Temporal Action Segmentation from Timestamp Supervision

Zhe Li, Yazan Abu Farha, Juergen Gall University of Bonn, Germany

s6zeliii@uni-bonn.de, {abufarha,gall}@iai.uni-bonn.de

Abstract

Temporal action segmentation approaches have been very successful recently. However, annotating videos with frame-wise labels to train such models is very expensive and time consuming. While weakly supervised methods trained using only ordered action lists require less annotation effort, the performance is still worse than fully supervised approaches. In this paper, we propose to use timestamp supervision for the temporal action segmentation task. Timestamps require a comparable annotation effort to weakly supervised approaches, and yet provide a more supervisory signal. To demonstrate the effectiveness of timestamp supervision, we propose an approach to train a segmentation model using only timestamps annotations. Our approach uses the model output and the annotated timestamps to generate frame-wise labels by detecting the action changes. We further introduce a confidence loss that forces the predicted probabilities to monotonically decrease as the distance to the timestamps increases. This ensures that all and not only the most distinctive frames of an action are learned during training. The evaluation on four datasets shows that models trained with timestamps annotations achieve comparable performance to the fully supervised approaches.

1. Introduction

Analyzing and understanding video content is very important for many applications, such as surveillance or intelligent advertisement. Recently, several approaches have been very successful in analyzing and segmenting activities in videos [20, 24, 1, 29, 41]. Despite the success of the previous approaches, they rely on fully annotated videos where the start and end frames of each action are annotated.

This level of supervision, however, is very time consuming and hard to obtain. Furthermore, as the boundaries between action segments are usually ambiguous, this might result in inconsistencies between annotations obtained from different annotators. To alleviate these problems, many researchers start exploring weaker levels of supervision in the form of transcripts [3, 35, 25] or even sets [34, 11, 26]. For



Figure 1. For fully supervised action segmentation, each frame in the training videos is annotated with an action label (top). This process is time-consuming since it requires an accurate annotation of the start and end frame of each action. To reduce the annotation effort, we propose to use timestamps as supervision (bottom). In this case, only one arbitrary frame needs to be annotated for each action and the annotators do not need to search for the start and end frames, which is the most time-consuming annotation part.

transcript-level supervision, the videos are annotated with an ordered list of actions occurring in the video without the starting and ending time of each action. Whereas for the set-level supervision, only the set of actions are provided without any information regarding the order or how many times each action occurs in the videos.

While transcript-level and set-level supervision significantly reduce the annotation effort, the performance is not satisfying and there is still a gap compared to fully supervised approaches. In this paper, inspired by the recently introduced timestamp supervision for action recognition [31], we propose to use timestamp supervision for the action segmentation task to address the limitations of the current weakly supervised approaches. For timestamp supervision, only one frame is annotated from each action segment as illustrated in Fig. 1. Such timestamps annotations can be obtained with comparable effort to transcripts, and yet it provides more supervision. Besides the ordered list of actions occurring in the video, timestamps annotations give partial information about the location of the action segments, which can be utilized to further improve the performance.

Given the timestamps annotations, the question is how to train a segmentation model with such level of supervision. A naive approach takes only the sparsely annotated frames for training. This, however, ignores most of the information in the video and does not achieve good results as we will show in the experiments. Another strategy is to iterate the process and consider frames with high confidence scores near the annotations as additional annotated frames and include them during training [31]. Furthermore, frames that are far away from the annotations can be considered as negative samples [28]. For temporal action segmentation, which is comparable to semantic image segmentation, however, all frames need to be annotated and there are no large parts of the video that can be used to sample negative examples. Furthermore, relying only on frames with high confidence discards many of the video frames that occur during an action and focuses only on the most distinctive frames of an action, which can be sufficient for action recognition or detection but not for action segmentation.

In this work, we therefore propose a different approach where all frames of the videos are used. Instead of detecting frames of high confidences, we aim to identify changes of actions in order to divide the videos into segments. Since for each action change the frames before the change should be assigned to the previous timestamp and after the change to the next timestamp, we find the action changes by minimizing the variations of the features within each of the two clusters of frames. While we can then train the model on all frames by assigning the label of the timestamp to the corresponding frames, it does not guarantee that all frames of an action are effectively used. We therefore introduce a loss function that enforces a monotonic decrease in the class probabilities as the distance to the timestamps increases. This loss encourages the model to predict higher probabilities for low confident regions that are surrounded by high confident frames and therefore to use all frames and not only the most distinctive frames.

Our contribution is thus three folded.

- We propose to use timestamp supervision for the temporal action segmentation task, where the goal is to predict frame-wise action labels for untrimmed videos.
- 2. We introduce an approach to train a temporal action segmentation model from timestamp supervision. The approach uses the model predictions and the annotated timestamps for estimating action changes.
- 3. We propose a novel confidence loss that forces the model confidence to decrease monotonically as the distance to the timestamp increases.

