

# SF-Net: Single-Frame Supervision for Temporal Action Localization

Fan Ma<sup>1</sup>, Linchao Zhu<sup>1</sup>, Yi Yang<sup>1</sup>, Shengxin Zha<sup>2</sup>, Gourab Kundu<sup>2</sup>,  
Matt Feiszli<sup>2</sup>, and Zheng Shou<sup>2</sup>

<sup>1</sup> University of Technology Sydney, Australia

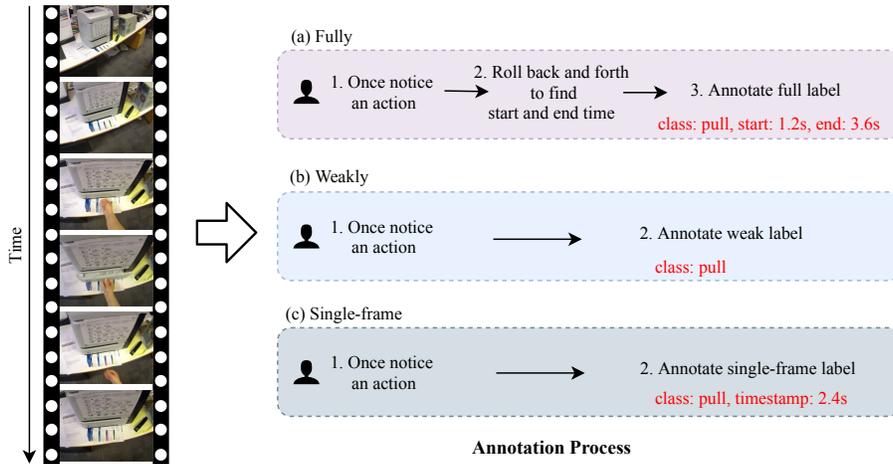
<sup>2</sup> Facebook AI

**Abstract.** In this paper, we study an intermediate form of supervision, i.e., **single-frame supervision**, for temporal action localization (TAL). To obtain the single-frame supervision, the annotators are asked to identify only a single frame *within* the temporal window of an action. This can significantly reduce the labor cost of obtaining full supervision which requires annotating the action *boundary*. Compared to the weak supervision that only annotates the video-level label, the single-frame supervision introduces extra temporal action signals while maintaining low annotation overhead. To make full use of such single-frame supervision, we propose a unified system called **SF-Net**. First, we propose to predict an actionness score for each video frame. Along with a typical category score, the actionness score can provide comprehensive information about the occurrence of a potential action and aid the temporal boundary refinement during inference. Second, we mine pseudo action and background frames based on the single-frame annotations. We identify pseudo action frames by adaptively expanding each annotated single frame to its nearby, contextual frames and we mine pseudo background frames from all the unannotated frames across multiple videos. Together with the ground-truth labeled frames, these pseudo-labeled frames are further used for training the classifier. In extensive experiments on THUMOS14, GTEA, and BEOID, SF-Net significantly improves upon state-of-the-art weakly-supervised methods in terms of both segment localization and single-frame localization. Notably, SF-Net achieves comparable results to its fully-supervised counterpart which requires much more resource intensive annotations. The code is available at <https://github.com/Flowerfan/SF-Net>.

**Keywords:** single-frame annotation action localization

## 1 Introduction

Recently, weakly-supervised Temporal Action Localization (TAL) has attracted substantial interest. Given a training set containing only video-level labels, we aim to detect and classify each action instance in long, untrimmed testing videos. In the fully-supervised annotation, the annotators usually need to rollback the



**Fig. 1.** Different ways of annotating actions while watching a video. (a) Annotating actions in the fully-supervised way. The start and end time of each action instance are required to be annotated. (b) Annotating actions in the weakly-supervised setting. Only action classes are required to be given. (c) Annotating actions in our single-frame supervision. Each action instance should have one timestamp. Note that the time is automatically generated by the annotation tool. Compared to the weakly-supervised annotation, the single-frame annotation requires only a few extra pauses to annotate repeated seen actions in one video.

video for repeated watching to give the precise temporal boundary of an action instance when they notice an action while watching the video [40]. For the weakly-supervised annotation, annotators just need to watch the video once to give labels. They can record the action class once they notice an unseen action. This significantly reduces annotation resources: video-level labels use fewer resources than annotating the start and end times in the fully-supervised setting.

Despite the promising results achieved by state-of-the-art weakly-supervised TAL work [24,26,28], their localization performance is still inferior to fully-supervised TAL work [5,17,27]. In order to bridge this gap, we are motivated to utilize single-frame supervision [22]: for each action instance, only one single positive frame is pointed out. The annotation process for single-frame supervision is almost the same as it in the weakly-supervised annotation. The annotators only watch the video once to record the action class and timestamp when they notice each action. It significantly reduces annotation resources compared to full supervision.

In the image domain, Bearman et al. [2] were the first to propose point supervision for image semantic segmentation. Annotating at point-level was extended by Mettes et al. [21] to video domain for spatio-temporal localization, where each action frame requires one spatial point annotation during training. Moltisanti et al. [22] further reduced the required resources by proposing single-frame super-

vision and developed a method that selects frames with a very high confidence score as pseudo action frames.

