels. Hence, the approach detects several disjoint action parts and not one spatio-temporal consistent local-

ization. Ma *et al.* [32] first extract a per-frame hierarchical segmentation, which is tracked over the videos. Using a foreground scoring, they obtain a hierarchy of spatio-temporal segments where the upper level corresponds to human body location candidates. More recently, Chen and Corso [33] propose to generate unsupervised proposals by clustering intentional motion based on dense trajectories. A classifier is then learned using the best proposal of a video as positive sample. The method is thus weakly-supervised but assumes only one action per video during training. Furthermore, this method is not robust to nearby motions and assumes significant motion.

Some action classification methods automatically discover discriminative parts to improve the performance, but do not aim at precisely localizing the action. For instance, Shapovalova *et al.* [34] extend latent SVM to model pairwise similarities between latent variables, aiming at discovering common parts for a particular action. Boyraz *et al.* [35] optimize jointly the classification error and the location of a fixed number of discriminative parts. Lan *et al.* [36] leverage a discriminative clustering approach for parsing complex actions into mid-level action elements.

3 Dataset and evaluation

Experiments are conducted on three existing action localization datasets described in Section 3.1: UCF-Sports, J-HMDB, and UCF-101. In Section 3.2, we introduce **DALY** (Daily Action Localization in YouTube), a new challenging action localization dataset.

3.1 Existing action localization datasets

- The **UCF-Sports** dataset [37] consists of 150 sports videos with 10 actions, such as *diving* or *running*. Videos are trimmed to the action and every frame is annotated with a bounding box. For each class, the sequences present similarities in background, camera viewpoint and actors, which is a limitation of this dataset. We use the train/test split defined in [21].
- The **J-HMDB** dataset [38] contains 928 short videos with 21 actions, including *stand up*, *run* and *pour*. The videos are trimmed to the action, are very short (1.4 sec on average) and contain only one human. The annotations are human silhouettes in every frame. We use the bounding boxes around these silhouettes as ground-truth. The dataset has 3 train/test splits.
- The **UCF-101** dataset [39] contains spatio-temporal annotation for 24 actions in 3207 sports videos. In contrast to UCF-Sports and J-HMDB, the detection is also temporal but the videos remain short; for half of the classes, the action lasts for more than 80% of the video duration. There are 3 train/test splits. Results are reported for the first split only.

3.2 The DALY dataset

The existing datasets for action localization are limited by the type of actions (mainly sports), the lack of diversity in the videos and the duration of the

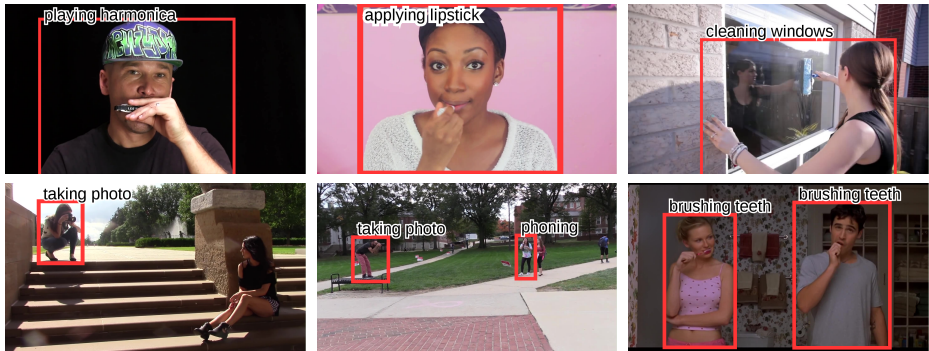


Fig. 2. Example frames from the DALY dataset.

videos, see Table 1 for details. There is a clear need for a realistic dataset for spatio-temporal action localization. We introduce DALY, a dataset for **D**aily **A**ction **L**ocalization in **Y**ouTube¹. The DALY dataset consists of 31 hours of YouTube videos, with spatial and temporal annotations for 10 everyday human actions: *applying make up on lips*, *brushing teeth*, *cleaning floor*, *cleaning windows*, *drinking*, *folding textile*, *ironing*, *phoning*, *playing harmonica* and *taking photos/videos*, see Figure 2. The videos are collected from YouTube using related queries. We collect 510 videos, *i.e.*, 51 videos per category, representing a total of around 3.3M frames. Each video lasts between 1 and 20 minutes with an average duration of 3min 45s. We generate a split with 31 training videos and 20 test videos for each class, ensuring that videos with the same characters or scenes are in the same set.

Temporal annotations of the 10 actions result in 3724 instances in total. Actions are short (8 seconds on average) with some classes having very brief instances (*e.g. drinking*) or somewhat longer (*e.g. brushing teeth*). The actions cover about 25% of the videos. For each instance in the test set, we annotate the spatial extent, *i.e.* a bounding box around the actor, for 5 regularly sampled frames. Note that videos can contain simultaneously multiple actions, see bottom row, middle and right column in Figure 2. More details on the DALY dataset are given in Appendix A.

4 Building human tubes

Given a video V with F frames, spatio-temporal action localization aims at detecting when and where actions are performed. More precisely, for each instance of an action, a tube T , *i.e.*, a set of bounding boxes $T = \{b_f\}_{f=f_s..f_e}$ around the actor is returned, with one box b_f per frame f during the predicted temporal extent of the action between frames f_s and f_e .

¹ The DALY dataset is available online at <http://pascal.inrialpes.fr/data/daly/>.

	DALY	UCF-Sports [37]	J-HMDB [38]	UCF-101 [39]	MSR-II [18]
#classes	10	10	21	24	3
action types	everyday	sports	everyday	sports	artificial
#videos	510	150	928	3207	54
avg resolution	1290x790	690x450	320x240	320x240	320x240
total #frames	3.3M	10k	32k	558k	41k
avg video dur.	3min 45s	5.8s	1.4s	5.8s	51s
avg action dur.	8s	5.8s	1.4s	4.5s	6s
#instances	3724	154	928	4030	203
spatial annotation	subset	all	all	all	cuboid

Table 1. Comparison of our DALY dataset with existing action localization datasets.

This section presents the human detector (Section 4.1) and the human-specific tracker (Section 4.2) used to obtain reliable human tubes. We denote by \mathcal{T}_V the set of human tubes proposals extracted for a video V . Section 4.3 presents an evaluation of the obtained human tubes and a comparison with state-of-the-art action proposals.