We evaluate our approach on four datasets: 50Salads [38], Breakfast [16], BEOID [7], and Georgia Tech Egocentric Activities (GTEA) [10]. We show that training an action segmentation model is feasible with only timestamp supervision without compromising the performance compared to the fully supervised approaches. On the 50Salads dataset, for instance, we achieve 97% of the accuracy compared to fully supervised learning, but at a tiny fraction of the annotation costs. ¹

2. Related Work

We briefly discuss the related work for the temporal action segmentation task at different levels of supervision.

Fully Supervised Action Segmentation. Temporal action segmentation has received an increasing interest recently. In contrast to action recognition where the goal is to classify trimmed videos [36, 4, 12], temporal action segmentation requires capturing long-range dependencies to classify each frame in the input video. To achieve this goal, many approaches combined frame-wise classifiers with grammars [40, 32] or with hidden Markov models (HMMs) [21, 17, 19]. Despite the success of these approaches, their performance was limited and they were slow both at training and inference time. Recent approaches, therefore, utilized temporal convolutional networks to capture long-range dependencies for the temporal action segmentation task [20, 24]. While such approaches managed to generate accurate predictions, they suffer from an over-segmentation problem. To alleviate this problem, current state-of-the-art methods follow a multi-stage architecture with dilated temporal convolutions [1, 41, 27, 15, 14]. These approaches rely on fully annotated datasets that are expensive to obtain. On the contrary, we address the temporal action segmentation task in a weakly supervised setup.

Weakly Supervised Action Segmentation. Weakly supervised action segmentation has been an active research area recently. Earlier approaches apply discriminative clustering to detect actions using movie scripts [3, 9]. However, these approaches ignored the action ordering information and only focused on distinguishing action segments from background, which is a common practice for the temporal action localization task [22, 23]. Bojanowski et al. [3] extended these ideas to segment actions in videos using transcripts in the form of an ordered list of actions as supervision. Recently, many researchers addressed the task by aligning the video frames and the transcripts using connectionist temporal classification [13], dynamic time warping [5], or energy-based learning [25]. Other approaches generate pseudo ground truth labels for the training videos and iteratively refine them [18, 33, 8, 19]. In [35], a framewise loss function is combined with the Viterbi algorithm

¹The source code for our model and the timestamps annotations are publicly available at https://github.com/ZheLi2020/ TimestampActionSeg.

to generate the target labels. While these approaches have been very successful, they suffer from a slow inference time as they iterate over all the training transcripts and select the one with the highest score. Souri et al. [37] addressed this issue by predicting the transcript besides the frame-wise scores at inference time. While these approaches rely on a cheap transcript supervision, their performance is much worse than fully supervised approaches. In contrast to these approaches, we propose a higher level of supervision in the form of timestamps that can be obtained with comparable effort to the transcript supervision, and yet reduces the gap to the fully supervised approaches. There is another line of research addressing the action segmentation task from setlevel supervision [34, 11, 26]. These approaches opt for a weaker level of supervision at the cost of performance. In contrast to these approaches, we propose a good compromise between supervision level and performance.

Timestamp Supervision for Recognizing Activities. Timestamp supervision has not yet received much attention from the action recognition community. Initial attempts were inspired by the success of point supervision for semantic segmentation [2]. Mettes et al. [30] apply multiple instance learning for spatio-temporal action localization using points annotation on a sparse subset of frames. Chéron et al. [6] use discriminative clustering to integrate different types of supervision for the spatio-temporal action localization task. Recently, Moltisanti et al. [31] proposed a sampling distribution based on a plateau function centered around temporal timestamps annotations to train a fine-grained action classifier. This approach relies on the classifier response to sample frames around the annotated timestamps and uses them for training. The method was tested for classifying trimmed videos and also showed promising results for temporal action localization. Ma et al. [28] extended the action localization setup by mining action frames and background frames for training.

3. Temporal Action Segmentation

Temporal action segmentation is the task of predicting frame-wise action labels for a given input video. Formally, given a sequence of video frames $X = [x_1, \ldots, x_T]$, where T is the number of frames, the goal is to predict a sequence of frame-wise action labels $[a_1, \ldots, a_T]$. In contrast to the fully supervised approaches, which assume that the frame-wise labels are given at training time, we consider a weaker level of supervision in the form of timestamps. In Section 3.1 we introduce the timestamp supervision for the temporal action segmentation task. Then, we describe the proposed framework for learning from timestamp supervision in Section 3.2. Finally, we provide the details of the loss function in Section 3.3.



Figure 2. The framework of the proposed approach for training with timestamp supervision. Given the output of the segmentation model and the timestamps annotations, we generate action labels for each frame in the input video by estimating where the action labels change. A loss function is then computed between the predictions and the generated labels.