However, [22] was designed for whole video classification. In order to make full use of single-frame supervision for our TAL task, the unique challenge of **localizing temporal boundaries of actions** remains unresolved. To address this challenge in TAL, we make three innovations to improve the localization model’s capability in distinguishing background frames and action frames. First, we predict “actionness” at each frame which indicates the probability of being any actions. Second, based on the actionness, we investigate a novel background mining algorithm to determine frames that are likely to be the background and leverage these pseudo background frames as additional supervision. Third, when labeling pseudo action frames, besides the frames with high confidence scores, we aim to determine more pseudo action frames and thus propose an action frame expansion algorithm.

In addition, for many real-world applications, detecting precise start time and end time is overkill. Consider a reporter who wants to find some car accident shots in an archive of street camera videos: it is sufficient to retrieve a single frame for each accident, and the reporter can easily truncate clips of desired lengths. Thus, in addition to evaluating traditional segment localization in TAL, we also propose a new task called single-frame localization, which requires only localizing one frame per action instance.

In summary, our contributions are three-fold:

(1) To our best knowledge, this is the first work to use single-frame supervision for the challenging problem of localizing temporal boundaries of actions. We show that the single-frame annotation significantly saves annotation time compared to fully-supervised annotation.

(2) We find that single-frame supervision can provide strong cue about the background. Thus, from frames that are not annotated, we propose two novel methods to mine likely background frames and action frames, respectively. These likely background and action timestamps are further used as pseudo ground truth for training.

(3) We conduct extensive experiments on three benchmarks, and the performances on both segment localization and single-frame localization tasks are largely boosted.

## 2 Related Work

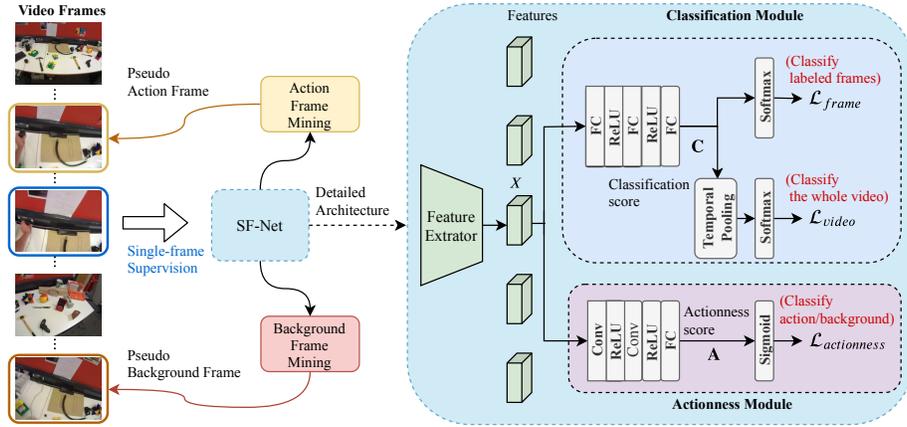
**Action recognition.** Action recognition has recently witnessed an increased focus on trimmed videos. Both temporal and spatial information is significant for classifying the video. Early works mainly employed hand-crafted features to solve this task. IDT [33] had been widely used across many video-related tasks. Recently, various deep neural networks were proposed to encode spatial-temporal video information. Two-stream network [30] adopted optical flow to learn temporal motion, which had been used in many latter works [3,30,35]. Many 3D convolutional networks [3,32,10] are also designed to learn action embeddings.

Beyond fully-supervised action recognition, a few works focus on self-supervised video feature learning [42] and few-shot action recognition [43]. In this paper, we focus on single-frame supervision for temporal action localization.

**Point supervision.** Bearman et al. [2] first utilized the point supervision for image semantic segmentation. Mettes et al. [21] extended it to spatio-temporal localization in video, where the action is pointed out by one spatial location in each action frame. We believe this is overkill for temporal localization, and demonstrate that single-frame supervision can achieve very promising results already. Recently, single-frame supervision has been used in [22] for video-level classification, but this work does not address identifying temporal boundaries. Note that Alwassel et al. [1] proposed to spot action in the video during inference time but targeted detecting one action instance per class in one video while our proposed single-frame localization task aims to detect every instance in one video.

**Fully-supervised temporal action localization.** Approaches of temporal action localization trained in full supervision have mainly followed a proposal-classification paradigm [5,7,11,17,29,27], where temporal proposals are generated first and then classified. Other categories of methods, including sequential decision-making [1] and single-shot detectors [16] have also been studied. Given full temporal boundary annotations, the proposal-classification methods usually filter out the background frames at the proposal stage via a binary actionness classifier. Activity completeness has also been studied in the temporal action localization task. Zhao et al. [41] used a structural temporal pyramid pooling followed by an explicit binary classifier to evaluate the completeness of an action instance. Yuan et al. [37] structured an action into three parts to model its temporal evolution. Chéron *et al.* [6] handled the spatio-temporal action localization with various supervisions. Long et al. [20] proposed a Gaussian kernel to dynamically optimize temporal scale of action proposals. However, these methods use fully temporal annotations, which are resource intensive.

**Weakly-supervised temporal action localization.** Multiple Instance Learning (MIL) has been widely used in weakly-supervised temporal action localization. Without temporal boundary annotations, temporal action score sequence has been widely used to generate action proposals [34,24,18,23]. Wang et al. [34] proposed UntrimmedNet composed of a classification module and a selection module to reason about the temporal duration of action instances. Nguyen *et al.* [24] introduced a sparsity regularization for video-level classification. Shou *et al.* [28] and Liu [19] investigated score contrast in the temporal dimension. Hide-and-Seek [31] randomly removed frame sequences during training to force the network to respond to multiple relevant parts. Liu *et al.* [18] proposed a multi-branch network to model the completeness of actions. Narayan *et al.* [23] introduced three-loss forms to guide the learning discriminative action features with enhanced localization capabilities. Nguyen [25] used attention modules to detect foreground and background for detecting actions. Despite the improvements over time, the performances of weakly-supervised methods are still inferior to the fully-supervised method.



