

# Robust Action Segmentation from Timestamp Supervision

Yaser Souri\*<sup>1, 2</sup>

yasersouri@microsoft.com

Yazan Abu Farha\*<sup>1, 3</sup>

yabufarha@birzeit.edu

Emad Bahrami\*<sup>1</sup>

bahrami@iai.uni-bonn.de

Gianpiero Francesca<sup>4</sup>

gianpiero.francesca@toyota-europe.com

Juergen Gall<sup>1</sup>

gall@iai.uni-bonn.de

<sup>1</sup> Computer Vision Group

University of Bonn

Bonn, Germany

<sup>2</sup> Microsoft

Redmond, United States

<sup>3</sup> Birzeit University

Birzeit, West Bank, Palestine

<sup>4</sup> Toyota Motor Europe

Brussels, Belgium

\* indicates equal contribution

---

## Abstract

Action segmentation is the task of predicting an action label for each frame of an untrimmed video. As obtaining annotations to train an approach for action segmentation in a fully supervised way is expensive, various approaches have been proposed to train action segmentation models using different forms of weak supervision, *e.g.*, action transcripts, action sets, or more recently timestamps. Timestamp supervision is a promising type of weak supervision as obtaining one timestamp per action is less expensive than annotating all frames, but it provides more information than other forms of weak supervision. However, previous works assume that every action instance is annotated with a timestamp, which is a restrictive assumption since it assumes that annotators do not miss any action. In this work, we relax this restrictive assumption and take missing annotations for some action instances into account. We show that our approach is more robust to missing annotations compared to other approaches and various baselines.

## 1 Introduction

Action segmentation is an important task for many applications like home monitoring systems [5], worker monitoring and guidance in the assembly line [19], or tutorial generation from instructional videos [69, 43]. It requires to identify for each frame in an untrimmed video what action is happening. Recently, fully supervised approaches for action segmentation [10, 27, 36, 42] have achieved very good results. However, a major bottleneck that prevents the widespread adoption of action segmentation technologies is the cost of obtaining annotations for fully supervised learning. Annotating every instance of all actions of interest with their exact temporal boundaries is a labor-intensive endeavor. In particular, the start and end of an action is subjective and can lead to annotation inconsistencies.




Figure 1: For an input video with 3 actions A, B, and C, the annotator has correctly annotated a timestamp (black line) for the actions A and C, but action B has been missed. In this setting, Li *et al.* [28] assume that all actions are annotated and generate a labeling from the annotated timestamps that is incorrect for the entire duration of action B. We propose an approach that takes into account the possibility of missing timestamps and correctly ignores the frames (gray) corresponding to the missed action B.

To address these issues, several approaches have been proposed that train networks for action segmentation using only weak supervision, *e.g.*, in form of action sequences without any time information [26, 32, 37] or video tags [9, 22, 33]. While weakly supervised learning significantly reduces the annotation effort, these approaches still perform considerably worse than fully supervised approaches. Recently, [50] proposed to annotate only one frame, called timestamp, per action and this type of annotation has been also used to train a network for action segmentation in [28]. Timestamp supervision is a form of weak supervision where the annotator annotates only a single frame and its class for each action in a video. While it requires to annotate only a very small fraction of the frames, this type of weak supervision provides a form of temporal information that is crucial for action segmentation and that is missing in other forms of weak supervision. While the cost of obtaining timestamps is only marginally higher than other annotations for weakly supervised learning, [28] showed that timestamp supervision results in a much higher action segmentation performance.

Despite the promising results, [28] assumes that every action instance in a video is annotated by a timestamp since the approach detects an action change between two consecutive timestamps and labels the frames before and after the detected change based on the action classes of the corresponding timestamps as it is illustrated Figure 1. Annotators, however, can easily miss an action. In this case, the frames of the missed action are wrongly annotated. Since the generated labels are then used for training a network for action segmentation, the accuracy of the network drastically decreases when only 80-90% of all action instances are annotated as shown in Figure 2.

In this paper, we thus address this limitation and propose an approach that is designed to deal with missing timestamps. In contrast to [28], our approach can label frames as unknown as illustrated in Figure 1. These frames will then be ignored when the network for action segmentation is trained. To generate a labeling of a training video from a few annotated timestamps, we optimize the segment boundaries for all timestamps jointly. During the optimization, we expand the segment boundaries from the timestamps in each direction such that the network is confident that the frames within a segment belong to the same class as the timestamp and such that the number of frames that are unknown is minimized. This is also illustrated in Figure 3.

We evaluate our approach on two action segmentation benchmarks and show that our approach is considerably more robust to missing timestamps compared to [28] and various baselines. On the 50Salads [58] dataset, the accuracy of our approach drops by only 4.5% when the percentage of annotated timestamps is reduced from 95% to 70% compared to

a drop of 18.2% for [28]. We further evaluate our approach on two datasets for action localization. While the goal of the proposed approach is to increase the robustness to missing timestamps, the approach also performs well when all action instances are annotated.

