

# Timestamp-Supervised Action Segmentation from the Perspective of Clustering

Dazhao Du<sup>1,2</sup>, Enhao Li<sup>2,3</sup>, Lingyu Si<sup>1,2</sup>, Fanjiang Xu<sup>1,\*</sup> and Fuchun Sun<sup>1,4</sup>

<sup>1</sup>Institute of Software, Chinese Academy of Science

<sup>2</sup>University of Chinese Academy of Sciences

<sup>3</sup>Institute of Computing Technology, Chinese Academy of Sciences

<sup>4</sup>Department of Computer Science and Technology, Tsinghua University

dudazhao20@mails.ucas.ac.cn, lienhan20g@ict.ac.cn, {lingyu,fanjiang}@iscas.ac.cn, fcsun@mail.tsinghua.edu.cn

## Abstract

Video action segmentation under timestamp supervision has recently received much attention due to lower annotation costs. Most existing methods generate pseudo-labels for all frames in each video to train the segmentation model. However, these methods suffer from incorrect pseudo-labels, especially for the semantically unclear frames in the transition region between two consecutive actions, which we call **ambiguous intervals**. To address this issue, we propose a novel framework from the perspective of clustering, which includes the following two parts. First, **pseudo-label ensembling** generates incomplete but high-quality pseudo-label sequences, where the frames in ambiguous intervals have no pseudo-labels. Second, **iterative clustering** iteratively propagates the pseudo-labels to the ambiguous intervals by clustering, and thus updates the pseudo-label sequences to train the model. We further introduce a clustering loss, which encourages the features of frames within the same action segment more compact. Extensive experiments show the effectiveness of our method.

## 1 Introduction

Recognizing and segmenting actions from long untrimmed videos is a challenging task. Action segmentation aims to classify actions for all frames in the video, resulting in several action segments. There have been many fully supervised approaches achieving good results. These approaches mainly use TCN [Farha and Gall, 2019; Li *et al.*, 2020] or Transformer [Du *et al.*, 2022] to model the long-term dependencies between frames. However, obtaining label sequences for long videos is costly. Researchers have proposed several weaker levels of supervision, e.g., set supervision [Richard *et al.*, 2018a], transcript supervision [Kuehne *et al.*, 2017], and timestamp supervision [Li *et al.*, 2021]. Among them, timestamp supervision has the most promising performance.

As shown in Figure 1, timestamp supervision only labels one frame in each action segment. Compared to other weakly supervised setups, timestamp supervision provides temporal

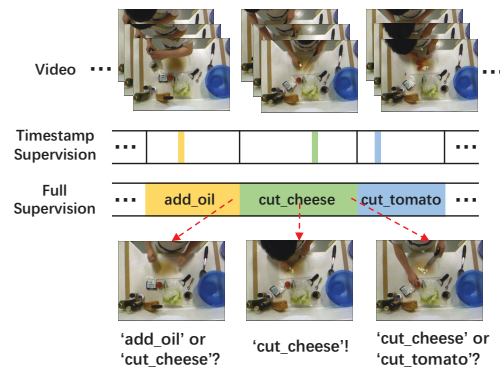


Figure 1: Only one frame in each action segment is labeled under timestamp supervision, while all frames are labeled under full supervision. Those frames near the end of action ‘add\_oil’ can be classified as either action ‘add\_oil’ or ‘cut\_cheese’. On the contrary, the frames near the middle of the action segment are usually representative of the action class.

information that is crucial for action segmentation and therefore has better performance. However, it is surprising that timestamped supervision sometimes outperforms the fully supervised setup. The reason is that the boundaries of action segments are usually ambiguous, which might result in inconsistencies between labels labeled by different annotators under full supervision. For example, the frames near the boundary between action ‘add\_oil’ and action ‘cut\_cheese’ in Figure 1, can be classified as either action ‘add\_oil’ or action ‘cut\_cheese’. We call the intervals close to the boundaries **ambiguous intervals**. In contrast, timestamp supervision usually labels the representative frame of an action, such as a frame near the middle of the action segment. The most popular method [Li *et al.*, 2021] iteratively generates full pseudo-label sequences by detecting action changes, as shown in Figure 2 (a). However, this might lead to error accumulations, i.e. training with wrong pseudo-labels of ambiguous intervals will cause the model to deviate further from expectations. Therefore, we do not expect to generate pseudo-labels for frames in ambiguous intervals at the beginning, as shown in Figure 2 (b). As the training progresses, we can gradually assign pseudo-labels with high certainty to ambiguous intervals based on the model outputs.