**Fig. 2.** Overall training framework of our proposed SF-Net. Given single-frame supervision, we employ two novel frame mining strategies to label pseudo action frames and background frames. The detailed architecture of SF-Net is shown on the right. SF-Net consists of a classification module to classify each labeled frame and the whole video, and an actionness module to predict the probability of each frame being action. The classification module and actionness module are trained jointly with three losses explained in Sec. 3.3.

### 3 Method

In this section, we define our tasks, present architecture of our SF-Net, and finally discuss details of training and inference, respectively.

#### 3.1 Problem Definition

A training video can contain multiple action classes and multiple action instances. Unlike the full supervision setting, which provides temporal boundary annotation of each action instance, in our single-frame supervision setting, each instance only has one frame pointed out by annotator with timestamp  $t$  and action class  $y$ . Note that  $y \in \{1, \dots, N_c\}$  where  $N_c$  is the total number of classes and we use index 0 to represent the background class.

Given a testing video, we perform two temporal localization tasks: (1) **Segment localization.** We detect the start time and end time for each action instance with its action class prediction. (2) **Single-frame localization.** We output the timestamp of each detected action instance with its action class prediction. The evaluation metrics for these two tasks are explained in Sec. 4.

#### 3.2 Framework

**Overview.** Our overall framework is presented in Fig. 2. During training, learning from single-frame supervision, SF-Net mines pseudo action and background

frames. Based on the labeled frames, we employ three losses to jointly train a classification module to classify each labeled frame and the whole video, and an actionness module to predict the probability of each frame being action. In the following, we outline the framework while details of frame mining strategies and different losses are described in Sec. 3.3.

**Feature extraction.** For a training batch of  $N$  videos, the features of all frames are extracted and stored in a feature tensor  $X \in R^{N \times T \times D}$ , where  $D$  is the feature dimension, and  $T$  is the number of frames. As different videos vary in the temporal length, we simply pad zeros when the number of frames in a video is less than  $T$ .

**Classification module.** The classification module outputs the score of being each action class for all frames in the input video. To classify each labeled frame, we feed  $X$  into three Fully-Connected (FC) layers to get the classification score  $C \in \mathcal{R}^{N \times T \times N_c + 1}$ . The classification score  $C$  is then used to compute frame classification loss  $\mathcal{L}_{frame}$ . We also pool  $C$  temporally as described in [23] to compute video-level classification loss  $\mathcal{L}_{video}$ .

**Actionness module.** As shown in Fig. 2, our model has an actionness branch of identifying positive action frames. Different from the classification module, the actionness module only produces a scalar for each frame to denote the probability of being contained in an action segment. To predict an actionness score, we feed  $X$  into two temporal convolutional layers followed by one FC layer, resulting in an actionness score matrix  $A \in \mathcal{R}^{N \times T}$ . We apply sigmoid on  $A$  and then compute a binary classification loss  $\mathcal{L}_{actionness}$ .

### 3.3 Pseudo Label Mining and Training Objectives

**Action classification at labeled frames.** We use cross entropy loss for the action frame classification. As there are  $NT$  frames in the input batch of videos and most of the frames are unlabeled, we first filter the labeled frames for classification. Suppose we have  $K$  labeled frames where  $K \ll NT$ . We can get classification activations of  $K$  labeled frames from  $C$ . These scores are fed to a Softmax layer to get classification probability  $\mathbf{p}^l \in \mathcal{R}^{K \times N_c + 1}$  for all labeled frames. The classification loss of annotated frames in the batch of videos is formulated as:

$$\mathcal{L}_{frame}^l = -\frac{1}{K} \sum_i^K \mathbf{y}_i \log \mathbf{p}_i^l, \quad (1)$$

where the  $\mathbf{p}_i^l$  denote the prediction for the  $i^{th}$  labeled action frame.

**Pseudo labeling of frames.** With only a single label per action instance, the total number of positive examples is quite small and may be difficult to learn from. While we do not use full temporal annotation, it is clear that actions are longer events spanning consecutive frames. To increase the temporal information available to the model, we design an action frame mining and a background frame mining strategy to introduce more frames into the training process.

(a) Action frame mining: We treat each labeled action frame as an anchor frame for each action instance. We first set the expand radius  $r$  to limit the maximum expansion distance to the anchor frame at  $t$ . Then we expand the past from  $t - 1$  frame and the future from  $t + 1$  frame, separately. Suppose the action class of the anchor frame is represented by  $y_i$ . If the current expanding frame has the same predicted label with the anchor frame, and the classification score at  $y_i$  class is higher than that score of the anchor frame multiplying a predefined value  $\xi$ , we then annotate this frame with label  $y_i$  and put it into the training pool. Otherwise, we stop the expansion process for the current anchor frame.