## 2 Related Work

**Fully Supervised Action Segmentation.** Fully supervised approaches rely on frame-level annotations during training. Earlier methods in this setup applied a sliding window approach with non-maximum suppression [14, 35]. Other approaches used context-free grammars [51, 40] or hidden Markov models (HMMs) [16, 18, 20] to capture long-range dependencies. Recently, temporal convolutional networks (TCNs) with large receptive fields have been very successful in segmenting actions in long videos [1, 21, 23, 27]. Building on the success of these approaches, several approaches have been proposed to refine the TCN predictions using graph convolutional networks [12], boundary-aware pooling [13, 41], or hierarchical modeling [2]. In [11], a neural network architecture search scheme is used to select the dilation factors for the TCN layers. Recently, [42] proposed a transformer-based architecture for the temporal action segmentation task. In contrast to these approaches, our approach relies on a weaker level of supervision in the form of timestamps.

**Weakly Supervised Action Segmentation.** While fully supervised action segmentation approaches achieve very good results, they rely on dense frame-level annotations. As acquiring such annotations is time-consuming and expensive, many approaches have been proposed to train action segmentation models using weaker forms of supervision. A popular form of weak supervision are sequences of actions, which describe the order of actions in a video without any time information. The sequences of actions are also called transcripts. One of the first approaches [3] for this setting used discriminative clustering. Other approaches addressed the segmentation task by aligning the network output to the transcript [4, 11, 26] or by generating pseudo ground truth using the Viterbi algorithm [2, 12, 18, 34]. At inference time, most of these approaches iterate over the transcripts seen during training and select the one with the highest alignment score to generate the final prediction. While such approaches generate good segmentation results, this comes at the cost of slow inference time. To alleviate this problem, [37] combined a TCN-based backbone with a sequence-to-sequence model that predicts the transcript and segment lengths at inference time. To supervise the segment length output, an additional loss that enforces consistency between the sequence-to-sequence predictions and the TCN output is used. Recently, a gradient descent based approximation for the inference stage of the action segmentation has been proposed [36]. Another form of supervision that has recently emerged is based on video tags or action sets [9, 24, 25, 33]. In contrast to transcripts, set supervision provides only the set of actions that occur in the videos without any information regarding the order or how many times each action occurs. However, the performance of these approaches is still inferior compared to approaches that use transcripts for supervision.

**Timestamp Supervision for Action Segmentation.** Timestamp supervision has been recently used for other activity understanding tasks. [30] proposed to sample frames around the annotated timestamps to train an action classifier based on a sampling function that is fitted to the classifier response. In the context of action localization, [29] proposed a mining strategy to sample action frames and background frames for training. In [22], this ap-

proach is extended in two ways. First, it uses a greedy approach to find for each action class the frames belonging to the action using sampled action frames and background frames as seeds. Second, it uses a contrastive loss assuming that there are at least two instances of the same action in a video. In the context of action segmentation, [28] has used timestamps for training and reported competitive results compared to fully supervised learning. For timestamp supervision, [28] proposed to annotate only a single frame for each action instance. To train an action segmentation model with this level of supervision, the model output and the timestamps are used to generate a labeling of an entire video, which is then used to update the model. While this approach has been effective, it assumes that all action instances are annotated by timestamps. By contrast, we relax this assumption in this paper and address the case where some actions have been missed.

## 3 Robust Action Segmentation from Timestamps

### 3.1 Timestamp Supervision for Action Segmentation

Given an untrimmed video, action segmentation is the task of predicting the action label for each frame. The video is represented as a set of  $D$ -dimensional pre-extracted framewise features  $X = (x_1, \dots, x_T)$  with length  $T$  and a network for action segmentation predicts the framewise probability estimates  $\tilde{Y} = (\tilde{y}_1, \dots, \tilde{y}_T)$ , where  $\tilde{y}_t \in [0, 1]^C$  represents the predicted probability for each action class and  $C$  is the set of classes. We denote the ground truth labels for the video by  $Y = (y_1, \dots, y_T)$ .

While in case of full supervision the labels  $Y$  are available for all frames during training, [28, 30] addressed the training setup where only one frame per action is annotated. The annotated frames are called timestamps, which we denote by  $P = (p_1, \dots, p_N)$  and their corresponding labels by  $y_{p_1}, \dots, y_{p_N}$ . We assume that the timestamps are already ordered, *i.e.*,

$$1 \leq p_1 \leq \dots \leq p_N \leq T.$$

Timestamps are very sparse annotations of the training videos. For instance, less than 0.5% of the frames are annotated for the 50Salads [38] and Breakfast [15] datasets in case of timestamp supervision.