\*Corresponding author.




Figure 2: Two types of pseudo-label sequences for a video containing  $N$  action segments. (a) In the full pseudo-label sequence without ambiguous intervals, action segments are connected head-to-tail.  $\{b_1, b_2, \dots, b_{N+1}\}$  contains the boundaries of all the actions, including the first frame and last frame of the video. (b) There exist ambiguous intervals between action segments.  $l_n$  and  $r_n$  are the left and right boundaries of the  $n$ -th action segment.

We implement this idea from a clustering perspective by considering each action segment as a cluster. Our framework consists of two main parts, namely **pseudo-label ensembling (PLE)** and **iterative clustering (IC)**. First, PLE generates high-quality pseudo-label sequences with ambiguous intervals in Figure 2 (b). Specifically, we utilize three different clustering algorithms to generate three full pseudo-label sequences for each video. For each frame, if the three algorithms generate inconsistent pseudo-labels, the frame belongs to the ambiguous interval. Frames in the ambiguous intervals have no pseudo-labels and therefore do not affect training. Afterward, IC gradually assigns the frames in ambiguous intervals to left or right action segments by feature clustering based on their distance from the surrounding cluster centers. During the iterative training, the model can generate better features, based on which propagating pseudo-labels to the ambiguous interval results in more consistent and accurate supervision. As this process involves feature clustering, we further introduce the clustering loss commonly used in deep clustering to help feature learning and clustering.

In summary, our main contributions are twofold: (1) We design a framework for timestamp-supervised action from the perspective of clustering. The framework mainly contains two novel algorithms, **pseudo-label ensembling (PLE)** and **iterative clustering (IC)**, to address the negative impact of ambiguous intervals on training. Besides, we introduce the clustering loss. (2) Extensive experiments and ablation studies demonstrate the effectiveness of our framework and illustrate the impact of each component.

## 2 Related Work

**Fully Supervised Action Segmentation.** In the fully supervised setting, all the frames are labeled during training. To model the temporal dynamics, some researchers utilize graphical models and sequential models such as HMM [Tang *et al.*, 2012; Kuehne *et al.*, 2016], RNN [Singh *et al.*, 2016]. With the development of deep learning, temporal convolutional networks (TCNs) have been successful due to their powerful temporal modeling ability [Lea *et al.*, 2017; Farha and Gall, 2019; Li *et al.*, 2020; Wang *et al.*, 2020; Ishikawa *et al.*, 2021]. Recently, ASFormer [Yi *et al.*, 2021] and TUT [Du *et al.*, 2022] have designed transformer-based architectures that outperform TCN-based models. In contrast

to these methods, our framework is model agnostic and relies on timestamp supervision instead of full supervision.

**Weakly Supervised Action Segmentation.** Many works have attempted to explore weaker levels of supervision to reduce annotation costs. Transcript supervision provides ordered action lists. Some researchers introduce connectionist temporal classification [Huang *et al.*, 2016], dynamic time warping [Chang *et al.*, 2019], and energy-based learning [Li *et al.*, 2019] to align the frame-wise outputs to transcripts. Other approaches utilize the Viterbi algorithm on the output probabilities to generate pseudo-label sequences [Kuehne *et al.*, 2017; Richard *et al.*, 2018b]. Another form of supervision, set supervision [Richard *et al.*, 2018a; Fayyaz and Gall, 2020; Li and Todorovic, 2021], provides only the unordered action sets. However, these weakly supervised methods perform considerably worse than fully supervised approaches.

**Timestamp-Supervised Action Segmentation.** In action segmentation, Li *et al.* [2021] firstly proposed a timestamp-supervised framework that detected action changes by the energy function and generated full pseudo-label sequences to train the segmentation model. Based on it, Khan *et al.* [2022] used learnable GNN instead of the energy function to obtain action boundaries, achieving better performance. UVAST [Behrmann *et al.*, 2022] extended the K-medoids clustering algorithm to yield pseudo-labels for all frames. Rahaman *et al.* [2022] utilized the Expectation-Maximization algorithm to solve this problem. However, the loss function is still calculated on all frames to train the model during the M-step. Different from these methods, we iteratively generate pseudo-label sequences with ambiguous intervals, rather than generating pseudo-labels for all frames.