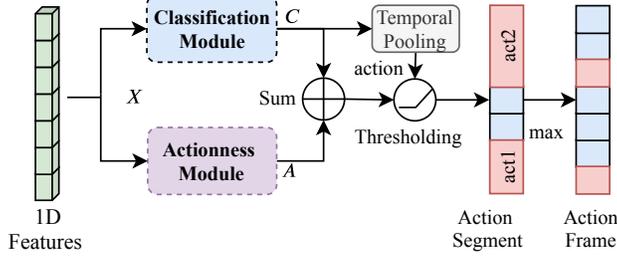
(b) Background frame mining: The background frames are also important and widely used in localization methods [18,25] to boost the model performance. Since there is no background label under the single-frame supervision, our proposed model manages to localize background frames from all the unlabeled frames in the  $N$  videos. At the beginning, we do not have supervision about where the background frames are. But explicitly introducing a background class can avoid forcing classifying a frame into one of the action classes. Our proposed background frame mining algorithm can offer us the supervision needed for training such a background class so as to improve the discriminability of the classifier. Suppose we try to mine  $\eta K$  background frames, we first gather the classification scores of all unlabeled frames from  $C$ . The  $\eta$  is the ratio of background frames to labeled frames. These scores are then sorted along background class to select the top  $\eta K$  scores  $\mathbf{p}^b \in \mathcal{R}^{\eta K}$  as the score vector of the background frames. The pseudo background classification loss is calculated on the top  $\eta K$  frames by,

$$\mathcal{L}_{frame}^b = -\frac{1}{\eta K} \sum \log \mathbf{p}^b, \quad (2)$$

The background frame classification loss assists the model with identifying irrelevant frames. Different from background mining in [18,25] which either require extra computation source to generate background frames or adopt a complicated loss for optimization, we mining background frames across multiple videos and use the classification loss for optimization. The selected pseudo background frames may have some noises in the initial training rounds. As the training evolves and the classifier’s discriminability improves, we are able to reduce the noises and detect background frames more correctly. With the more correct background frames as supervision signals, the classifier’s discriminability can be further boosted. In our experiments, we observed that this simple background mining strategy allows for better action localization results. We incorporate the background classification loss with the labeled frame classification loss to formulate the single-frame classification loss

$$\mathcal{L}_{frame} = \mathcal{L}_{frame}^l + \frac{1}{N_c} \mathcal{L}_{frame}^b \quad (3)$$

where  $N_c$  is the number of action classes to leverage the influence from background class.



**Fig. 3.** The inference framework of SF-Net. The classification module outputs the classification score  $C$  of each frame for identifying possible target actions in the given video. The action module produces the actionness score determining the possibility of a frame containing target actions. The actionness score together with the classification score are used to generate action segment based on the threshold.

**Actionness prediction at labeled frames.** In fully-supervised TAL, various methods learn to generate action proposals that may contain the potential activities [36,17,5]. Motivated by this, we design the actionness module to make the model focus on frames relevant to target actions. Instead of producing the temporal segment in proposal methods, our actionness module produces the actionness score for each frame. The actionness module is in parallel with the classification module in our SF-Net. It offers extra information for temporal action localization. We first gather the actionness score  $A^l \in \mathcal{R}^K$  of labeled frames in the training videos. The higher the value for a frame, the higher probability of that frame belongs to a target action. We also use the background frame mining strategy to get the actionness score  $A^b \in \mathcal{R}^{nK}$ . The actionness loss is calculated by,

$$\mathcal{L}_{actionness} = -\frac{1}{K} \sum \log \sigma(A^l) - \frac{1}{\eta K} \sum \log(1 - \sigma(A^b)), \quad (4)$$

where  $\sigma$  is the sigmoid function to scale the actionness score to  $[0, 1]$ .

**Full objective.** We employ video-level loss as described in [23] to tackle the problem of multi-label action classification at video-level. For the  $i^{th}$  video, the top-k activations per category (where  $k = T_i/8$ ) of the classification activation  $C(i)$  are selected and then are averaged to obtain a class-specific encoding  $r_i \in \mathcal{R}^{C+1}$  as in [26,23]. We average all the frame label predictions in the video  $v_i$  to get the video-level ground-truth  $q_i \in \mathcal{R}^{N_c+1}$ . The video-level loss is calculated by

$$\mathcal{L}_{video} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{N_c} q_i(j) \log \frac{\exp(r_i(j))}{\sum_{N_c+1} \exp(r_i(k))}, \quad (5)$$

where  $q_i(j)$  is the  $j^{th}$  value of  $q_i$  representing the probability mass of video  $v_i$  belong to  $j^{th}$  class.

Consequently, the total training objective for our proposed method is

$$\mathcal{L} = \mathcal{L}_{frame} + \alpha\mathcal{L}_{video} + \beta\mathcal{L}_{actionness}, \quad (6)$$

where  $\mathcal{L}_{frame}$ ,  $\mathcal{L}_{video}$ , and  $\mathcal{L}_{actionness}$ , denote the frame classification loss, video classification loss, and actionness loss, respectively.  $\alpha$  and  $\beta$  are the hyper-parameters leveraging different losses.

### 3.4 Inference

During the test stage, we need to give the temporal boundary for each detected action. We follow previous weakly-supervised work [24] to predict video-level labels by temporally pooling and thresholding on the classification score. As shown in Fig. 3, we first obtain the classification score  $C$  and actionness score  $A$  by feeding input features of a video to the classification module and actionness module. Towards segment localization, we follow the thresholding strategy in [24,25] to keep the action frames above the threshold and consecutive action frames constitute an action segment. For each predicted video-level action class, we localize each action segment by detecting an interval that the sum of classification score and actionness score exceeds the preset threshold at every frame inside the interval. We simply set the confidence score of the detected segment to the sum of its highest frame classification score and the actionness score. Towards single frame localization, for the action instance, we choose the frame with the maximum activation score in the detected segment as the localized action frame.