While a naive approach for action segmentation just uses the timestamps for training the network, Li *et al.* [28] proposed to generate a labeling of all frames from the timestamps. This is done by finding action boundaries between two consecutive timestamps in a forward and backward pass over the video, where each pass minimizes the variance of the framewise features that are separated by the action boundary. The approach, however, assumes that there is no other action between two consecutive timestamps. In other words, the approach assumes that the annotators did not miss any action during the labeling process. If an action is missed, the labeling generation process of [28] assigns the frames either to the previous or next timestamp, which is incorrect as it is illustrated in Figure 1. As a consequence, the approach is not very robust to missed annotations and the accuracy drastically decreases if only 80-90% of the actions are annotated as shown in Figure 2. While a recall of 90% is realistic for larger datasets, 70% is already very low given that 100% annotated timestamps cover already less than 0.5% of the frames. In this work, we therefore address the research question of how can we make the learning of networks for action segmentation more robust to missing timestamps.

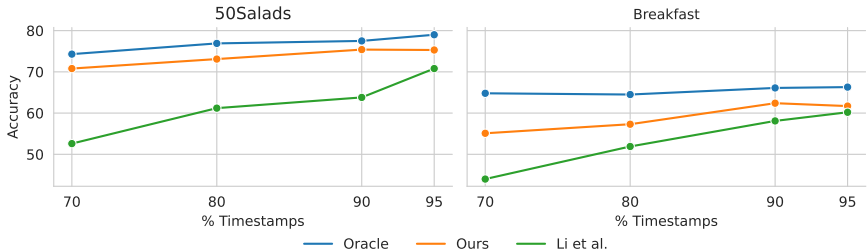


Figure 2: Accuracy of different action segmentation approaches trained using timestamp supervision with different percentages of annotated timestamps. Note that 100% timestamps corresponds to one timestamp per action segment, which already covers less than 0.5% of all frames. The two plots show the framewise accuracy for the 50Salads [58] and Breakfast dataset [45], respectively. The ‘Oracle’ is an upper bound of our approach and it corresponds to training in a fully supervised setting, but it ignores all frames of a ground-truth segment without timestamp. Our approach is more robust to missed timestamps than Li *et al.* [28] and consistently achieves a higher accuracy. The gap between our approach and the approach of [28] increases as the percentage of missed timestamps increases.




Figure 3: Given the annotated timestamps, we aim to identify the start and end of the actions that correspond to the timestamps. This is done by a joint optimization over all timestamps in a training video where we divide the frames between two consecutive timestamps  $p_1$  and  $p_2$  into three parts. The first part  $r_1$  will take the same label as  $p_1$  and the last part  $l_2$  will take the same label as  $p_2$ . The gray colored frames  $g_1$  will be labeled as unknown and ignored when training the network.

### 3.2 Robust Action Segmentation from Timestamps

In order to make the labeling process, and thus the training of the network, more robust, we need to consider that some timestamps might have been missed during the annotation process. This can be achieved by labeling some frames as unknown as shown in Figure 1. Instead of detecting an action change between two consecutive timestamps as in [28], we thus propose to expand the segment boundaries from the timestamps in each direction such that the network is confident that the frames within a segment belong to the same class as the timestamp and such that the number of frames that are unknown is minimized. The latter is important since there is otherwise no incentive to extend the segment boundaries.

Figure 3 shows the notation and illustrates the outcome of the labeling process for two timestamps. For each timestamp  $p_i$ , the corresponding action segment is defined by  $[p_i - l_i, p_i + r_i]$ . Between two timestamps  $p_i$  and  $p_{i+1}$ , there can be  $g_i$  frames that are labeled as unknown, *i.e.*,  $p_i + r_i + g_i + l_{i+1} = p_{i+1}$ . We also allow the frames at the beginning and at the end of a video to be labeled as unknown, which we denote by  $g_0$  and  $g_N$ , respectively.

As [28], we first train a network for action segmentation for a few epochs using only the annotated timestamps. During training, the network provides framewise probability estimates  $\tilde{Y} = (\tilde{y}_1, \dots, \tilde{y}_T)$  for all frames. We then optimize the parameters  $l_i$ ,  $r_i$ , and  $g_i$  such that on the one hand the predictions for the class label of the timestamp  $y_{p_i}$  are highly confident within  $[p_i - l_i, p_i + r_i]$  (colored regions in Figure 3) and such that on the other hand the number of unknown frames  $g_i$  (gray regions in Figure 3) is as small as possible. This goal

can be formulated for each video as the following constrained minimization problem:

$$\begin{aligned} \{r_i^*, g_i^*, l_i^*\} = \operatorname{argmin}_{r_i, g_i, l_i} & \sum_{i=1}^N \left( \sum_{t=p_i-l_i}^{p_i+r_i} -\log \tilde{y}_t[y_{p_i}] \right) + \beta \sum_{i=0}^N g_i \\ \text{s.t.} & p_{i+1} - p_i = r_i + g_i + l_{i+1} \\ & r_i \geq 0, \quad g_i \geq 0, \quad l_i \geq 0. \end{aligned} \quad (1)$$

The first constraint ensures that all frames are either labeled by the class of the associated timestamp or as unknown, and the last three constraints ensure that the number of frames for each part is not negative. Due to the constraints, the segments cannot overlap and the number of unknown frames is correctly counted. The hyper-parameter  $\beta$  controls how strong the number of unknown frames should be penalized. We study the effect of this hyper-parameter in Section 4.3.