4.1 Human detector

Faster R-CNN detector. We use the state-of-the-art detector Faster R-CNN [13] to train our human detector. Faster R-CNN integrates a Region Proposals Network (RPN) to produce 300 proposals per image, that are classified as background or a particular class (here ‘human’). Detection is extremely fast (around 200ms per image on a GPU) since convolutions are only computed once and then used both for generating proposals and for scoring each proposal using a region pooling layer. Faster R-CNN also includes bounding box regression to overcome the stride of the network which limits the precision of the estimated localization. We use the variant with end-to-end training² using VGG net with 16 layers [40]. Given a frame f , let \mathcal{B}_f be the set of bounding box detections output by the network and $s_H(b)$ be the human score after softmax of a box $b \in \mathcal{B}_f$. We now present the data we leverage for training the network.

External training data. We use the MPII Human Pose dataset [14] to leverage training data with sufficient variability. It contains more than 40k annotated poses, including a bounding box around the head and joint positions for the full body. The images represent several frames from around 4000 videos, selected to contain around 500 different activities. We use the training set for which annotations are publicly available. It represents 28778 annotated poses in 17372 images. We obtain a bounding box for each person by taking the box containing the head and all visible joints, with a fixed additional margin of 20 pixels. The bounding boxes are thus not perfect, see Figure 3. For instance, they can be slightly too large (top of bounding boxes from left image) or may not cover the extremity of the limbs (hands and feet in the second image). In particular, the bounding boxes are also cropped if some joints are not visible (right image).

² <https://github.com/rbgirshick/py-faster-rcnn>

Nevertheless, this dataset remains large enough and offers a huge variability in the poses. It is thus well suited for training an accurate human detector.



Fig. 3. Examples of bounding boxes (green) estimated from annotated joints and head (yellow) on the MPII Human Pose dataset.



Fig. 4. Example results of our human detector.

Qualitative results. Example detections for different action localization datasets (UCF-Sports and J-HMDB) are shown in Figure 4. One can see that the obtained human detector is robust to unusual poses (first two examples) and to humans that are not fully-visible (third example). Nevertheless, detection are sometimes imprecise, for example not completely covering the human (second example). Note that there can be multiple detections when different people are present (fourth example).

4.2 Human-specific tracker

Once humans are detected, the second step consists in tracking these humans to build tube proposals. To this end, we design a human-specific tracking-by-detection approach. In the following, we present the different steps of our approach. It is summarized in Algorithm 1.

Boxes to track. The list of boxes to track is initialized with all detected candidates \mathcal{B}_f in all frames f . We start by building a first track $T = \{b_f\}_{f=1..F}$ using the detection with the highest human score in the entire video sequence. Once we have tracked this detection over the whole sequence, we remove all detections that have an Intersection Over Union (IoU) above 0.3 with any box b_f in this track. We then run the tracker a second time starting from the remaining

detection with the highest human score and repeat the process until no boxes are left.

Initial box refinement. Let \mathbf{b} be the selected box from frame f . As a first step, we refine its position by performing a search for a higher scoring location in its neighborhood. To this end, we perform a forward pass of the network on the frame f using the neighborhood $N(\mathbf{b})$ of \mathbf{b} , *i.e.*, we use the network without the region proposal part and the bounding box regression branch. The neighborhood $N(\mathbf{b})$ of \mathbf{b} simulates a sliding window in scale and space. More precisely, we use the box \mathbf{b} plus its translation with $0, \pm\mathfrak{s}, \pm2\mathfrak{s}, \pm3\mathfrak{s}, \pm4\mathfrak{s}$ where \mathfrak{s} is the stride of the network, and similarly for rescaled versions of the box by a factor $0, \pm10\%, \pm20\%$. The refined position b_f of the initialization box is set to the region of interest $b \in N(\mathbf{b})$ that maximizes the human score $s_H(b)$.

Instance-level detector initialization. Based on this box b_f , we learn an instance-level detector using the features from the last fully connected layer denoted by $fc7$. More precisely, we learn a linear SVM using as positive \mathcal{P} the feature from the refined box $\mathcal{P} = \{b_f\}$ and as negatives \mathcal{N} the features from boxes in \mathcal{B}_f that have (almost) no overlap with b_f : $\mathcal{N} = \{b \mid b \in \mathcal{B}_f \text{ s.t. } \text{IoU}(b, b_f) < 0.1\}$.

Tracking procedure. Starting from the refined box b_f , we first track it forward until the end of the sequence. The box b_i of the track at a frame i is set using a sliding window to optimize the human score and the instance-level score. More precisely, at each frame i , we perform a forward pass of the network using the boxes in the neighborhood $N(b_{i-1})$ of the tracked box from the previous frame. b_i is then set to the box b that maximizes the sum of the human score $s_H(b)$ and the instance-level detector probability $s_I(b)$, computed using a sigmoid on the SVM score, among all boxes $b \in N(b_{i-1})$. We continue until we reach the end of the sequence updating the instance-level detector on the fly, see next paragraph. At the end of the forward pass, we reset the instance-level detector and track the initial refined box backward until the beginning of the sequence.

Instance-level detector update. At each frame, we also update the instance-level detector by adding b_i to the set of positive features \mathcal{P} and boxes from \mathcal{B}_i with (almost) no overlap with b_i to the negatives \mathcal{N} . At each iteration, we furthermore restrict \mathcal{N} to the hard negatives.

Efficient computation. Launching the tracker multiple times on a sequence can be performed efficiently. To extract human tubes in a video V , we first run Faster R-CNN in every frame f to detect humans. We also keep in cache the results of the last convolution layer denoted by $conv5_f$, which will be used for the sliding window, and the last fully connected layer $fc7_f$, which will be used as feature descriptors for the negatives of the instance-level detector. Now, when performing a sliding window at frame f using a forward pass of the network, we do not need to re-compute the convolutions, and we can directly start from $conv5_f$. Similarly, the $fc7$ feature descriptors of all detections are cached, so updating the set of negatives does not require any additional computation.

Algorithm 1 Extracting human tubes from a video.

Input: a video sequence V with F frames.