3.1. Timestamp Supervision

In a fully supervised setup, the frame-wise labels $[a_1, \ldots, a_T]$ of the training videos are available. On the contrary, for timestamp supervision, only a single frame for each action segment is annotated. Given a training video X with T frames and N action segments, where $N \ll T$, only N frames are annotated with labels $A_{TS} = [a_{t_1}, \ldots, a_{t_N}]$, where frame t_i belongs to the *i*-th action segment. To annotate timestamps, one can go fast forward through a video and press a button when an action occurs. This does not take more time than annotating transcripts. Whereas annotating the start and end frames of each action requires going slowly back and forth between the frames. As reported in [28], annotators need 6 times longer to annotate the start and end frame compared to annotating a single timestamp. While timestamps are much easier to obtain compared to the full annotation of the video frames, they provide much more information compared to weaker forms of supervision such as transcripts. Fig. 1 illustrates the difference between timestamp supervision and full supervision.

3.2. Action Segmentation from Timestamp Supervision

Given an action segmentation model \mathcal{M} and a set of training videos with timestamps annotations, the goal is to train the model \mathcal{M} to predict action labels for each frame in the input video. If the frame-wise labels are available during the training, as in the fully supervised case, then it is possible to apply a classification loss on the output of the model \mathcal{M} for each frame in the input video. However, in timestamp supervision, only a sparse set of frames are annotated. To alleviate this problem, we propose to generate frame-wise labels for the training videos, and use them as a target for the loss function as illustrated in Fig. 2.

Detecting Action Changes. Given the timestamps annotations $A_{TS} = [a_{t_1}, \ldots, a_{t_N}]$ for a video X, we want to generate frame-wise action labels $A = [\hat{a}_1, \dots, \hat{a}_T]$ for each frame in that video such that $\hat{a}_{t_i} = a_{t_i}$ for $i \in [1, N]$. As for each action segment there is an annotated frame, finding the frame-wise labels can be reduced to finding the action change between each consecutive annotated timestamps. To this end, we pass the input video X to the segmentation model \mathcal{M} , which will be described in Section 4.2, and use the output of the penultimate layer H combined with the timestamps annotations to estimate where the action labels change between the timestamps. To generate the labels, all the frames that lie between an annotated timestamp and an estimated time of action change are assigned with the same action label as the annotated timestamp as illustrated in Fig. 3. To detect the action change between two timestamps t_i and t_{i+1} , we find the time t_{b_i} that minimizes the following stamp-to-stamp energy function

$$t_{b_i} = \arg\min_{\hat{t}} \sum_{t=t_i}^{\hat{t}} d(h_t, c_i) + \sum_{t=\hat{t}+1}^{t_{i+1}} d(h_t, c_{i+1}),$$

s.t.

$$c_{i} = \frac{1}{\hat{t} - t_{i} + 1} \sum_{t=t_{i}}^{\hat{t}} h_{t},$$

$$c_{i+1} = \frac{1}{t_{i+1} - \hat{t}} \sum_{t=\hat{t}+1}^{t_{i+1}} h_{t},$$

$$t_{i} \le \hat{t} < t_{i+1},$$
(1)

where d(.,.) is the Euclidean distance, h_t is the output of the penultimate layer at time t, c_i is the average of the output between the first timestamp t_i and the estimate \hat{t} , and c_{i+1} is the average of the output between the estimate \hat{t} and the second timestamp t_{i+1} . *I.e.*, we find the time t_{b_i} that divides the frames between two timestamps into two clusters with the minimum distance between frames and the corresponding cluster center.

Forward-Backward Action Change Detection. In (1), the stamp-to-stamp energy function considers only the frames between the annotated timestamps to estimate where the actions change. Nonetheless, if we already have an estimate for $t_{b_{i-1}}$, then we already know that frames between the estimate $t_{b_{i-1}}$ and the next timestamp t_i will be assigned to action label a_{t_i} . This information can be used to estimate the time of action change for the next action segment t_{b_i} . The same argument also holds if we start estimating the boundaries in reverse order. *I.e.*, if we already know $t_{b_{i+1}}$, then frames between t_{i+1} and $t_{b_{i+1}}$ can be used to estimate t_{b_i} . We call the former estimate a forward estimate for the *i-th* action change, whereas the later is called the backward



Figure 3. Given the timestamps annotations, we first estimate where the actions change between consecutive timestamps. To generate the frame-wise labels, all the frames that lie between an annotated timestamp and an estimated time of action change are assigned with the same action label as the annotated timestamp.

estimate. The final estimate for t_{b_i} is the average of these two estimates. Formally

$$t_{b_i} = \frac{t_{b_i,FW} + t_{b_i,BW}}{2}$$

s.t.

$$t_{b_{i},FW} = \arg\min_{\hat{t}} \sum_{t=t_{b_{i-1}}}^{\hat{t}} d(h_{t},c_{i}) + \sum_{t=\hat{t}+1}^{t_{i+1}} d(h_{t},c_{i+1}),$$

$$t_{b_{i},BW} = \arg\min_{\hat{t}} \sum_{t=t_{i}}^{\hat{t}} d(h_{t},c_{i}) + \sum_{t=\hat{t}+1}^{t_{b_{i+1}}} d(h_{t},c_{i+1}).$$

(2)

3.3. Loss Function

Recent fully supervised approaches for action segmentation use a combination of a classification loss and a smoothing loss [1, 15, 41]. Besides these losses, we further introduce a novel confidence loss for the timestamp supervision. In the following, we describe in detail each loss function.