## 4 Experiment

### 4.1 Datasets

**THUMOS14.** There are 1010 validation and 1574 test videos from 101 action categories in THUMOS14 [12]. Out of these, 20 categories have temporal annotations in 200 validation and 213 test videos. The dataset is challenging, as it contains an average of 15 activity instances per video. Similar to [14, 16], we use the validation set for training and test set for evaluating our framework.

**GTEA.** There are 28 videos of 7 fine-grained types of daily activities in a kitchen contained in GTEA [14]. An activity is performed by four different subjects, and each video contains about 1800 RGB frames, showing a sequence of 7 actions, including the background action.

**BEOID.** There are 58 videos in BEOID [8]. There is an average of 12.5 action instances per video. The average length is about 60s, and there are 30 action classes in total. We randomly split the untrimmed videos in an 80-20% proportion for training and testing, as described in [22].

**Table 1.** Comparison between different methods for simulating single-frame supervision on THUMOS14. “Annotation” means that the model uses human annotated frame for training. “TS” denotes that the single-frame is sampled from action instances using a uniform distribution, while “TS in GT” is using a Gaussian distribution near the mid timestamp of each activity. The AVG for segment localization is the average mAP from IoU 0.1 to 0.7.

Position	mAP@hit	Segment mAP@IoU			
		0.3	0.5	0.7	AVG
Annotation	<b>60.2±0.70</b>	<b>53.3±0.30</b>	28.8±0.57	9.7±0.35	<b>40.6±0.40</b>
TS	57.6±0.60	52.0±0.35	<b>30.2±0.48</b>	<b>11.8±0.35</b>	40.5±0.28
TS in GT	52.8±0.85	47.4±0.72	26.2±0.64	9.1±0.41	36.7±0.52

## 4.2 Implementation Details

We use I3D network [3] trained on the Kinetics [4] to extract video features. For the RGB stream, we rescale the smallest dimension of a frame to 256 and perform the center crop of size  $224 \times 224$ . For the flow stream, we apply the TV-L1 optical flow algorithm [38]. We follow the two-stream fusion operation in [23] to integrate predictions from both appearance (RGB) and motion (Flow) branches. The inputs to the I3D models are stacks of 16 frames.

On all datasets, we set the learning rate to  $10^{-3}$  for all experiments, and the model is trained with a batch size of 32 using the Adam [13]. Loss weight hyper-parameters  $\alpha$  and  $\beta$  are set to 1. The model performance is not sensitive to these hyper-parameters. For the hyper-parameter  $\eta$  used in mining background frames, we set it to 5 on THUMOS14 and set it to 1 on the other two datasets. The number of iterations is set to 500, 2000 and 5000 for GTEA, BEOID and THUMOS14, respectively.

## 4.3 Evaluation Metrics

(1) **Segment localization:** We follow the standard protocol, provided with the three datasets, for evaluation. The evaluation protocol is based on mean Average Precision (mAP) for different intersection over union (IoU) values for the action localization task.

(2) **Single-frame localization:** We also use mAP to compare performances. Instead of measuring IoU, the predicted single-frame is regarded as correct when it lies in the temporal area of the ground-truth segment, and the class label is correct. We use mAP@hit to denote the mean average precision of selected action frame falling in the correct action segment.

## 4.4 Annotation Analysis

**Single-frame supervision simulation.** First, to simulate the single-frame supervision based on ground-truth boundary annotations existed in the above three datasets, we explore the different strategies to sample a single-frame for

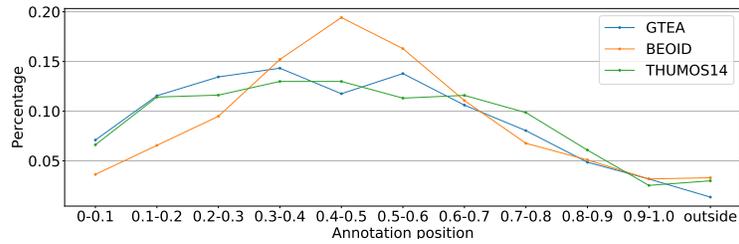
**Table 2.** Single-frame annotation differences between different annotators on three datasets. We show the number of action segments annotated by Annotator 1, Annotator 2, Annotator 3, and Annotator 4. In the last column, we report the total number of the ground-truth action segments for each dataset.

Datasets	Annotator 1	Annotator 2	Annotator 3	Annotator 4	# of total segments
GTEA	369	366	377	367	367
BEOID	604	602	589	599	594
THUMOS14	3014	2920	2980	2986	3007

each action instance. We follow the strategy in [22] to generate single-frame annotations with uniform and Gaussian distribution (**Denoted by TS and TS in GT**). We report the segment localization at different IoU thresholds and frame localization results on THUMOS14 in Table 1. The model with each single-frame annotation is trained five times. The mean and standard deviation of mAP is reported in the Table. Compared to models trained on sampled frames, the model trained on human annotated frames achieves the highest mAP@hit. As the the action frame is the frame with the largest prediction score in the prediction segment, the model with higher mAP@hit can assist with localizing action timestamp more accurately when people need to retrieve the frame of target actions. When sampling frames are from near middle timestamps to the action segment (TS in GT), the model performs inferior to other models as these frames may not contain informative elements of complete actions. For the segment localization result, the model trained on truly single-frame annotations achieves higher mAP at small IoU thresholds, and the model trained on frames sampled uniformly from the action instance gets higher mAP at larger IoU thresholds. It may be originated by sampled frames of uniform distribution containing more boundary information for the given action instances.