Output: a list of human tubes $\mathcal{T}_V = \{T\}$ with $T = \{b_f\}_{f=1..F}$.

 $\mathcal{T}_V \leftarrow \emptyset$

 ToTrack $\leftarrow \emptyset$
For $f = 1 \dots F$

 — $\mathcal{B}_f \leftarrow \text{FasterRCNN}(f)$

 (*human detection; cache conv5_f and fc7_f*)

 — ToTrack $\leftarrow \text{ToTrack} \cup \{(b, f) \mid b \in \mathcal{B}_f\}$
While ToTrack $\neq \emptyset$

 — $b, f \leftarrow \text{argmax}_{(b,i) \in \text{ToTrack}} s_H(b)$

 (*select highest scoring box to track*)

 — $b_f \leftarrow \text{argmax}_{b \in N(b)} s_H(b)$

 (*initial refinement; forward from conv5_f*)

 — $\mathcal{P} \leftarrow \{b_f\}$

 — $\mathcal{N} \leftarrow \{b \mid b \in \mathcal{B}_f \text{ s.t. } \text{IoU}(b, b_f) < 0.1\}$

 (*using cached fc7_f*)

 — **For** $i = f + 1 \dots F$ **and** $i = f - 1 \dots 1$

 ——— Learn a linear SVM with positives \mathcal{P} and negatives \mathcal{N}

 ——— $b_i \leftarrow \text{argmax}_{b \in N(b_{i \pm 1})} s_H(b) + s_I(b)$

 (*sliding window; forward from conv5_i*)

 ——— $\mathcal{P} \leftarrow \mathcal{P} \cup \{b_i\}$

 ——— $\mathcal{N} \leftarrow \mathcal{N} \cup \{b \mid b \in \mathcal{B}_i \text{ s.t. } \text{IoU}(b, b_i) < 0.1\}$

 (*using cached fc7_i*)

 ——— $\mathcal{N} \leftarrow \{b \mid b \in \mathcal{N} \text{ s.t. } s_I(b) \geq -1\}$

 (*restrict to hard negatives*)

 — $\mathcal{T}_V \leftarrow \mathcal{T}_V \cup \{T = \{b_i\}_{i=1..F}\}$

 (*add the track to the tube list*)

 — ToTrack $\leftarrow \{(b, i) \mid (b, i) \in \text{ToTrack} \text{ s.t. } \text{IoU}(b_i, b) < 0.3\}$

 (*remove overlapping detections*)

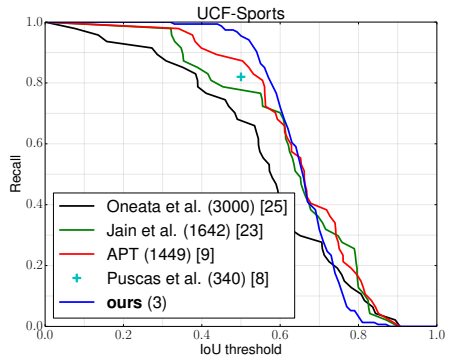
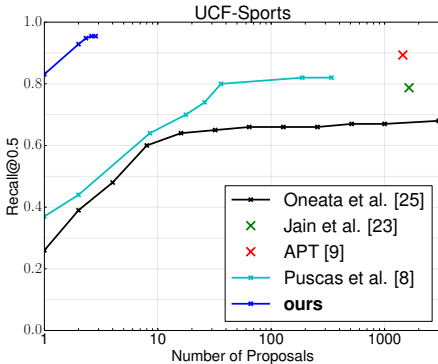


Fig. 5. Comparison of our human tubes to state-of-the-art proposals on the UCF-Sports dataset. Left: Recall@0.5 when varying the number of proposals. Right: Recall for varying IoU thresholds for a number of extracted proposals (indicated in parentheses).

4.3 Evaluation of our human tube proposals

We compare our human tubes to the state-of-the-art action proposals on the UCF-Sports dataset. Two different metrics are used. The first one is recall@0.5, *i.e.*, the ratio of ground-truth tubes covered by at least one proposal with average spatial IoU over 0.5, while varying the number of proposals. The second metric is the recall, computed over all proposals, for varying IoU thresholds. Results are presented in Figure 5. We can clearly see that our human tubes outperform the other approaches by a large margin. Given few proposals (3 on average), we obtain a recall of 95% at 0.5 IoU, whereas state-of-the-art approaches do not reach this recall with significantly more proposals (hundreds or thousands).

On J-HMDB, we also reach a recall of 95% for an IoU threshold at 0.5 with 3 proposals on average per video, demonstrating the effectiveness of our approach. For UCF-101, we obtain a recall of 80% at IoU 0.2 and of 48% at IoU 0.5. The significant decrease in performance is due to the low quality of the videos, sequences have low-resolution and are strongly compressed, and humans tend to be small. For our new DALY dataset, we obtain a spatial recall of 82% at an IoU threshold of 0.5. This is measured on the test set for which spatial annotations are available for a subset of frames. The excellent quality of our human tubes is, thus, confirmed on a realistic challenging dataset.

Figure 6 shows a few examples of the highest scoring human tube for several sequences of the DALY dataset. The first four examples show that the human tube extraction performs well despite motion of one arm (first row), turning of the person (second and third row), camera motion (third row) or presence of an animal close to the human (fourth row).

Nevertheless, there are some failure cases due to the fact that the full body disappears (fifth row). In this case, only the feet remain visible causing the failure of the human tracker, as the human detector performs poorly and the instance-level detector is trained on previous frames where the full body is visible. Other failure cases can be explained by a poor performance of the human detector, and thus of the human tube (last two rows). This typically happens when only a small part of the body is visible (sixth row) and in the case of occlusion (last row).

5 Weakly-supervised human tube classifier

In this section, we first introduce the feature descriptor of the human tubes (Section 5.1). We then present the weakly-supervised MIL learning (Section 5.2) and the detection stage (Section 5.3). We finally evaluate the MIL training procedure (Section 5.4).