Classification Loss. We use a cross entropy loss between the predicted action probabilities and the corresponding generated target label

$$\mathcal{L}_{cls} = \frac{1}{T} \sum_{t} -log(\tilde{y}_{t,\hat{a}}), \tag{3}$$

where $\tilde{y}_{t,\hat{a}}$ is the predicted probability for the target label \hat{a} at time t.

Smoothing Loss. As the classification loss treats each frame independently, it might result in an undesired oversegmentation effect. To encourage a smooth transition between frames and reduce over-segmentation errors, we use the truncated mean squared error [1] as a smoothing loss

$$\mathcal{L}_{T-MSE} = \frac{1}{TC} \sum_{t,a} \tilde{\Delta}_{t,a}^2, \qquad (4)$$

$$\tilde{\Delta}_{t,a} = \begin{cases} \Delta_{t,a} & : \Delta_{t,a} \le \tau \\ \tau & : otherwise \end{cases},$$
(5)

$$\Delta_{t,a} = |\log \tilde{y}_{t,a} - \log \tilde{y}_{t-1,a}|, \tag{6}$$



Figure 4. The confidence loss penalizes increases in the model confidence for label a_{t_i} as we move away from the annotated timestamp t_i .

where T is the video length, C is the number of action classes, and $\tilde{y}_{t,a}$ is the probability of action a at time t.

Confidence Loss. Our approach relies on the model output to detect action changes. Nonetheless, as some frames are more informative than others, the model confidence might alternate between high and low values within the same action segment. Such behavior might result in ignoring regions with low confidence within the segments. To alleviate this problem, we apply the following loss

$$\mathcal{L}_{conf} = \frac{1}{T'} \sum_{a_{t_i} \in A_{TS}} \left(\sum_{t=t_{i-1}}^{t_{i+1}} \delta_{a_{t_i}, t} \right), \tag{7}$$

 $\delta_{a_{t_i},t} = \begin{cases} \max(0, \log \tilde{y}_{t,a_{t_i}} - \log \tilde{y}_{t-1,a_{t_i}}) & if \ t \ge t_i \\ \max(0, \ \log \tilde{y}_{t-1,a_{t_i}} - \log \tilde{y}_{t,a_{t_i}}) & if \ t < t_i \end{cases},$ (8)

where $\tilde{y}_{t,a_{t_i}}$ is the probability of action a_{t_i} at time t, and $T' = 2(t_N - t_1)$ is the number of frames that contributed to the loss. For the first and last timestamps, we set $t_0 = t_1$ and $t_{N+1} = t_N$. This loss penalizes an increase in confidence as we deviate from the annotations as illustrated in Fig. 4.

Enforcing monotonicity on the model confidence has two effects as shown in Fig. 6. First, it encourages the model to predict higher probabilities for low confident regions that are surrounded by regions with high confidence. Second, it suppresses outlier frames with high confidence that are far from the timestamps and not supported by high confident regions.

The final loss function to train the segmentation model is the sum of these three losses

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \alpha \mathcal{L}_{T-MSE} + \beta \mathcal{L}_{conf}, \qquad (9)$$

where α and β are hyper-parameters to balance the contribution of each loss.

4. Experiments

4.1. Datasets and Metrics

Datasets. We evaluate our approach on four datasets: 50Salads [38], Breakfast [16], BEOID [7], and Georgia Tech Egocentric Activities (GTEA) [10].

The **50Salads** dataset contains 50 videos with roughly 0.6M frames, where the frames are annotated with 17 action

classes. The videos show actors preparing different kind of salads. We use five-fold cross validation for evaluation and report the average.

The **Breakfast** dataset contains 1712 videos with roughly 3.6M frames, where the frames are annotated with 48 action classes. All actions are related to breakfast preparation activities. We use the standard four splits for evaluation and report the average.

The **BEOID** dataset contains 58 videos, where the frames are annotated with 34 actions classes. For evaluation, we use the same training-testing split as in [28].

The **GTEA** dataset contains 28 videos with roughly 32K frames, where the frames are annotated with 11 action classes. For evaluation, we report the average of four splits.

To generate the timestamps annotations, we randomly select one frame from each action segment in the training videos. We further evaluate our approach using human and noisy annotations in Section 4.7. Additional settings are evaluated in the supplementary material.

Metrics. We use the standard metrics for fully supervised action segmentation and report frame-wise accuracy (Acc), segmental edit distance (Edit) and segmental F1 scores at overlapping thresholds 10%, 25% and 50%.

Baselines. We implement two baselines: a Naive and a Uniform baseline. The **Naive** baseline computes the loss at the annotated timestamps only and does not generate framewise labels. Whereas the **Uniform** baseline generates the frame-wise labels by assuming that action labels change at the center frame between consecutive timestamps.