In order to optimize (1) efficiently, we can re-write the objective function as

$$\begin{aligned} & \sum_{i=1}^N \left( \sum_{t=p_i-l_i}^{p_i+r_i} -\log \tilde{y}_t[y_{p_i}] \right) + \beta \sum_{i=0}^N g_i = \\ & \sum_{i=1}^N \left( \sum_{t=1}^T -\log \tilde{y}_t[y_{p_i}] \mathcal{I}(t|p_i-l_i \leq t \leq p_i+r_i) \right) \\ & + \beta \sum_{t=1}^T \left( 1 - \sum_{i=1}^N \mathcal{I}(t|p_i-l_i \leq t \leq p_i+r_i) \right) \end{aligned} \quad (2)$$

where  $\mathcal{I}(t|p_i-l_i \leq t \leq p_i+r_i)$  is the indicator function with value 1 if  $t$  is within the left and right bounds of the timestamp and 0 otherwise. As the  $\mathcal{I}$  function is a non-differentiable function, we approximate it with a plateau function [30, 36]. Further details are given in the supplementary document.

To address the constraints of (1), we re-parameterize  $r_i$ ,  $g_i$ , and  $l_{i+1}$  by  $r'_i$ ,  $g'_i$ , and  $l'_{i+1}$ , respectively, such that

$$z = \frac{(p_{i+1} - p_i) \exp(z')}{\exp(r'_i) + \exp(g'_i) + \exp(l'_{i+1})} \quad \text{for } z \in \{r_i, g_i, l_{i+1}\}. \quad (3)$$

In this way, we can optimize (2) for any value of  $r'_i$ ,  $g'_i$ , and  $l'_{i+1}$  without constraints since the re-parameterization ensures that  $r_i \geq 0$ ,  $g_i \geq 0$ ,  $l_i \geq 0$ , and  $p_{i+1} - p_i = r_i + g_i + l_{i+1}$ .

With the optimization problem becoming unconstrained and differentiable using the plateau function approximation, we can solve this optimization using gradient descent similar to FIFA [36]. This gradient-based optimization is solved only during training for each video independently. After the optimization, we use the generated labeling for each training video to continue the training of the network for action segmentation where the frames that are labeled as unknown are ignored. More details are provided in the supplementary material.

## 4 Experiments

### 4.1 Datasets and Metrics

**Datasets.** We evaluate our approach for action segmentation on the 50Salads [33] and Breakfast [15] datasets. Furthermore, we provide additional results for the BEOID [6] and Georgia

Tech Egocentric Activities (GTEA) [8] datasets for action localization.

The 50Salads dataset consists of 50 videos of people preparing mixed salads. The average length of the videos is around 6.4 minutes. The dataset provides fine-granular annotations of 17 low-level activities such as *peel cucumber* and *mix ingredients*. For evaluation, we use five-fold cross-validation and report the average.

The Breakfast dataset contains 1712 videos of breakfast-related activities. The average length of the videos is around 2.3 minutes. The frames are annotated with 48 fine-grained action classes such as *take bowl* and *pour cereals*. For evaluation, we use the standard four splits and report the average.

For both datasets, we use the timestamp annotations provided by [28] and simulate missing annotations by randomly removing timestamps. Otherwise, the setup is the same as [28]. We use the same network for action segmentation and train it for 30 epochs using only the annotated timestamps. We then generate the labels for the training videos as described in Section 3.2 and continue to train the network for 20 epochs on the labeled sequences. For frames that are labeled as unknown, the training loss is not computed, *i.e.*, the frames are ignored for training the network. If not otherwise mentioned, we set the value of  $\beta$  in (2) to 0.7. More details regarding the optimization of (2) are provided in the supplemental material. As in previous works for action segmentation, we report frame-wise accuracy (Acc), a score based on the segmental edit distance (Edit), and segmental F1 scores at overlapping thresholds 10%, 25% and 50%. For all metrics, a higher value is better.