5.1 Human tube features

For each human tube, we extract improved dense trajectories and aggregate them with a Fisher Vector representation [1]. We start by extracting the tra-

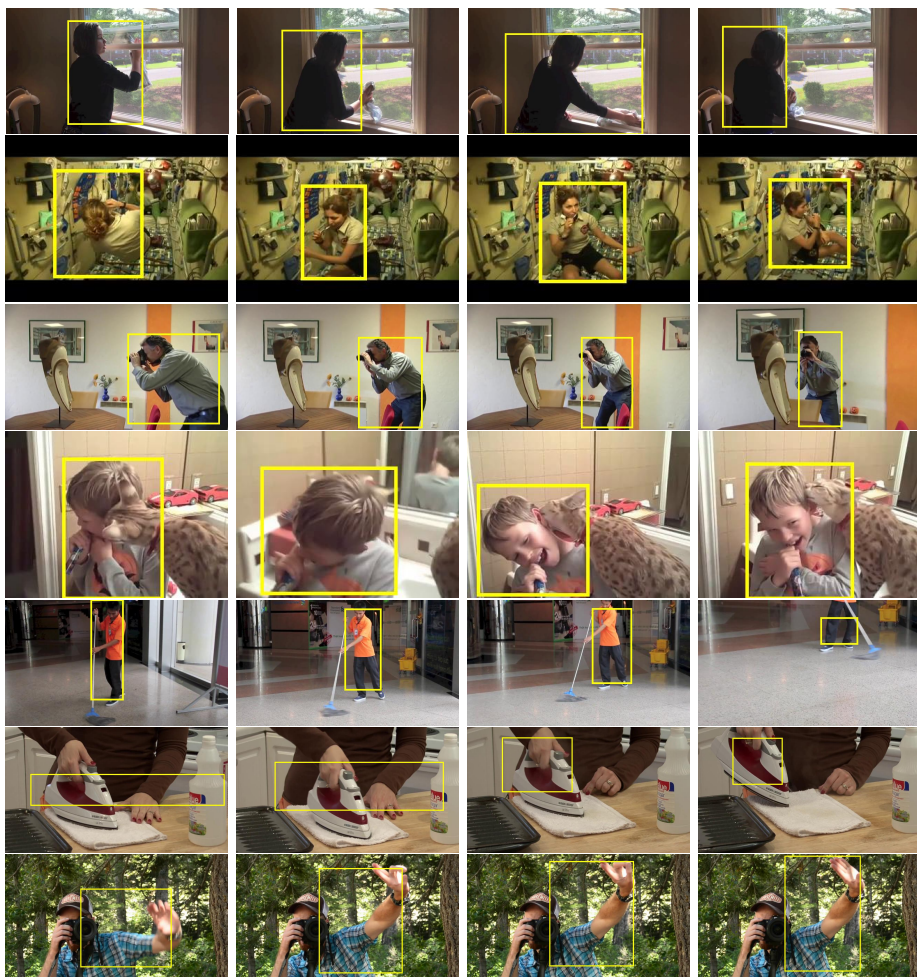


Fig. 6. Example of human tubes with successful human tube extraction in the first four rows, and some failure cases in the last three rows. Failures are caused by partial visibility of the human (end of fifth row and sixth row) and missed human detection caused by an occluding camera (last row).

jectories features for the entire video³. For each descriptor type (HOG, HOF, MBH_x, MBH_y), we first reduce its initial dimension by a factor of 2 using PCA and learn a codebook of 256 Gaussians. For each tube, we build a Fisher Vector per descriptor type, using only the trajectories that start inside the track (increased by a margin of 10%). Each of the 4 Fisher Vectors is then independently power-normalized and L2-normalized [41]. A tube is finally described by the concatenation of the 4 normalized Fisher Vectors, resulting in 102400 dimensions.

5.2 Multi-fold multiple instance learning

We use a Multiple Instance Learning (MIL) formulation to learn a detector in a weakly-supervised setting, *i.e.*, given only information about the presence/absence of a class in the training videos. For a given label, let \mathcal{V}_P (resp. \mathcal{V}_N) be the set of positive (resp. negative) videos. MIL alternates between inferring the localization of the action in the positive videos \mathcal{V}_P and using these locations to train a detector. One issue is that it tends to lock onto the initial locations [17], which is particularly the case for high dimensional descriptors. We therefore resort to the multi-fold variant proposed in [17]. In the following, we present the different steps of our approach. It is summarized in Algorithm 2. **Initialization.** For each video V , we use the track $T \in \mathcal{T}_V$ with the highest average human score $s_H(T) = \sum_{b_f \in T} s_H(b_f)/F$ as initial tube. Videos are positive and negative according to the annotated video labels.

Iteration. At each iteration, we learn a linear SVM using the current positive and negative tubes, denoted by \mathcal{T}_P and \mathcal{T}_N respectively. We then perform two hard negative mining iterations with negatives being mined only in negative videos \mathcal{V}_N . The next step consists in re-estimating the tube used as positive in each positive video. We randomly split the positive videos \mathcal{V}_P into $K = 4$ folds. For each fold, we learn a linear SVM using the positives from all other folds and all negatives. This classifier is then run on videos from this fold. The new estimated location in each positive video of this fold is set to the tube with the highest score and will be used as positive in the next iteration. We perform 10 iterations of multi-fold MIL.

5.3 Temporal supervision and detection

Given a test video, we first extract human tubes using the method presented in Section 4. We then extract a feature vector for each tube as described in Section 5.1. Next, we score each tube for all classes using the classifier learned in a weakly-supervised setting (Section 5.2).

For the UCF-101 dataset, which also requires temporal localization, we perform in addition a multi-scale temporal sliding window inside each tube. When training, since the videos are almost trimmed, we first perform multi-fold MIL iterations using features from the whole video. Next, we assume temporal supervision in the training set and train a new classifier using as positive features the

³ https://lear.inrialpes.fr/people/wang/improved_trajectories

Algorithm 2 Learning a classifier on human tubes with weak supervision.

Input: a set of positive (\mathcal{V}_P) and negative (\mathcal{V}_N) videos, each with a list of tubes.

Output: a linear SVM.