4.2. Implementation Details

We use a multi-stage temporal convolutional network [1] as a segmentation model \mathcal{M} . Following [39], we use two parallel stages for the first stage with kernel size 5 and 3 and pass the sum of the outputs to next stages. We train our model for 50 epochs with Adam optimizer. To minimize the impact of initialization, only the annotated timestamps are used for the classification loss in the first 30 epochs, and the generated labels are used afterwards. The learning rate is set to 0.0005 and the batch size is 8. For the loss function, we set $\tau = 4$, $\alpha = 0.15$ as in [1] and set $\beta = 0.075$. As input for our model, we use the same I3D [4] features that were used in [1].

4.3. Comparison with the Baselines

In this section, we compare the proposed approach for action segmentation from timestamp supervision with the naive and uniform baselines. The results on the three datasets are shown in Table 1. Our approach outperforms these baselines with a large margin in all the evaluation metrics. While the naive baseline achieves a good frame-wise accuracy, it suffers from a severe over-segmentation problem as indicated by the low F1 and Edit scores. This is



Figure 5. Qualitative results on (a) 50Salads, (b) Breakfast, and (c) GTEA datasets. As the naive baseline only trains on the sparse annotations, it suffers from an over-segmentation problem. While the uniform baseline reduces this problem by uniformly assigning labels to the frames, the durations of the predicted segments are not accurate and the predictions tend towards a uniform segmentation of the videos. On the contrary, our approach generates better predictions by utilizing the model output to detect where the action labels change.

because it only uses the sparse timestamps annotations for training, which leaves a lot of ambiguity for frames without annotations. Using the uniform baseline reduces the oversegmentation by uniformly assigning a label for each frame. However, this results in inferior frame-wise accuracy as the uniform assignment generates many wrong labels. On the contrary, our approach utilizes the model predictions to generate much better target labels, which is reflected in the performance as illustrated in Fig. 5. We also compare the performance of our approach to the fully supervised setup in Table 1. Our approach achieves comparable performance to the fully supervised case.

4.4. Impact of the Loss Function

The loss function to train our model consists of three losses: a classification loss, a smoothing loss, and a confidence loss. Table 2 shows the impact of each loss on both the 50Salads and the Breakfast dataset. While either of the smoothing loss and the confidence loss gives an additional boost in performance, the best performance is achieved when both of the losses are combined with the

	F1	@{10,2	25, 50}	Edit	Acc
50Salads					
Naive	47.9	43.3	34.0	37.2	69.6
Uniform	62.9	58.2	42.3	60.4	63.4
Ours	73.9	70.9	60.1	66.8	75.6
Full Supervision	70.8	67.7	58.6	63.8	77.8
Breakfast					
Naive	34.1	29.1	20.1	37.4	56.8
Uniform	66.2	56.3	36.4	68.1	51.0
Ours	70.5	63.6	47.4	69.9	64.1
Full Supervision	69.9	64.2	51.5	69.4	68.0
GTEA					
Naive	59.7	55.3	39.6	51.1	56.5
Uniform	78.9	72.5	50.9	73.1	56.5
Ours	78.9	73.0	55.4	72.3	66.4
Full Supervision	85.1	82.7	69.6	79.6	76.1

Table 1. Comparison with the baselines on the three datasets.

classification loss with a frame-wise accuracy improvement of 2.8% and 3.9% on 50Salads and the Breakfast dataset respectively, and roughly 10% on the F1 score at 50% overlapping threshold.

While the smoothing loss forces a smooth transition between consecutive frames, it does not take the annotations

	F1@	{10, 25	, 50}	Edit	Acc
50Salads					
\mathcal{L}_{cls}	65.7	62.6	50.7	57.7	72.8
$\mathcal{L}_{cls} + \alpha \mathcal{L}_{T-MSE}$	70.1	66.8	55.3	62.6	74.6
$\mathcal{L}_{cls} + \beta \mathcal{L}_{conf}$	73.2	70.6	60.1	65.2	75.3
$\mathcal{L}_{cls} + \alpha \mathcal{L}_{T-MSE} + \beta \mathcal{L}_{conf}$	73.9	70.9	60.1	66.8	75.6
Breakfast					
\mathcal{L}_{cls}	60.3	52.8	36.7	64.2	60.2
$\mathcal{L}_{cls} + \alpha \mathcal{L}_{T-MSE}$	67.5	60.1	44.3	68.9	63.7
$\mathcal{L}_{cls} + \beta \mathcal{L}_{conf}$	67.6	60.4	43.7	68.0	61.6
$\mathcal{L}_{cls} + \alpha \mathcal{L}_{T-MSE} + \beta \mathcal{L}_{conf}$	70.5	63.6	47.4	69.9	64.1

Table 2. Contribution of the different loss functions on the 50Salads and Breakfast datasets.