## 4.2 Impact of Missing Annotations

We first compare our approach with [28] for different percentages of annotated timestamps. While we evaluate our approach for a setting where we use all provided timestamps, *i.e.*, 100% of the timestamps, in Section 4.4, we focus in this section on a setting where some actions have been missed during the annotation process. The results for the 50Salads and Breakfast datasets are reported in Table 1. Our approach is robust to missing annotations and outperforms the approach of [28] for all percentages and both datasets. The frame-wise accuracy of our approach drops by only 4.5% on 50Salads and by 6.6% on the Breakfast dataset when the percentage of annotated timestamps is reduced from 95% to 70%. On the contrary, the accuracy of [28] drops considerably by 18.2% on 50Salads and by 16.2% on the Breakfast dataset. The frame-wise accuracy is also plotted in Figure 2, which shows that the gap between our approach and [28] increases as more timestamps are missing. Using 70% of the timestamps for training, our approach outperforms [28] by 15.2% and 11.5% for the F1@25 metric on 50Salads and Breakfast, respectively.

These results demonstrate that our proposed approach is much more robust to missing timestamp annotations than [28]. While the approach [28] generates dense labels, which results in wrong labels in the case of many missing timestamps, our approach explicitly handles missing annotations by ignoring some frames during training. This effect is also visible in the qualitative result shown in Figure 4. While the approach of [28] assigns wrong labels to the frames of actions that have been missed during the annotation process, our approach labels these frames as unknown.

Since the number of labeled actions decreases as more timestamps are missing, we also report the results for an oracle in Table 1. This oracle defines an upper bound that can be achieved by our method if  $l_i$  and  $r_i$  (2) are perfectly estimated for each annotated timestamp. To this end, we use the frame-wise ground-truth annotations of the training videos to get  $l_i$  and  $r_i$ . Note that frames that belong to an action that has not been annotated by a timestamp

% Timestamps	Method	Breakfast					50Salads				
		F1@{10, 25, 50}			Edit	Acc	F1@{10, 25, 50}			Edit	Acc
95%	Li <i>et al.</i> [28]	67.3	59.7	42.7	68.2	60.2	70.9	67.4	53.4	63.8	70.8
	Ours	<b>70.2</b>	<b>62.4</b>	<b>44.8</b>	<b>69.7</b>	<b>61.7</b>	<b>72.9</b>	<b>69.6</b>	<b>57.5</b>	<b>64.2</b>	<b>75.3</b>
	Oracle	71.0	65.3	51.4	70.2	66.3	74.8	72.0	64.1	67.9	79.0
90%	Li <i>et al.</i> [28]	65.0	56.5	39.7	66.8	58.1	63.9	59.6	44.3	57.6	63.8
	Ours	<b>69.8</b>	<b>62.2</b>	<b>44.5</b>	<b>69.7</b>	<b>62.4</b>	<b>70.0</b>	<b>65.1</b>	<b>55.2</b>	<b>62.1</b>	<b>75.4</b>
	Oracle	69.0	63.0	49.1	68.6	66.1	73.9	71.6	62.5	66.9	77.5
80%	Li <i>et al.</i> [28]	59.8	50.4	33.0	62.8	51.9	62.7	56.9	40.3	54.2	61.2
	Ours	<b>67.3</b>	<b>58.9</b>	<b>40.8</b>	<b>68.5</b>	<b>57.3</b>	<b>70.9</b>	<b>67.8</b>	<b>53.7</b>	<b>61.4</b>	<b>73.1</b>
	Oracle	69.2	62.8	48.7	68.4	64.5	73.3	70.2	61.0	65.8	76.9
70%	Li <i>et al.</i> [28]	53.9	43.5	26.8	59.2	44.0	50.2	44.0	29.5	44.7	52.6
	Ours	<b>65.0</b>	<b>55.0</b>	<b>35.7</b>	<b>66.5</b>	<b>55.1</b>	<b>64.1</b>	<b>59.2</b>	<b>44.8</b>	<b>56.9</b>	<b>70.8</b>
	Oracle	69.5	63.2	49.0	69.0	64.8	67.4	63.3	53.2	59.0	74.3

Table 1: Impact of missing timestamps on the Breakfast and 50Salads datasets.




Figure 4: Qualitative results for comparing the generated labels by our approach and [28] on one video from the 50Salads dataset. The top row shows the ground truth frame-wise labels and the annotated timestamps as black lines. Note that only the 17 timestamps are given as annotation for training. The missed actions are highlighted with red dashed ellipses. The labels that are generated by our approach from these 17 timestamps are shown in the second row, where the frames labeled as unknown are colored in gray. The labeling generated by [28] is shown in the third row. While our approach correctly labels the frames of the missed actions correctly as unknown, [28] labels these frames wrongly.

are still ignored by the oracle. Therefore, the accuracy of the oracle also decreases as the percentage of annotated timestamps decreases. As shown in Figure 2, the gap between our approach and its upper bound is quite constant for the 50Salads dataset while it increases slightly on Breakfast as the number of annotated timestamps decreases.