**Single-frame annotation.** We also invite four annotators with different backgrounds to label a single frame for each action segment on three datasets. More details of annotation process can be found in the supplementary material. In Table 2, we have shown the action instances of different datasets annotated by different annotators. The ground-truth in the Table denotes the action instances annotated in the fully-supervised setting. From the Table, we obtain that the number of action instances by different annotators have a very low variance. The number of labeled frames is very close to the number of action segments in the fully-supervised setting. This indicates that annotators have common justification for the target actions and hardly miss the action instance despite that they only pause once to annotate single-frame of each action.

We also present the distribution of the relative position of single-frame annotation to the corresponding action segment. As shown in Fig. 4, there are rare frames outside of the temporal range of action instances from the ground-truth in the fully-supervised setting. As the number of annotated single frames is almost the same as the number of action segments, we can draw the inference that



**Fig. 4.** Statistics of human annotated single-frame on three datasets. X-axis: single-frame falls in the relative portion of the whole action; Y-axis: percentage of annotated frames. We use different colors to denote annotation distribution on different datasets.

the single frame annotation includes all almost potential action instances. We obtain that annotators prefer to label frames near to the middle part of action instances. This indicates that humans can identify an action without watching the whole action segment. On the other hand, this will significantly reduce the annotation time compared with fully-supervised annotation as we can quickly skip the current action instance after single-frame annotation.

**Annotation speed for different supervision.** To measure the required annotation resource for different supervision, we conduct a study on GTEA. Four annotators are trained to be familiar with action classes in GTEA. We ask the annotator to indicate the video-level label, single-frame label and temporal boundary label of 21 videos lasting 93 minutes long. While watching, the annotator is able to skim quickly, pause, and go to any timestamp. On average, the annotation time used by each person to annotate 1-minute video is 45s for the video-level label, 50s for the single-frame label, and 300s for the segment label. The annotation time for single-frame label is close to the annotation time for video-level label but much fewer than time for the fully-supervised annotation.

## 4.5 Analysis

**Effectiveness of each module, loss, and supervision.** To analyze the contribution of the classification module, actionness module, background frame mining strategy, and the action frame mining strategy, we perform a set of ablation studies on THUMOS14, GTEA and BEOID datasets. The segment localization mAP at different thresholds is presented in Table 3. We also compare the model with only weak supervision and the model with full supervision. The model with weak supervision is implemented based on [23].

We observe that the model with single-frame supervision outperforms the weakly-supervised model. And large performance gain is obtained on GTEA and BEOID datasets as the single video often contains multiple action classes, while action classes in one video are fewer in THUMOS14. Both background frame mining strategy and action frame mining strategy boost the performance

**Table 3.** Segment localization mAP results at different IoU thresholds on three datasets. Weak denotes that only video-level labels are used for training. All action frames are used in the full supervision approach. SF uses extra single frame supervision with frame level classification loss. SFB means that pseudo background frames are added into the training, while the SFBA adopts the actionness module, and the SFBAE indicates the action frame mining strategy added in the model. For models trained on single-frame annotations, we report mean and standard deviation results of five runs. AVG is the average mAP from IoU 0.1 to 0.7.

Dataset	Models	mAP@IoU				
		0.1	0.3	0.5	0.7	AVG
GTEA	Full	58.1	40.0	22.2	14.8	31.5
	Weak	14.0	9.7	4.0	3.4	7.0
	SF	50.0±1.42	35.6±2.61	<b>21.6±1.67</b>	<b>17.7±0.96</b>	30.5±1.23
	SFB	52.9±3.84	34.9±4.72	17.2±3.46	11.0±2.52	28.0±3.53
	SFBA	52.6±5.32	32.7±3.07	15.3±3.63	8.5±1.95	26.4±3.61
	SFBAE	<b>58.0±2.83</b>	<b>37.9±3.18</b>	19.3±1.03	11.9±3.89	<b>31.0±1.63</b>
BEOID	Full	65.1	38.6	22.9	7.9	33.6
	Weak	22.5	11.8	1.4	0.3	8.7
	SF	54.1±2.48	24.1±2.37	6.7±1.72	1.5±0.84	19.7±1.25
	SFB	57.2±3.21	26.8±1.77	9.3±1.94	1.7±0.68	21.7±1.43
	SFBA	<b>62.9±1.68</b>	36.1±3.17	12.2±3.15	2.2±2.07	27.1±1.44
	SFBAE	<b>62.9±1.39</b>	<b>40.6±1.8</b>	<b>16.7±3.56</b>	<b>3.5±0.25</b>	<b>30.1±1.22</b>
THUMOS14	Full	68.7	54.5	34.4	16.7	43.8
	Weak	55.3	40.4	20.4	7.3	30.8
	SF	58.6±0.56	41.3±0.62	20.4±0.55	6.9±0.33	31.7±0.41
	SFB	60.8±0.65	44.5±0.37	22.9±0.38	7.8±0.46	33.9±0.31
	SFBA	68.7±0.33	52.3±1.21	28.2±0.42	<b>9.7±0.51</b>	39.9±0.43
	SFBAE	<b>70.0±0.64</b>	<b>53.3±0.3</b>	<b>28.8±0.57</b>	<b>9.7±0.35</b>	<b>40.6±0.40</b>

on BEOID and THUMOS14 by putting more frames into the training, the performance on GTEA decreases mainly due to that GTEA contains almost no background frame. In this case, it is not helpful to employ background mining and the actionness module which aims for distinguishing background against action. The actionness module works well for the BEOID and THUMOS14 datasets, although the actionness module only produces one score for each frame.