			>		
	F1	@{10, 2	5, 50}	Edit	Acc
$\beta = 0$	70.1	66.8	55.3	62.6	74.6
$\beta = 0.025$	70.9	68.8	57.4	63.4	76.2
$\beta = 0.05$	73.1	70.2	58.7	65.4	75.6
$\beta = 0.075$	73.9	70.9	60.1	66.8	75.6
$\beta = 0.1$	73.2	70.6	60.1	66.1	74.6

Table 3. Impact of β on the 50Salads dataset.

into account. On the contrary, the confidence loss forces the predicted probabilities to monotonically decrease as the distance to the timestamps increases. This encourages the model to have a high confidence for all frames within an action segment, and yet it suppresses outlier frames that are far from the annotations and not supported by regions with high confidence as illustrated in Fig. 6.

To balance the contribution of the different losses, we set the weight of the smoothing loss to 0.15 as in [1], and the weight of the confidence loss $\beta = 0.075$. In Table 3, we study the impact of β on the performance on the 50Salads dataset. As shown in the table, good results are achieved for β between 0.05 and 0.1.

4.5. Impact of the Energy Function for Action Change Detection

Our approach generates target labels by estimating where the action labels change using the forward-backward estimate as in (2). To analyze the impact of this estimate, we train another model that directly uses the stamp-tostamp estimate (Stamp-to-Stamp (Features)) as in (1). As shown in Table 4, our approach performs better. We also tried another variant of the stamp-to-stamp energy function that maximizes the average probabilities of the action segments (Stamp-to-Stamp (Prob.)) instead of minimizing the distances to cluster centers. However, the performance is worse than the proposed energy function.

4.6. Impact of the Segmentation Model \mathcal{M}

In all experiments, we used a multi-stage temporal convolutional architecture based on [1] and [39]. In this section we study the impact of the segmentation model on the performance. To this end, we apply the proposed training scheme on the original MS-TCN [1] and the recently introduced MS-TCN++ [27]. As shown in Table 5, our approach



Figure 6. Impact of the confidence loss. Forcing monotonicity encourages the model to have a high confidence for all frames within an action segment (a). It also suppresses outlier frames with high confidence (b).

	$F1@{10, 25, 50}$			Edit	Acc
50Salads					
Stamp-to-Stamp (Prob.)	67.5	61.8	48.6	61.1	68.9
Stamp-to-Stamp (Features)	73.4	70.5	59.9	66.7	74.2
Ours	73.9	70.9	60.1	66.8	75.6
Breakfast					
Stamp-to-Stamp (Prob.)	65.7	55.9	35.9	68.0	58.8
Stamp-to-Stamp (Features)	66.3	59.6	44.4	67.9	60.1
Ours	70.5	63.6	47.4	69.9	64.1

Table 4. Impact of the energy function for action change detection on the 50Salads and Breakfast datasets.

Dataset	Seg. Model \mathcal{M}	F1@	{10, 25	, 50}	Edit	Acc
50Salads	MS-TCN [1]	71.7	68.7	57.0	64.0	74.7
	MS-TCN++ [27]	75.0	71.1	55.8	67.2	72.9
	Ours	73.9	70.9	60.1	66.8	75.6
GTEA	MS-TCN [1]	79.8	73.3	47.7	76.3	59.3
	MS-TCN++ [27]	78.3	72.2	49.1	74.5	62.2
	Ours	78.9	73.0	55.4	72.3	66.4

Table 5. Impact of the segmentation model \mathcal{M} on the 50Salads and GTEA datasets.

is agnostic to the segmentation model and performs well with all these models.

4.7. Comparison with the State-of-the-Art

In this section, we compare our approach with recent state-of-the-art approaches for timestamp supervision. To the best of our knowledge, timestamp supervision has not been studied for the temporal action segmentation task. We, therefore, compare with similar methods in the context of action recognition [31] and action localization [28].

Since the approach of [31] assumes the testing videos are trimmed and does not work for long untrimmed videos, we replaced their backbone network with our segmentation model for a fair comparison. To this end, we initialized the plateau functions around the timestamps annotations of the training videos and iteratively update their parameters based on the segmentation model output as in [31]. Results for the 50Salads, Breakfast, and GTEA datasets are shown in

Supervision	Method	F1@	{10, 25	, 50}	Edit	Acc
Full	MS-TCN [1]	76.3	74.0	64.5	67.9	80.7
	MS-TCN++ [27]	80.7	78.5	70.1	74.3	83.7
	BCN [41]	82.3	81.3	74.0	74.3	84.4
	ASRF [15]	84.9	83.5	77.3	79.3	84.5
Timestamps	Seg. Model \mathcal{M} + plateau [31] Ours	71.2 73.9	68.2 70.9	56.1 60.1	62.6 66.8	73.9 75.6
Transcripts	CDFL [25]	-	-	-	-	54.7
	NN-Viterbi [35]	-	-	-	-	49.4
	HMM-RNN [33]	-	-	-	-	45.5

Table 6. Comparison with different levels of supervision on the 50Salads dataset.