**Weakly Supervised Learning.** When faced with the challenge of training samples lacking labels, one common approach is to assign pseudo-labels to all unlabeled samples and use them to train the model [Lee and others, 2013]. However, we do not generate pseudo-labels for all unlabeled frames, considering the presence of ambiguous intervals. Iscen *et al.* [2019] constructed a similarity graph among samples based on features to propagate labels, while we propagate labels by iterative clustering. There are also some works learning with coarse-grained labels [Qiang *et al.*, 2023].

## 3 Methodology

### 3.1 Problem Definition

Given an input long video which is represented by a sequence of frame-wise visual features:  $X = [x_1, \dots, x_T]$ , where  $x_t$  is the visual feature of the  $t$ -th frame. Our goal is to predict the action class label of each frame  $[\hat{c}_1, \dots, \hat{c}_T]$ . Under timestamp supervision, only a single frame in each action segment is annotated. For a video having  $N$  action segments, timestamp supervision provides a set of annotated timestamps  $\{c_{\tau_1}, \dots, c_{\tau_N}\}$ , where the  $\tau_n$ -th frame is in the  $n$ -th segment.

### 3.2 Overview

As shown in Figure 3, we design a novel framework for timestamp-supervised action segmentation, which contains pseudo-label ensembling (PLE) and iterative clustering (IC).

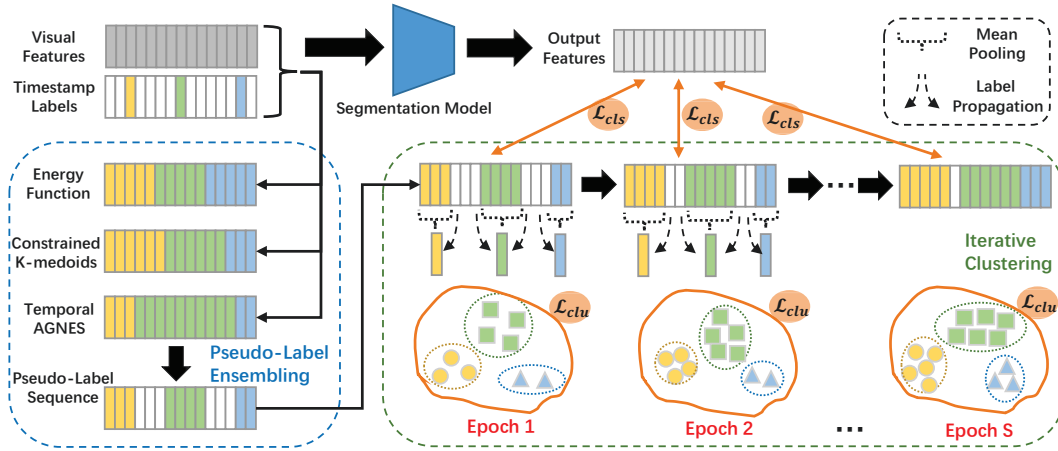


Figure 3: The overall framework of our method. Given frame-wise visual features and timestamp supervision, pseudo-label ensembling can generate pseudo-label sequences with ambiguous intervals. Subsequently, in each epoch of iterative clustering, the pseudo-labels are propagated into ambiguous intervals by clustering, and the resulting new pseudo-label sequences will train the segmentation model.  $\mathcal{L}_{cls}$  and  $\mathcal{L}_{ctu}$  are classification loss and clustering loss, respectively.

Given the visual features and timestamp labels of each video, PLE combines three different clustering algorithms to generate pseudo-label sequences, where the frames in ambiguous intervals have no pseudo-labels while other frames have. We regard PLE as a pre-processing step since it relies on only the visual features of frames. During the first training epoch in IC, we utilize the pseudo-label sequences generated by PLE to train the segmentation model. After training, we cluster the output features of the model. Specifically, we consider frames with the same pseudo-label to be in the same cluster, which means that an action segment is a cluster. The center vector of each cluster can be calculated based on the output features of all frames within the action segment. We classify the frames within ambiguous intervals according to the distances between their output features and surrounding cluster centers, which can propagate pseudo-labels and narrow down the ambiguous intervals. During the next training epoch, the model is trained by the updated pseudo-label sequences. We iterate this process for  $S$  epochs. The additional clustering loss gathers features of the same class. We describe PLE, IC, and the clustering loss in detail in the following subsections.