### 4.3 Sensitivity to $\beta$

In (2), the number of frames that will be labeled as unknown is controlled by  $\beta$ . This is a hyper-parameter for our approach and has been set so far to 0.7 in all experiments. In this section, we study the impact of  $\beta$  on the accuracy, and Figure 5 shows the frame-wise accuracy of our approach using different values for  $\beta$ . The results are reported for different percentages of annotated timestamps on the 50Salads dataset. We plot the mean and standard deviation of three runs for each setting. As shown in the figure, lower values for  $\beta$  tend to achieve better results if annotations are missing. The accuracy drops for large values of  $\beta$ , except for the setting where all timestamps are available (100%). This is expected since with a lower percentage of annotations more frames should be labeled as unknown, which is achieved by lowering the value of  $\beta$ . When there are no missing timestamps, no frame should be labeled as unknown, and higher values of  $\beta$  are better.






Figure 5: Impact of  $\beta$  on the frame-wise accuracy for different percentages of annotated timestamps on the 50Salads dataset. For each experiment, the mean and standard deviation of three runs are shown.

#### 4.4 Comparison without Missing Annotations

So far, we mainly considered the case where some actions have been missed since the goal of the proposed approach is to increase the robustness with respect to missing timestamp annotations. We finally evaluate our approach for the protocol used in [28] where for each action a timestamp is provided. Table 2 shows the results on both the 50Salads and the Breakfast dataset. While our approach is designed to handle missing annotations, it also works well when there are no missing timestamps. Our approach outperforms previous approaches on 50Salads and achieves competitive results on the Breakfast dataset. As discussed in Section 4.3, if we know that there are no missing timestamps, then we can directly use a higher value of  $\beta$  (for instance  $\beta = 2$ ). By increasing  $\beta$ , we encourage our approach to decrease the number of frames that are labeled as unknown, which improves the accuracy. For completeness, we also compare to fully supervised approaches and approaches that are trained with less supervision.

Supervision Method	Breakfast				50Salads				Supervision Method	Breakfast 50Salads	
	F1@ [10, 25, 50]	Edit	Acc		F1@ [10, 25, 50]	Edit	Acc			Acc	
Full	MS-TCN [16]	52.6	48.1	37.9	61.7	66.3	76.3	74.0	64.5	67.9	80.7
	MS-TCN++ [16]	64.1	58.6	45.9	65.6	67.6	80.7	78.5	70.1	74.3	83.7
	BCN [16]	68.7	65.5	55.0	66.2	70.4	82.3	81.3	74.0	74.3	84.4
	ASRF [16]	74.3	68.9	56.1	72.4	67.6	84.9	83.5	77.3	79.3	84.5
	FIFA [16]	75.5	70.2	54.8	78.5	68.6	-	-	-	-	-
	MuCon [16]	73.2	66.1	48.4	76.3	62.8	-	-	-	-	-
	ASFormer [16]	76.0	70.6	57.4	75.0	73.5	85.1	83.4	76.0	79.6	85.6
Timestamps	Plateau [16]*	65.5	59.1	43.2	65.9	63.5	71.2	68.2	56.1	62.6	73.9
	Li et al. [16]	70.5	63.6	47.4	69.9	64.1	73.9	70.9	60.1	66.8	75.6
	Ours	67.0	60.0	43.8	66.7	60.8	75.2	72.1	61.5	67.7	78.4
	Ours ( $\beta = 2$ )	71.5	64.3	47.3	70.9	62.9	77.0	74.2	62.2	69.8	79.3
Transcripts	FIFA [16]	-	-	-	-	-	51.3	-	-	-	-
	CDFL [16]	-	-	-	-	-	50.2	-	54.7	-	-
	HMM-RNN [16]	-	-	-	-	-	-	-	45.5	-	-
	MuCon [16]	-	-	-	-	-	48.5	-	-	-	-
	D <sup>3</sup> TW [16]	-	-	-	-	-	45.7	-	-	-	-
Sets	NN-Viterbi [16]	-	-	-	-	-	43.0	-	49.4	-	-
	TCFPN [16]	-	-	-	-	-	38.4	-	-	-	-
	SCT [16]	-	-	-	-	-	30.4	-	-	-	-
	SCV [16]	-	-	-	-	-	30.2	-	-	-	
	Action Sets [16]	-	-	-	-	-	23.3	-	-	-	

Table 2: Results for timestamp annotations (100%) on the Breakfast and 50Salads datasets. For completeness, we also compare to fully supervised approaches and approaches that are trained with less supervision (transcripts or action sets). \* are results reported in [28].