Method	F1@	{10, 25	, 50}	Edit	Acc
MS-TCN [1]	52.6	48.1	37.9	61.7	66.3
MS-TCN++ [27]	64.1	58.6	45.9	65.6	67.6
BCN [41]	68.7	65.5	55.0	66.2	70.4
ASRF [15]	74.3	68.9	56.1	72.4	67.6
Seg. Model \mathcal{M} + plateau [31]	65.5	59.1	43.2	65.9	63.5
Ours	70.5	63.6	47.4	69.9	64.1
CDFL [25]	-	-	-	-	50.2
MuCon [37]	-	-	-	-	47.1
D ³ TW [5]	-	-	-	-	45.7
NN-Viterbi [35]	-	-	-	-	43.0
TCFPN [8]	-	-	-	-	38.4
HMM-RNN [33]	-	-	-	-	33.3
ECTC [13]	-	-	-	-	27.7
SCT [11]	-	-	-	-	30.4
SCV [26]	-	-	-	-	30.2
Action Sets [34]	-	-	-	-	23.3
	Method MS-TCN [1] MS-TCN++ [27] BCN [41] ASRF [15] Seg. Model <i>M</i> + plateau [31] Ours CDFL [25] MuCon [37] D ³ TW [5] NN-Viterbi [35] TCFPN [8] HMM-RNN [33] ECTC [13] SCT [11] SCV [26] Action Sets [34]	Method F1@ MS-TCN [1] 52.6 MS-TCN++ [27] 64.1 BCN [41] 68.7 ASRF [15] 74.3 Seg. Model M + plateau [31] 65.5 Ours 70.5 CDFL [25] - MuCon [37] - D ³ TW [5] - NN-Viterbi [35] - TCFPN [8] - HMM-RNN [33] - ECTC [13] - SCV [26] - Action Sets [34] -	Method F1@{10, 25} MS-TCN [1] 52.6 48.1 MS-TCN++ [27] 64.1 58.6 BCN [41] 68.7 65.5 ASRF [15] 74.3 68.9 Seg. Model \mathcal{M} + plateau [31] 65.5 59.1 Ours 70.5 63.6 CDFL [25] - - MuCon [37] - - D ³ TW [5] - - NN-Viterbi [35] - - HMM-RNN [33] - - ECTC [13] - - SCT [11] - - SCV [26] - - Action Sets [34] - -	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$\begin{array}{c c c c c c c c c c c c c c c c c c c $

Table 7. Comparison with different levels of supervision on the Breakfast dataset.

Supervision	Method	F1@	{10, 25	, 50}	Edit	Acc
Full	MS-TCN [1]	85.8	83.4	69.8	79.0	76.3
	MS-TCN++ [27]	88.8	85.7	76.0	83.5	80.1
	BCN [41]	88.5	87.1	77.3	84.4	79.8
	ASRF [15]	89.4	87.8	79.8	83.7	77.3
Timestamps	Seg. Model \mathcal{M} + plateau [31] Ours	74.8 78.9	68.0 73.0	43.6 55.4	72.3 72.3	52.9 66.4

Table 8. Comparison with different levels of supervision on the GTEA dataset.

Tables 6-8, respectively. Our approach outperforms [31] on all datasets with a large margin of up to 13.5% frame-wise accuracy and 11.8% for the F1 score with 50% overlapping threshold on the GTEA dataset.

We also compare our approach for timestamp supervision with other levels of supervision for the temporal action segmentation task. As shown in Tables 6-8, timestamp supervision outperforms weaker levels of supervision in the form of transcripts or sets with a large margin. Our approach provides a good compromise between annotation effort and performance, and further reduces the gap to fully supervised approaches.

Timestamp supervision has recently been studied for action localization in [28]. In their approach, they use the model confidence to sample foreground action frames and background frames for training. To compare with [28], we use the same setup and the provided human annotations to train our model and report mean average precision at different overlapping thresholds. Table 9 shows the results on the GTEA and BEOID [7] datasets. Our approach outperforms

mAP@IoU	0.1	0.3	0.5	0.7	Avg
GTEA					
SF-Net [28]	58.0	37.9	19.3	11.9	31.0
Ours	60.2	44.7	28.8	12.2	36.4
BEOID					
SF-Net [28]	62.9	40.6	16.7	3.5	30.1
Ours	71.5	40.3	20.3	5.5	34.4

Table 9. Comparison with SF-Net [28] for action localization with timestamp supervision on the GTEA and BEOID datasets.

Fraction	Method	Acc
0.1	HMM-RNN [19] Ours	60.9 68.4
0.01	HMM-RNN [19] Ours	58.8 67.4

Table 10. Comparison with Kuehne *et al.* [19] on the Breakfast dataset with semi-supervised setup.

Ma *et al.* [28] with a large margin of 5.4% average mAP on GTEA and 4.3% on the BEOID dataset. In contrast to [28] where only the frames with high confidence are used for training, our approach detects action changes and generates a target label for each frame in the training videos.