### 3.3 Pseudo-Label Ensembling

We find that action segmentation under timestamp supervision can be viewed as a special clustering problem. For a video containing  $N$  action segments, we can cluster the features of frames into  $N$  clusters. Since each timestamp will belong to a certain cluster, we label all frames in that cluster with the action label of the timestamp. In PLE, we ensemble three different clustering algorithms, including energy function [Li *et al.*, 2021], constrained k-medoids [Behrmann *et al.*, 2022], and temporal AGNES we propose.

**Energy Function.** Given frame-wise visual features and labeled timestamps in a video, the energy function aims to find action changes  $\{b_2, b_3, \dots, b_N\}$  by minimizing the variations of the features within each of two consecutive segments, as

shown in Figure 2 (a).

**Constrained K-medoids.** Similar to the energy function, the constrained k-medoids algorithm also aims to find the temporal boundaries of each segment such that the accumulative distance of each cluster to the current cluster center is minimized. Different from the typical k-medoids algorithm, it forces the clusters to be temporally continuous.

**Temporal AGNES.** AGNES (Agglomerative Nesting) [Han and Kamber, 2006] is a bottom-up hierarchical clustering algorithm: each frame starts in its own cluster, and clusters are successively merged together. We propose temporal AGNES, which only merges temporally adjacent clusters. The clustering algorithm stops when the number of clusters equals the number of timestamps. The detailed steps of the temporal AGNES algorithm are shown in Appendix A.

**PLE.** For each video sample, we can generate three different full pseudo-label sequences by the above algorithms. The frames within ambiguous intervals are semantically ambiguous, and different clustering algorithms might assign different action labels to these frames. Inspired by ensembling learning [Dietterich, 2002], we take the intersection of these three pseudo-label sequences. If the three algorithms assign the same pseudo-label to a frame, then that frame is labeled as this action class; otherwise, the frame belongs to the ambiguous interval and has no pseudo-label. So we can obtain a pseudo-label sequence with ambiguous intervals, as shown in Figure 2 (b). The pseudo-label sequence generated by PLE achieves a trade-off between label rate and accuracy. Compared to the timestamp labels, it has a higher label rate. Compared to the full pseudo-label sequence generated by any clustering algorithm, it has a higher accuracy. More analysis can be found in Figure 4 and Table 4.

### 3.4 Iterative Clustering

From the perspective of clustering, IC iteratively generates pseudo-labels for frames within ambiguous intervals to up-

---

**Algorithm 1** The iterative clustering algorithm
 

---

**Input:** visual features:  $[x_1, x_2, \dots, x_T]$ , pseudo-label sequence:  $[\hat{c}_1, \hat{c}_2, \dots, \hat{c}_T]$ , timestamp supervision:  $\{c_{\tau_1}, c_{\tau_2}, \dots, c_{\tau_N}\}$ , the number of epochs  $S$ , initial segmentation model:  $\mathcal{M}$