$\mathcal{T}_P \leftarrow \{\operatorname{argmax}_{T \in \mathcal{T}_V} S_H(T) \mid V \in \mathcal{V}_P\}$ *(initialization with the tube with*

$\mathcal{T}_N \leftarrow \{\operatorname{argmax}_{T \in \mathcal{T}_V} S_H(T) \mid V \in \mathcal{V}_N\}$ *highest human detection score)*

For each iteration

— Learn a linear SVM using all positives \mathcal{T}_P and all negatives \mathcal{T}_N

— Perform 2 hard negative iterations in \mathcal{V}_N *(update \mathcal{T}_N)*

— Randomly split \mathcal{V}_P in K folds

— **For** each fold k

— Learn a linear SVM using positives from other folds and all negatives

— Re-estimate the tube in the videos from fold k *(update \mathcal{T}_P)*

Learn a linear SVM using all positives \mathcal{T}_P and all negatives \mathcal{T}_N

descriptors from the track selected by multi-fold MIL, restricted to the ground-truth temporal extent. At test time, we perform a sliding window with the same temporal lengths as [6] ($\{20, 30, 40, 50, 60, 70, 80, 90, 100, 150, 300, 450, 600\}$ frames) and the same stride (10 frames). In order to penalize short tubes, we score each tube using the SVM score minus α/L where α is a parameter experimentally set to 20 and L is the length of the detection.

For the DALY dataset, we assume temporal supervision during training, *i.e.*, the start and end time of each action. For training, we thus extract human tubes only for the temporal extent of the actions. We then run multi-fold MIL to select the relevant tubes and train the detector. Additional temporal localization is beyond the scope of this paper and would require an additional step for pre-selecting relevant shots. At test time, we do not suppose temporal localization given and perform spatio-temporal localization on the entire test set. We first divide the videos into shots using a shot detector⁴. We generate human tubes inside each shot and perform a temporal sliding window inside each tube. Since videos/actions are longer, we add $\{900, 1200, 1500, 1800, 2400, 3000, \dots, 12000\}$ to the temporal scales used on UCF-101.

5.4 Evaluation of our MIL learning

Figure 7 shows an evaluation of our multi-fold MIL algorithm. We report the mAP@0.5 on the test set, as well as the CorLoc@0.5 on the training set defined as the ratio of selected tubes (*i.e.*, the tube with highest score in each positive video) which have an IoU with any ground-truth over 0.5. When using MIL, the CorLoc remains constant: the locations of the positive tubes do not change. This can be explained by the small number of positive videos and the high dimensional features. Multi-fold MIL overcomes this limitation, the CorLoc@0.5 increases and so does the mAP. The difference is moderate, as in many cases only one human is present in the test videos. An example for an improved selection of the human performing the action is shown in Figure 7 (right).

⁴ <https://github.com/johmathe/Shotdetect>

Learning	UCF-Sports	
	CorLoc@0.5 (training set)	mAP@0.5 (test set)
Initialization	80.14%	-
MIL ($K = 1$)	80.14%	83.38%
multi-fold MIL ($K = 4$)	81.99%	83.87%

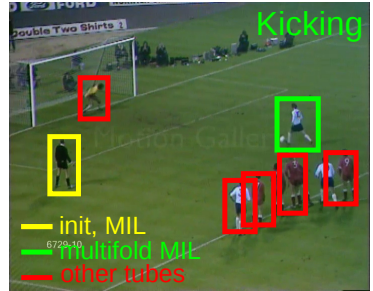
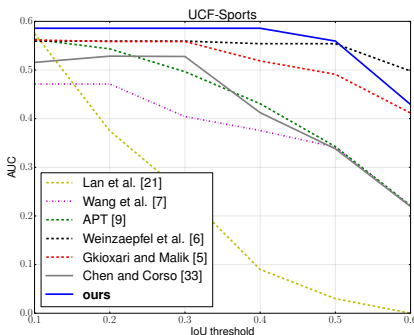


Fig. 7. *Left:* Comparison of multi-fold MIL to standard MIL on UCF-Sports. *Right:* an example where multi-fold MIL allows to select the correct tube. The yellow box is the tube with highest human score used as initialization. The location of the positive tube remains locked to this initialization when using MIL, while multi-fold MIL allows to move to the human performing the kicking action (green). Red boxes indicate other human tubes.

6 Experimental results

We have already assessed the quality of the human tubes in Section 4.3 and the impact of multi-fold MIL in Section 5.4. In this section, we evaluate the overall performance of our weakly-supervised action localization approach and compare to the state-of-the-art weakly- and fully-supervised approaches (Section 6.1). We then propose an evaluation of our approach on the DALY dataset (Section 6.2).

6.1 Comparison to the state of the art



	ours	[31]	[32]
Diving	56.2	43.7	44.3
Golf	50.6	52.3	50.5
Kicking	68.7	52.9	48.3
Lifting	62.4	63.5	41.4
Riding	63.5	32.5	30.6
Run	57.2	30.1	33.1
Skateboard	69.1	43.2	38.5
Swing1	64.3	57.5	54.3
Swing2	68.6	44.1	20.6
Walkg	68.3	47.1	39.0
avg.	62.9	46.7	41.0

Fig. 8. Evaluation on the UCF-Sports dataset. AUC at various IoU thresholds (*left*) and mean-IoU in percentage (*right*).

UCF-Sports. We first evaluate our method and compare to the state of the art on UCF-Sports. Most of previous works plot ROC curves at different IoU

thresholds until a false positive rate of 0.6 and report the Area Under the Curve (AUC). We show a comparison to the state of the art in Figure 8 (left). Plain lines indicate weakly-supervised methods and dashed ones full supervision. Our method outperforms all approaches based on this metric, even if they leverage supervision. We beat by a large margin the only other weakly supervised approach [33] which reports results for this metric. Interestingly, our performance decreases only slowly with respect to the IoU threshold. This indicates the high IoU of our detections with respect to the ground-truth, thanks to accurate human detection and tracking. We have also evaluated the mean-IoU, which is the average IoU between the best detection in a test video and the ground-truth tube. A comparison with other weakly-supervised approaches [31, 32] reporting this metric is shown in Figure 8 (right). We outperform them by more than 15%. Finally, we compare with the more meaningful mAP metric in Table 2. We obtain a mAP of 84% which is close to the state-of-the-art fully-supervised method [6] and better than [5], despite the fact that we use significantly less supervision.