**Comparisons with state-of-the-art.** Experimental results on THUMOS14 testing set are shown in Table 4. Our proposed single-frame action localization method is compared to existing methods for weakly-supervised temporal action localization, as well as several fully-supervised ones. Our model outperforms the previous weakly-supervised methods at all IoU thresholds regardless of the choice of feature extraction network. The gain is substantial even though only one single-frame for each action instance is provided. The model trained on human annotated frames achieves higher mAP at lower IoU compared to model trained

**Table 4.** Segment localization results on THUMOS14 dataset. The mAP values at different IoU thresholds are reported, and the column AVG indicates the average mAP at IoU thresholds from 0.1 to 0.5. \* denotes the single-frame labels are simulated based on the ground-truth annotations. # denotes single-frame labels are manually annotated by human annotators.

Supervision	Method	mAP @IoU							
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	AVG
Full	S-CNN [29]	47.7	43.5	36.3	28.7	19.0	-	5.3	35.0
Full	CDC [27]	-	-	40.1	29.4	23.3	-	7.9	-
Full	R-C3D [36]	54.5	51.5	44.8	35.6	28.9	-	-	43.1
Full	SSN [41]	60.3	56.2	50.6	40.8	29.1	-	-	47.4
Full	Faster- [5]	59.8	57.1	53.2	48.5	42.8	<b>33.8</b>	<b>20.8</b>	52.3
Full	BMN [15]	-	-	56.0	47.4	38.8	29.7	20.5	-
Full	P-GCN [39]	<b>69.5</b>	<b>67.8</b>	<b>63.6</b>	<b>57.8</b>	<b>49.1</b>	-	-	<b>61.6</b>
Weak	Hide-and-Seek [31]	36.4	27.8	19.5	12.7	6.8	-	-	20.6
Weak	UntrimmedNet [34]	44.4	37.7	28.2	21.1	13.7	-	-	29.0
Weak	W-TALC [9]	49.0	42.8	32.0	26.0	18.8	-	6.2	33.7
Weak	AutoLoc [28]	-	-	35.8	29.0	21.2	13.4	5.8	-
Weak	STPN [24]	52.0	44.7	35.5	25.8	16.9	9.9	4.3	35.0
Weak	W-TALC [26]	55.2	49.6	40.1	31.1	22.8	-	7.6	39.7
Weak	Liu <i>et al.</i> [18]	57.4	50.8	41.2	32.1	23.1	15.0	7.0	40.9
Weak	Nguyen <i>et al.</i> [25]	<b>60.4</b>	<b>56.0</b>	<b>46.6</b>	<b>37.5</b>	<b>26.8</b>	<b>17.6</b>	<b>9.0</b>	<b>45.5</b>
Weak	3C-Net [23]	59.1	53.5	44.2	34.1	26.6	-	8.1	43.5
Single-frame simulation*	Moltisanti <i>et al.</i> [22]	24.3	19.9	15.9	12.5	9.0	-	-	16.3
Single-frame simulation*	SF-Net	68.3	62.3	52.8	<b>42.2</b>	<b>30.5</b>	<b>20.6</b>	<b>12.0</b>	51.2
Single-frame#	SF-Net	<b>71.0</b>	<b>63.4</b>	<b>53.2</b>	40.7	29.3	18.4	9.6	<b>51.5</b>

on sampling frames uniformly from action segments. The differences come from the fact that the uniform sampling frames from ground-truth action segments contain more information about temporal boundaries for different actions. As there are many background frames in the THUMOS14 dataset, the single frame supervision assists the proposed model with localizing potential action frames among the whole video. Note that the supervised methods have the regression module to refine the action boundary, while we simply threshold on the score sequence and still achieve comparable results.

## 5 Conclusions

In this paper, we have investigated how to leverage single-frame supervision to train temporal action localization models for both segment localization and single-frame localization during inference. Our SF-Net makes full use of single-frame supervision by predicting actionness score, pseudo background frame mining and pseudo action frame mining. SF-Net significantly outperforms weakly-supervised methods in terms of both segment localization and single-frame localization on three standard benchmarks.