Finally, we compare our approach with the semisupervised setup on the Breakfast dataset proposed in Kuehne *et al.* [19]. In this setup, the training videos are annotated with the transcript of actions and a fraction of the frames as well. Compared to the timestamp supervision, this setup provides annotations for much more frames. Since the timestamps are randomly sampled from the video, there are sometimes multiple timestamps for one action and not all actions are annotated as reported in the supplementary material. As shown in Table 10, our approach outperforms [19] with a large margin. While the approach of [19] relies on an expensive Viterbi decoding during inference, our approach directly predicts the frame-wise labels.

5. Conclusion

In this paper, we proposed an approach to train a temporal action segmentation model using only timestamps annotations. Our approach combines the model predictions with the timestamps annotations for estimating where the action labels change. We further introduced a confidence loss that enforces monotonicity on the model confidence. The loss encourages high confidence values for all frames within an action segment and suppresses outlier frames. Results on four datasets show that models trained with timestamp supervision achieve comparable performance to the fully supervised setup. The proposed approach is model agnostic and can be applied to any segmentation model.

Acknowledgements: The work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – GA 1927/4-2 (FOR 2535 Anticipating Human Behavior) and the ERC Starting Grant ARCA (677650).

References

- 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. 1, 2, 4, 5, 7, 8
- [2] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What's the point: Semantic segmentation with point supervision. In *European Conference on Computer Vision* (ECCV), pages 549–565. Springer, 2016. 3
- [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. 1, 2
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6299–6308, 2017. 2, 5
- [5] Chien-Yi Chang, De-An Huang, Yanan Sui, Li Fei-Fei, and Juan Carlos Niebles. D³TW: 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. 2, 8
- [6] Guilhem Chéron, Jean-Baptiste Alayrac, Ivan Laptev, and Cordelia Schmid. A flexible model for training action localization with varying levels of supervision. In Advances in Neural Information Processing Systems (NeurIPS), pages 942–953, 2018. 3
- [7] 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. 2, 5, 8
- [8] 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. 2, 8
- [9] Olivier Duchenne, Ivan Laptev, Josef Sivic, Francis Bach, and Jean Ponce. Automatic annotation of human actions in video. In *IEEE International Conference on Computer Vi*sion (ICCV), pages 1491–1498, 2009. 2
- [10] 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. 2, 5
- [11] 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. 1, 3, 8
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [13] 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. 2, 8

- [14] 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. 2
- [15] 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. 2, 4, 8
- [16] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goaldirected human activities. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 780– 787, 2014. 2, 5
- [17] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-toend generative framework for video segmentation and recognition. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016. 2
- [18] Hilde Kuehne, Alexander Richard, and Juergen Gall. Weakly supervised learning of actions from transcripts. *Computer Vision and Image Understanding*, 163:78–89, 2017. 2
- [19] 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. 2, 8
- [20] 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. 1, 2
- [21] 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. 2
- [22] Pilhyeon Lee, Youngjung Uh, and Hyeran Byun. Background suppression network for weakly-supervised temporal action localization. In AAAI Conference on Artificial Intelligence, pages 11320–11327, 2020. 2
- [23] Pilhyeon Lee, Jinglu Wang, Yan Lu, and Hyeran Byun. Weakly-supervised temporal action localization by uncertainty modeling. In AAAI Conference on Artificial Intelligence, 2021. 2
- [24] 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. 1, 2
- [25] 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. 1, 2, 8
- [26] 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. 1, 3, 8

- [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. 2, 7, 8
- [28] 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. 2, 3, 5, 7, 8
- [29] Khoi-Nguyen C. Mac, Dhiraj Joshi, Raymond A. Yeh, Jinjun Xiong, Rogerio S. Feris, and Minh N. Do. Learning motion in feature space: Locally-consistent deformable convolution networks for fine-grained action detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 1
- [30] Pascal Mettes, Jan C Van Gemert, and Cees GM Snoek. Spot on: Action localization from pointly-supervised proposals. In *European Conference on Computer Vision (ECCV)*, pages 437–453. Springer, 2016. 3
- [31] 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. 1, 2, 3, 7, 8
- [32] 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. 2
- [33] 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. 2, 8
- [34] 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. 1, 3, 8
- [35] 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. 1, 2, 8
- [36] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In Advances in Neural Information Processing Systems (NeurIPS), pages 568–576, 2014. 2
- [37] Yaser Souri, Mohsen Fayyaz, Luca Minciullo, Gianpiero Francesca, and Juergen Gall. Fast weakly supervised action segmentation using mutual consistency. arXiv preprint arXiv:1904.03116, 2019. 3, 8
- [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. 2, 5
- [39] Kanav Vats, Mehrnaz Fani, Pascale Walters, David A Clausi, and John Zelek. Event detection in coarsely annotated sports videos via parallel multi-receptive field 1D convolutions. In IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 882–883, 2020. 5, 7

- [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. 2
- [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. 1, 2, 4, 8