Method	Supervision	UCF-Sports	J-HMDB	UCF-101		DALY
		mAP@0.5	mAP@0.5	mAP@0.05	mAP@0.2	mAP@0.2
Gkioxari and Malik [5]	fully	75.8%	53.3%	-	-	-
Weinzaepfel <i>et al.</i> [6]	fully	90.5%	60.7%	54.3%	46.8%	-
APT [9]	fully	-	-	58.0%	37.8%	-
Yu and Yuan [12]	fully	-	-	42.8%	-	-
ours	weakly	83.9%	54.1%	62.8%	45.4%	10.8%

Table 2. Comparison to the state of the art with mAP@0.5 on the UCF-Sports and J-HMDB datasets and mAP@0.2 on the UCF-101 (split 1) and DALY datasets. For UCF-101 we also report mAP@0.05 to compare to [12].

J-HMDB. Mean Average-Precision (mAP) on the J-HMDB dataset is shown in Table 2. Similar to the UCF-Sports dataset, our method is in between the two supervised methods from Weinzaepfel *et al.* [6] and Gkioxari and Malik [5]. Since the J-HMDB dataset contains only one human per video, the main challenge actually consists in classifying the tubes. Thus, a tube description that also leverages CNN features, in the same spirit as [5, 6] would probably increase the performance.

UCF-101. We report results with the standard mAP@0.2 metric for this dataset in Table 2. Our method also obtains a performance slightly below the state-of-the-art fully-supervised method [6]. We perform better than the 37.8% mAP@0.2 reported by van Gemert *et al.* [9], which extract thousands of proposals before using supervision for learning a classifier. Yu and Yuan [12] report a mAP@0.05 of 42.8% compared to our 62.8% at this threshold. At this low overlap threshold, we obtain better results than all other reported approaches (fully-supervised). This can be explained by the fact that our approach obtains less false positive detections due to the quality of our human tubes. Nevertheless, the IoU between

our detections and the ground-truth tubes are often below 0.2, due to the low quality of the videos which makes the human tube extraction challenging.

6.2 Evaluation on the DALY dataset




Fig. 9. Example of highly-ranked drinking detections on the DALY dataset. The two left examples are correct, whereas the two examples on the right correspond to confusions with playing harmonica and applying make up on lips.

Spatio-temporal action localization. For spatio-temporal action localization, we obtain a $\text{mAP}@0.2$ of 10.8% and a $\text{mAP}@0.1$ of 15.8%, see Tables 2 and 3. These results are obtained by training in a spatially weakly-supervised setting. Note that the IoU is defined as the temporal IoU multiplied by the spatial IoU obtained by averaging over the annotated keyframes, since spatial annotations are not available at every frame. The relatively low performance can be explained by several factors. First, temporal detection is difficult as actions are short in long untrimmed videos. Moreover, there are many actions that are close in time and our method tends to return one detection that covers multiple ground-truth short actions. Second, the different labels have similarities, leading to confusion between classes, as shown in Figure 9. For instance, drinking, playing harmonica and applying make up on lips all involve poses where the hands come near the mouth. The relatively low performance highlights the difficulty of the introduced DALY dataset and demonstrates that significant improvements are necessary to tackle challenging real-world videos.

In order to tease apart the challenges of the DALY dataset, we separately evaluate class confusion, temporal detection and spatio-temporal detection, see Table 3.

	$\text{mAP}@0.1$	$\text{mAP}@0.2$
clip classification	72.9%	
action localization in trimmed clips	49.7%	48.6%
action localization in shots with actions	23.5%	16.1%
action localization	15.8%	10.8%

Table 3. Detailed performance analysis on the DALY dataset. See text for more details.

Clip classification. We first evaluate the performance of clip classification, *i.e.* of video clips trimmed to the duration of the annotated actions. Each clip is

represented by a Fisher Vector of improved dense trajectories and a linear SVM is learned for clip classification. When learning a linear SVM for one class, we do not use the clips from other actions that temporally overlap with this class. For evaluation, we use the mAP metric since one clip can have multiple labels. As during training, we remove the clips from other actions that temporally overlap with this class during evaluation. We report a mAP of 73%. This is lower than the IDT+FV results for video classification on UCF-Sports (88%) and UCF-101 (88%), and in the same order of magnitude on J-HMDB (65%). Distinguishing the classes is not straightforward since many classes contain similar motion patterns (hand moving near the head).

Action localization in trimmed clips. We also evaluate action localization in these trimmed clips. This setting is similar to action localization datasets such as J-HMDB or UCF-Sports where clips are trimmed to the action. We obtain a mAP@0.2 of 48.6% and a mAP@0.5 of 31.7%. This relatively low performance can be explained by the fact that many instances are short and can easily be confused.

Action localization in shots containing actions. We run the spatio-temporal localization only in shots that contain an action. This removes, compared to the standard setting, the incorrect detections in shots that do not contain any action and reduces the total duration from 31 to 21 hours. The mAP@0.2 decreases to 16.1% when adding this temporal component, but it is above the result with all shots (10.8%). Once again, the challenge is to obtain a precise temporal localization without fragmenting long actions into multiple small temporal parts, which leads to false detections. Note that this is significantly different from the UCF-101 setting. Here, the shots can last for multiple minutes while the actions only last several seconds.

7 Conclusion

We present a novel approach for extracting human tubes that outperforms state-of-the-art proposals by a large margin. We show that they can be used effectively for weakly-supervised action localization, with a performance close to fully-supervised approaches. We also introduce a new challenging dataset, DALY, that overcomes the limitations of the existing datasets and will allow to measure progress in the field over the next few years. Future work includes improving the tracker with long-term constraints, enhancing the tube description, and dealing with temporal detection without supervision.

Acknowledgements

This work was partially supported by the advanced ERC “Allegro”, the MSR-Inria Joint Centre, as well as Google and Facebook Research Awards. We also acknowledge the NVIDIA Corporation for the donation of GPUs used for this research.

A The DALY Dataset

In this appendix, we describe how DALY was collected. The first part details action class selection, video filtering and spatio-temporal annotation of action instances. The second part presents per-class statistics for the dataset.

A.1 Dataset collection

Picking action classes. In order to allow precise annotation, we choose action classes with clearly defined temporal boundaries. For instance, the *brushing teeth* action is defined as “*toothbrush inside the mouth*”. Another example is *cleaning windows* for which the moment where “*the tool is in contact with the window*” is annotated.