**Acknowledgements.** This research was partially supported by ARC DP200100938 and Facebook.

## References

1. Alwassel, H., Caba Heilbron, F., Ghanem, B.: Action search: Spotting actions in videos and its application to temporal action localization. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 251–266 (2018)
2. Bearman, A., Russakovsky, O., Ferrari, V., Fei-Fei, L.: What’s the point: Semantic segmentation with point supervision. In: European conference on computer vision. pp. 549–565. Springer (2016)
3. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
4. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
5. Chao, Y.W., Vijayanarasimhan, S., Seybold, B., Ross, D.A., Deng, J., Sukthankar, R.: Rethinking the faster r-cnn architecture for temporal action localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1130–1139 (2018)
6. Chéron, G., Alayrac, J.B., Laptev, I., Schmid, C.: A flexible model for training action localization with varying levels of supervision. In: Advances in Neural Information Processing Systems. pp. 942–953 (2018)
7. Dai, X., Singh, B., Zhang, G., Davis, L.S., Qiu Chen, Y.: Temporal context network for activity localization in videos. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5793–5802 (2017)
8. Damen, D., Leelasawassuk, T., Haines, O., Calway, A., Mayol-Cuevas, W.W.: Youdo, i-learn: Discovering task relevant objects and their modes of interaction from multi-user egocentric video. In: BMVC. vol. 2, p. 3 (2014)
9. Ding, L., Xu, C.: Weakly-supervised action segmentation with iterative soft boundary assignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6508–6516 (2018)
10. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
11. Gao, J., Yang, Z., Nevatia, R.: Cascaded boundary regression for temporal action detection. arXiv preprint arXiv:1705.01180 (2017)
12. Idrees, H., Zamir, A.R., Jiang, Y.G., Gorban, A., Laptev, I., Sukthankar, R., Shah, M.: The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding* **155**, 1–23 (2017)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
14. Lei, P., Todorovic, S.: Temporal deformable residual networks for action segmentation in videos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6742–6751 (2018)
15. Lin, T., Liu, X., Li, X., Ding, E., Wen, S.: Bmn: Boundary-matching network for temporal action proposal generation. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
16. Lin, T., Zhao, X., Shou, Z.: Single shot temporal action detection. In: Proceedings of the 25th ACM international conference on Multimedia. pp. 988–996. ACM (2017)

17. Lin, T., Zhao, X., Su, H., Wang, C., Yang, M.: Bsn: Boundary sensitive network for temporal action proposal generation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 3–19 (2018)
18. Liu, D., Jiang, T., Wang, Y.: Completeness modeling and context separation for weakly supervised temporal action localization. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
19. Liu, Z., Wang, L., Zhang, Q., Gao, Z., Niu, Z., Zheng, N., Hua, G.: Weakly supervised temporal action localization through contrast based evaluation networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3899–3908 (2019)
20. Long, F., Yao, T., Qiu, Z., Tian, X., Luo, J., Mei, T.: Gaussian temporal awareness networks for action localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 344–353 (2019)
21. Mettes, P., Van Gemert, J.C., Snoek, C.G.: Spot on: Action localization from pointly-supervised proposals. In: European conference on computer vision. pp. 437–453. Springer (2016)
22. Moltisanti, D., Fidler, S., Damen, D.: Action recognition from single timestamp supervision in untrimmed videos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9915–9924 (2019)
23. Narayan, S., Cholakkal, H., Khan, F.S., Shao, L.: 3c-net: Category count and center loss for weakly-supervised action localization. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
24. Nguyen, P., Liu, T., Prasad, G., Han, B.: Weakly supervised action localization by sparse temporal pooling network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6752–6761 (2018)
25. Nguyen, P.X., Ramanan, D., Fowlkes, C.C.: Weakly-supervised action localization with background modeling. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
26. Paul, S., Roy, S., Roy-Chowdhury, A.K.: W-talc: Weakly-supervised temporal activity localization and classification. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 563–579 (2018)
27. Shou, Z., Chan, J., Zareian, A., Miyazawa, K., Chang, S.F.: Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
28. Shou, Z., Gao, H., Zhang, L., Miyazawa, K., Chang, S.F.: Autoloc: Weakly-supervised temporal action localization in untrimmed videos. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 154–171 (2018)
29. Shou, Z., Wang, D., Chang, S.F.: Temporal action localization in untrimmed videos via multi-stage cnns. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
30. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: Advances in neural information processing systems. pp. 568–576 (2014)
31. Singh, K.K., Lee, Y.J.: Hide-and-peek: Forcing a network to be meticulous for weakly-supervised object and action localization. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 3544–3553. IEEE (2017)
32. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: The IEEE International Conference on Computer Vision (ICCV) (December 2015)

33. Wang, H., Schmid, C.: Action recognition with improved trajectories. In: The IEEE International Conference on Computer Vision (ICCV) (December 2013)
34. Wang, L., Xiong, Y., Lin, D., Van Gool, L.: Untrimmednets for weakly supervised action recognition and detection. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 4325–4334 (2017)
35. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Val Gool, L.: Temporal segment networks: Towards good practices for deep action recognition. In: ECCV (2016)
36. Xu, H., Das, A., Saenko, K.: R-c3d: Region convolutional 3d network for temporal activity detection. In: Proceedings of the IEEE international conference on computer vision. pp. 5783–5792 (2017)
37. Yuan, Z., Stroud, J.C., Lu, T., Deng, J.: Temporal action localization by structured maximal sums. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3684–3692 (2017)
38. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime tv-l 1 optical flow. In: Joint pattern recognition symposium. pp. 214–223. Springer (2007)
39. Zeng, R., Huang, W., Tan, M., Rong, Y., Zhao, P., Huang, J., Gan, C.: Graph convolutional networks for temporal action localization. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
40. Zhao, H., Torralba, A., Torresani, L., Yan, Z.: Hacs: Human action clips and segments dataset for recognition and temporal localization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 8668–8678 (2019)
41. Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., Lin, D.: Temporal action detection with structured segment networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2914–2923 (2017)
42. Zhu, L., Yang, Y.: Actbert: Learning global-local video-text representations. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
43. Zhu, L., Yang, Y.: Label independent memory for semi-supervised few-shot video classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020). <https://doi.org/10.1109/TPAMI.2020.3007511>