#### 4.5 Action Localization

While our approach is designed for action segmentation, it can also be applied to the action localization task. We report the results on the BEOID [16] and the Georgia Tech Egocentric Activities (GTEA) [8] dataset. Following [28], we use our label generation approach to train an action localization model using the human-annotated timestamps on the GTEA and BEOID datasets provided by [29]. Results for both datasets are shown in Table 3. Our approach outperforms [28] and [29] by a large margin and achieves competitive results com-

pared to [24], which is a specific approach for action localization.

mAP@IoU	0.1	0.3	0.5	0.7	Avg
<b>GTEA</b>					
SF-Net [24]	58.0	37.9	19.3	11.9	31.0
Li <i>et al.</i> [28]	60.2	44.7	28.8	12.2	36.4
Lee <i>et al.</i> [22]	<b>63.9</b>	<b>55.7</b>	<u>33.9</u>	<b>20.8</b>	<b>43.5</b>
Ours	<u>63.6</u>	<u>54.1</u>	<b>36.4</b>	<u>20.3</u>	<u>43.4</u>
<b>BEOID</b>					
SF-Net [24]	62.9	40.6	16.7	3.5	30.1
Li <i>et al.</i> [28]	71.5	40.3	20.3	5.5	34.4
Lee <i>et al.</i> [22]	<u>76.9</u>	<u>61.4</u>	<b>42.7</b>	<b>25.1</b>	<u>51.8</u>
Ours	<b>79.0</b>	<b>68.3</b>	<u>42.1</u>	<u>17.2</u>	<b>52.9</b>

Table 3: Comparison with other approaches for action localization with timestamp supervision on the GTEA and BEOID dataset.

## 5 Conclusion

In this paper, we proposed a robust action segmentation approach for timestamp supervision. In contrast to the previous method [28] that assumes that all actions are annotated by a timestamp, we proposed a label generation method that handles missing timestamp annotations and is thus more suitable for real-world applications. While our approach achieves competitive results when all actions are annotated, it is much more robust to missing annotations. Our approach achieves considerably higher accuracy, edit score, and F1 score compared to [28] when actions have been missed during the annotation process.

**Acknowledgement** The work has been supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) GA 1927/4-2 (FOR 2535 Anticipating Human Behavior), MKW NRW iBehave, and the ERC Consolidator Grant FORHUE (101044724).

## References

- [1] Yazan Abu Farha and Juergen Gall. MS-TCN: Multi-stage temporal convolutional network for action segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3575–3584, 2019.
- [2] Hyemin Ahn and Dongheui Lee. Refining action segmentation with hierarchical video representations. In *IEEE International Conference on Computer Vision (ICCV)*, pages 16302–16310, 2021.
- [3] Piotr Bojanowski, Rémi Lajugie, Francis Bach, Ivan Laptev, Jean Ponce, Cordelia Schmid, and Josef Sivic. Weakly supervised action labeling in videos under ordering constraints. In *European Conference on Computer Vision (ECCV)*, pages 628–643. Springer, 2014.

- [4] Chien-Yi Chang, De-An Huang, Yanan Sui, Li Fei-Fei, and Juan Carlos Niebles. D<sup>3</sup>TW: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3546–3555, 2019.
- [5] Rui Dai, Srijan Das, Saurav Sharma, Luca Minciullo, Lorenzo Garattoni, Francois Bremond, and Gianpiero Francesca. Toyota smarhome untrimmed: Real-world untrimmed videos for activity detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2022.
- [6] Dima Damen, Teesid Leelasawassuk, Osian Haines, Andrew Calway, and Walterio Mayol-Cuevas. You-do, I-learn: Discovering task relevant objects and their modes of interaction from multi-user egocentric video. In *British Machine Vision Conference (BMVC)*, 2014.
- [7] Li Ding and Chenliang Xu. Weakly-supervised action segmentation with iterative soft boundary assignment. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6508–6516, 2018.
- [8] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3281–3288, 2011.
- [9] Mohsen Fayyaz and Jurgen Gall. SCT: Set constrained temporal transformer for set supervised action segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [10] Shang-Hua Gao, Qi Han, Zhong-Yu Li, Pai Peng, Liang Wang, and Ming-Ming Cheng. Global2local: Efficient structure search for video action segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [11] De-An Huang, Li Fei-Fei, and Juan Carlos Niebles. Connectionist temporal modeling for weakly supervised action labeling. In *European Conference on Computer Vision (ECCV)*, pages 137–153. Springer, 2016.
- [12] Yifei Huang, Yusuke Sugano, and Yoichi Sato. Improving action segmentation via graph-based temporal reasoning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14024–14034, 2020.
- [13] Yuchi Ishikawa, Seito Kasai, Yoshimitsu Aoki, and Hirokatsu Kataoka. Alleviating over-segmentation errors by detecting action boundaries. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2322–2331, 2021.
- [14] Svebor Karaman, Lorenzo Seidenari, and Alberto Del Bimbo. Fast saliency based pooling of fisher encoded dense trajectories. In *In European Conference on Computer Vision (ECCV) THUMOS Workshop*, 2014.
- [15] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 780–787, 2014.