Some of the classes are chosen to contain similar motion patterns, in order to make the class distinction difficult. Several of our action classes imply motion of the hands near the head (*taking photos*, *phoning*) or the mouth (*playing harmonica*, *drinking*, *brushing teeth*, *applying make up on lips*).

In summary, we kept the following 10 actions: *applying make up on lips*, *brushing teeth*, *cleaning floor*, *cleaning windows*, *drinking*, *folding textile*, *ironing*, *phoning*, *playing harmonica* and *taking photos/videos*, see Figure 10.



Fig. 10. One example frame for each of the 10 classes of the DALY dataset.

Video collection. The videos are gathered from YouTube using manually designed queries related to the selected action classes. For the class *cleaning floor*, the queries include “*sweeping floor*”, “*mopping floor*”, “*cleaning floor*”, etc. We only keep videos with a duration between 1 and 20 minutes. A minimum duration of 1 min ensures that temporal localization will be meaningful (shorter videos contain only one action from the beginning to the end in most cases), and a maximum duration of 20 minutes avoids issues related to computational time or memory consumption.

Videos are filtered to remove cartoons, slideshows, actions performed by animals and first-person viewpoints. We also remove videos in which the human is

not visible when the action occurs, for instance when the camera focuses on the mop while performing the *cleaning floor* action.

51 videos are selected for each action class such that they contain at least one instance of the action class. In total this corresponds to 31 hours of video, or 3.3 million frames. Videos from a given action class often contain multiple instances of the main action, and may contain instances of other action classes, which we annotate exhaustively.

Temporal annotation. Selected videos are carefully watched by members of our research team in order to catch all actions, including those happening in the background. The *begin* and *end* time is annotated for all instances found. Precise guidelines are established prior to annotation. For example, the *phoning* action lasts as long as the phone remains close to the actor’s ears. In case of a shot change during an action, we annotate it as two separate instances and set a “*shotcut*” flag. DALY contains 3724 action instances in total, with an average duration of 8 seconds.

Spatial annotation. There are more than 700k frames containing at least one action. Annotating all of them is a tall order, as it would take a year and a half for one annotator to complete this task, assuming 15 seconds per annotation and a 40-hour week. Of course this does not include verifying the results which, albeit quicker, is also time consuming when you need to outsource the job. Thus, we subsample frames to be annotated, such that enough information is present for a reliable evaluation of spatio-temporal detections.

For each temporal instance in the test set, we pick 5 uniformly sampled frames, with a maximum of 1 frame per second. For each frame, annotators are asked to draw a bounding box around the actor. Some of the spatial annotations are carried out by external workers. In all cases, the bounding boxes are reviewed and adjusted by members of our research team.

A.2 Dataset statistics

The selected action classes are sufficiently common such that multiple action classes can be found in a single video, see Figure 11, left. The matrix should be read row by row, each row considers only the 51 videos selected for a given class. For example, out of the 51 videos selected for the class *brushing teeth*, 51 videos contain *brushing teeth* instances, 25 contain *drinking* instances, etc.

Some classes have expected overlap such as *brushing teeth* and *drinking*, *ironing* and *folding textile*. There is also overlap between *taking photos/videos* and *phoning*, which can be explained by the fact that taking photos is mostly performed outdoors, where other people are *phoning*.

Figure 12 shows the histogram of video duration and instance duration. Per-class statistics are presented in Figure 11, right. One can see that most of the videos last several minutes (less than 10). Videos are longest in average for *ApplyingMakeUpOnLips*, mainly because this action is present in a multitude of full-face make-up tutorials. Concerning the instances, most of them are shorter than 10 seconds. Nevertheless, DALY also contains instances of several minutes, especially for actions such as *brushing teeth*, *playing harmonica* or *folding textile*.

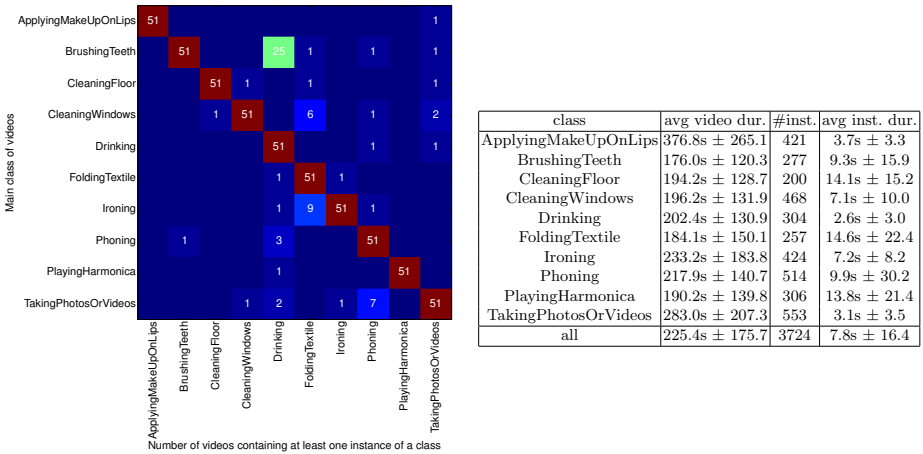


Fig. 11. *Left:* Statistics of multiple classes per video. Each row considers the 51 videos downloaded for a given class, each column counts the videos containing at least one instance of the column class. *Right:* Statistics for each class on the video duration (average and standard deviation), the number of instances, and the instance duration (average and standard deviation).

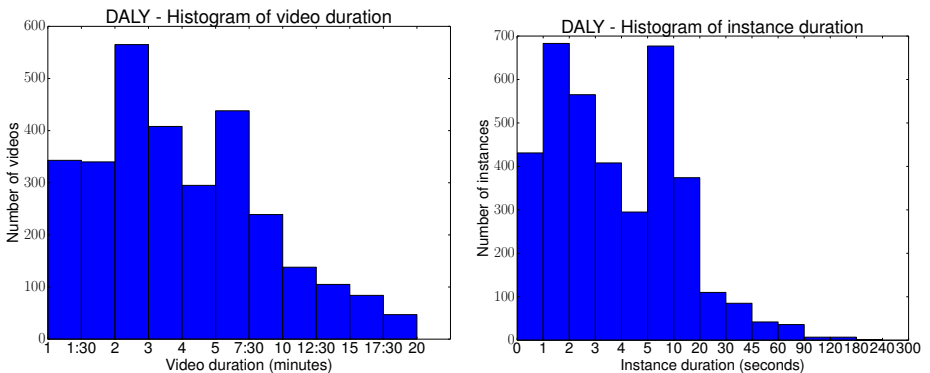


Fig. 12. Histogram of duration of the videos (left) and instances (right).

In some cases, short instance duration for actions can be explained by video editing. The uploader may cut the action to a few seconds and include it in a long video. Instances are shortest on average for *drinking* and *taking photos*, simply because drinking and taking a photo tend to take a short time (put the cup to the mouth and back for drinking, press the button to take a photo).

References

1. Wang, H., Oneata, D., Verbeek, J., Schmid, C.: A robust and efficient video representation for action recognition. *IJCV* (2015)
2. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: *NIPS*. (2014)
3. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: *ICCV*. (2015)
4. Yue-Hei Ng, J., Hausknecht, M., Vijayanarasimhan, S., Vinyals, O., Monga, R., Toderici, G.: Beyond short snippets: Deep networks for video classification. In: *CVPR*. (2015)
5. Gkioxari, G., Malik, J.: Finding action tubes. In: *CVPR*. (2015)
6. Weinzaepfel, P., Harchaoui, Z., Schmid, C.: Learning to track for spatio-temporal action localization. In: *ICCV*. (2015)
7. Wang, L., Qiao, Y., Tang, X.: Video action detection with relational dynamic-poselets. In: *ECCV*. (2014)
8. Marian Puscas, M., Sangineto, E., Culibrk, D., Sebe, N.: Unsupervised tube extraction using transductive learning and dense trajectories. In: *ICCV*. (2015)
9. van Gemert, J.C., Jain, M., Gati, E., Snoek, C.G.: Apt: Action localization proposals from dense trajectories. In: *BMVC*. (2015)
10. Kläser, A., Marszalek, M., Schmid, C., Zisserman, A.: Human Focused Action Localization in Video. In: *International Workshop on Sign, Gesture, and Activity (SGA)*. (2010)
11. Laptev, I., Pérez, P.: Retrieving actions in movies. In: *ICCV*. (2007)
12. Yu, G., Yuan, J.: Fast action proposals for human action detection and search. In: *CVPR*. (2015)
13. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: *NIPS*. (2015)
14. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: *CVPR*. (2014)
15. Hare, S., Saffari, A., Torr, P.: Struck: Structured output tracking with kernels. In: *ICCV*. (2011)
16. Kalal, Z., Mikolajczyk, K., Matas, J.: Tracking-learning-detection. *IEEE Trans. PAMI* (2012)
17. Cinbis, R.G., Verbeek, J., Schmid, C.: Weakly Supervised Object Localization with Multi-fold Multiple Instance Learning. *IEEE Trans. PAMI* (2016)
18. Cao, L., Liu, Z., Huang, T.S.: Cross-dataset action detection. In: *CVPR*. (2010)
19. Yuan, J., Liu, Z., Wu, Y.: Discriminative subvolume search for efficient action detection. In: *CVPR*. (2009)
20. Gaidon, A., Harchaoui, Z., Schmid, C.: Temporal Localization of Actions with Actoms. *IEEE Trans. PAMI* (2013)
21. Lan, T., Wang, Y., Mori, G.: Discriminative figure-centric models for joint action localization and recognition. In: *ICCV*. (2011)
22. Kläser, A., Marszaek, M., Schmid, C.: A spatio-temporal descriptor based on 3D-gradients. In: *BMVC*. (2008)
23. Jain, M., van Gemert, J.C., Jégou, H., Bouthemy, P., Snoek, C.G.M.: Action localization by tubelets from motion. In: *CVPR*. (2014)
24. Chen, W., Xiong, C., Xu, R., Corso, J.: Actionness ranking with lattice conditional ordinal random fields. In: *CVPR*. (2014)

25. Oneata, D., Revaud, J., Verbeek, J., Schmid, C.: Spatio-Temporal Object Detection Proposals. In: ECCV. (2014)
26. Bojanowski, P., Lajugie, R., Bach, F., Laptev, I., Ponce, J., Schmid, C., Sivic, J.: Weakly supervised action labeling in videos under ordering constraints. In: ECCV. (2014)
27. Duchenne, O., Laptev, I., Sivic, J., Bach, F., Ponce, J.: Automatic annotation of human actions in video. In: ICCV. (2009)
28. Hoai, M., Torresani, L., De la Torre, F., Rother, C.: Learning discriminative localization from weakly labeled data. Pattern Recognition (2014)
29. Siva, P., Xiang, T.: Weakly supervised action detection. In: BMVC. (2011)
30. Laptev, I.: On space-time interest points. IJCV (2005)
31. Mosabbeh, E.A., Cabral, R., De la Torre, F., Fathy, M.: Multi-label discriminative weakly-supervised human activity recognition and localization. In: ACCV. (2014)
32. Ma, S., Zhang, J., Ikizler-Cinbis, N., Sclaroff, S.: Action recognition and localization by hierarchical space-time segments. In: ICCV. (2013)
33. Chen, W., Corso, J.J.: Action detection by implicit intentional motion clustering. In: ICCV. (2015)
34. Shapovalova, N., Vahdat, A., Cannons, K., Lan, T., Mori, G.: Similarity constrained latent support vector machine: An application to weakly supervised action classification. In: ECCV. (2012)
35. Boyraz, H., Masood, S.Z., Liu, B., Tappen, M., Foroosh, H.: Action recognition by weakly-supervised discriminative region localization. In: BMVC. (2014)
36. Lan, T., Zhu, Y., Roshan Zamir, A., Savarese, S.: Action recognition by hierarchical mid-level action elements. In: ICCV. (2015)
37. Rodriguez, M.D., Ahmed, J., Shah, M.: Action mach: a spatio-temporal maximum average correlation height filter for action recognition. In: CVPR. (2008)
38. Jhuang, H., Gall, J., Zuffi, S., Schmid, C., Black, M.J.: Towards understanding action recognition. In: ICCV. (2013)
39. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. In: CRCV-TR-12-01. (2012)
40. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR. (2015)
41. Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J.: Image classification with the fisher vector: Theory and practice. IJCV (2013)