- [16] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-to-end generative framework for video segmentation and recognition. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016.
- [17] Hilde Kuehne, Alexander Richard, and Juergen Gall. Weakly supervised learning of actions from transcripts. *Computer Vision and Image Understanding*, 163:78–89, 2017.
- [18] Hilde Kuehne, Alexander Richard, and Juergen Gall. A Hybrid RNN-HMM approach for weakly supervised temporal action segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 42(04):765–779, 2020.
- [19] Sateesh Kumar, Sanjay Haresh, Awais Ahmed, Andrey Konin, M Zeeshan Zia, and Quoc-Huy Tran. Unsupervised activity segmentation by joint representation learning and online clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [20] Colin Lea, Austin Reiter, René Vidal, and Gregory D Hager. Segmental spatiotemporal CNNs for fine-grained action segmentation. In *European Conference on Computer Vision (ECCV)*, pages 36–52. Springer, 2016.
- [21] Colin Lea, Michael D. Flynn, Rene Vidal, Austin Reiter, and Gregory D. Hager. Temporal convolutional networks for action segmentation and detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [22] Pilhyeon Lee and Hyeran Byun. Learning action completeness from points for weakly-supervised temporal action localization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 13648–13657, 2021.
- [23] Peng Lei and Sinisa Todorovic. Temporal deformable residual networks for action segmentation in videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6742–6751, 2018.
- [24] Jun Li and Sinisa Todorovic. Set-constrained viterbi for set-supervised action segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10820–10829, 2020.
- [25] Jun Li and Sinisa Todorovic. Anchor-constrained viterbi for set-supervised action segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9806–9815, 2021.
- [26] Jun Li, Peng Lei, and Sinisa Todorovic. Weakly supervised energy-based learning for action segmentation. In *IEEE International Conference on Computer Vision (ICCV)*, pages 6243–6251, 2019.
- [27] Shijie Li, Yazan Abu Farha, Yun Liu, Ming-Ming Cheng, and Juergen Gall. MS-TCN++: Multi-stage temporal convolutional network for action segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.
- [28] Zhe Li, Yazan Abu Farha, and Juergen Gall. Temporal action segmentation from timestamp supervision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8365–8374, 2021.

- [29] Fan Ma, Linchao Zhu, Yi Yang, Shengxin Zha, Gourab Kundu, Matt Feiszli, and Zheng Shou. SF-Net: Single-frame supervision for temporal action localization. In *European Conference on Computer Vision (ECCV)*, 2020.
- [30] Davide Moltisanti, Sanja Fidler, and Dima Damen. Action recognition from single timestamp supervision in untrimmed videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9915–9924, 2019.
- [31] Hamed Pirsiavash and Deva Ramanan. Parsing videos of actions with segmental grammars. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 612–619, 2014.
- [32] Alexander Richard, Hilde Kuehne, and Juergen Gall. Weakly supervised action learning with RNN based fine-to-coarse modeling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 754–763, 2017.
- [33] Alexander Richard, Hilde Kuehne, and Juergen Gall. Action sets: Weakly supervised action segmentation without ordering constraints. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5987–5996, 2018.
- [34] Alexander Richard, Hilde Kuehne, Ahsan Iqbal, and Juergen Gall. NeuralNetwork-Viterbi: A framework for weakly supervised video learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7386–7395, 2018.
- [35] Marcus Rohrbach, Sikandar Amin, Mykhaylo Andriluka, and Bernt Schiele. A database for fine grained activity detection of cooking activities. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [36] Yaser Souri, Yazan Abu Farha, Fabien Despinoy, Gianpiero Francesca, and Juergen Gall. FIFA: Fast Inference Approximation for Action Segmentation. In *German Conference on Pattern Recognition (GCPR)*, 2021.
- [37] Yaser Souri, Mohsen Fayyaz, Luca Minciullo, Gianpiero Francesca, and Juergen Gall. Fast Weakly Supervised Action Segmentation Using Mutual Consistency. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2021.
- [38] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 729–738, 2013.
- [39] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. Coin: A large-scale dataset for comprehensive instructional video analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [40] Nam N Vo and Aaron F Bobick. From stochastic grammar to bayes network: Probabilistic parsing of complex activity. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2641–2648, 2014.
- [41] Zhenzhi Wang, Ziteng Gao, Limin Wang, Zhifeng Li, and Gangshan Wu. Boundary-aware cascade networks for temporal action segmentation. In *European Conference on Computer Vision (ECCV)*, 2020.

- [42] Fangqiu Yi, Hongyu Wen, and Tingting Jiang. Asformer: Transformer for action segmentation. In *British Machine Vision Conference (BMVC)*, 2021.
- [43] Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. Cross-task weakly supervised learning from instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.