**Output:** segmentation model after iterative training:  $\mathcal{M}$

- 1:  $[\tilde{c}_1, \tilde{c}_2, \dots, \tilde{c}_T] \leftarrow [\hat{c}_1, \hat{c}_2, \dots, \hat{c}_T]$
- 2: Obtain the sets of left and right action segment boundaries from the pseudo-label sequence:  $\{l_1, l_2, \dots, l_N\}$  and  $\{r_1, r_2, \dots, r_N\}$
- 3: **for** epoch = 1, 2, ...,  $S$  **do**
- 4:  $[h_1, \dots, h_T] \leftarrow \mathcal{M}([x_1, \dots, x_T])$
- 5: **for**  $n = 1, 2, \dots, N$  **do**
- 6:  $m_n \leftarrow \frac{1}{r_n - l_n} \sum_{t=l_n}^{r_n} h_t$
- 7: **end for**
- 8: **for**  $n = 1, 2, \dots, N - 1$  **do**
- 9: **while**  $r_n < l_{n+1}$  and  $d(h_{r_n}, m_n) < d(h_{r_n}, m_{n+1})$  **do**
- 10:  $r_n \leftarrow r_n + 1$
- 11: **end while**
- 12: **while**  $r_n \leq l_{n+1}$  and  $d(h_{l_{n+1}}, m_n) > d(h_{l_{n+1}}, m_{n+1})$  **do**
- 13:  $l_{n+1} \leftarrow l_{n+1} - 1$
- 14: **end while**
- 15:  $l_{n+1} \leftarrow l_{n+1} + 1$
- 16: **end for**
- 17: **for**  $n = 1, 2, \dots, N$  **do**
- 18:  $[\tilde{c}_{l_n}, \tilde{c}_{l_n+1}, \dots, \tilde{c}_{r_n-1}] \leftarrow [c_{\tau_n}, c_{\tau_n}, \dots, c_{\tau_n}]$
- 19: **end for**
- 20: Train the model for one epoch with the pseudo-label sequence  $[\tilde{c}_1, \dots, \tilde{c}_T]$ , and obtain the new segmentation model  $\mathcal{M}$  after parameter update
- 21: **end for**

---

date the pseudo-label sequences for training the segmentation model  $\mathcal{M}$ . We clarify some notations before introducing IC in detail.  $x_t$  and  $h_t$  are the visual feature and output feature of the  $t$ -th frame, respectively. As shown in Figure 2,  $l_n$  and  $r_n$  are the left and right boundaries of the  $n$ -th action segment in the pseudo-label sequence.  $m_n$  is the mean vector of all the output features in the  $n$ -th action segment, which is equivalent to the cluster center vector of the  $n$ -th cluster. Function  $d(\cdot, \cdot)$  calculates the Euclidean distance between the output features of two frames.

The algorithm 1 summarizes the steps of our iterative clustering algorithm. Steps 4 to 20 describe the procedure for each epoch. First, the visual features of the video frames are fed into the segmentation model to generate output features (step 4). The center vector of each action cluster is calculated based on output features and the pseudo-label sequences from the previous epoch (steps 5-7). Afterward, for every ambiguous interval, we first iterate through each frame in the ambiguous interval from left to right and compare the distance between that frame and the center vectors of the adjacent left and right action clusters (steps 9-11). If it is closer to the center vector of the left cluster, the frame belongs to the left action segment. We stop the iteration when the frame is closer to the center vector of the right cluster, which means that the frame either belongs to the right action segment or remains in the ambiguous interval. This iterative process is pseudo-label propagation, which propagates the pseudo-labels from left to right to the frames in the ambiguous interval. Subsequently, we propagate the pseudo-labels from right to left (steps 12-

15), which is symmetric to the above process. After finishing pseudo-label propagation, the range of ambiguous intervals is reduced and the resulting supervision signal is more intensive. Finally, the segmentation model is trained by the new pseudo-label sequence (step 20).

### 3.5 Loss Function

During the training phase, we combine four different losses.

**Classification Loss.** We use a cross-entropy loss between the predicted action probabilities and the corresponding generated pseudo-labels:

$$\mathcal{L}_{cls} = \frac{1}{\sum_{n=1}^N (r_n - l_n)} \sum_{n=1}^N \left( \sum_{t=l_n}^{r_n-1} -\log y_t(c_{\tau_n}) \right), \quad (1)$$

where  $y_t(c_{\tau_n})$  is the predicted probability that  $x_t$  belongs to the  $c_{\tau_n}$ -th class, and the denominator of the former term,  $\sum_{n=1}^N (r_n - l_n)$ , represents the total number of frames with pseudo-labels that contribute to the loss. It is worth noting that the frames within the ambiguous intervals have no pseudo-labels and are not calculated for cross-entropy loss.

**Smoothing Loss.** We use the truncated mean squared error [Farha and Gall, 2019] as the smoothing loss to improve over-segmentation errors:

$$\mathcal{L}_{smo} = \frac{1}{TC} \sum_{t=1}^T \sum_{c=1}^C (\max(|\log y_t(c) - \log y_{t-1}(c)|, \theta))^2, \quad (2)$$

where  $y_t(c)$  is the predicted probability that  $x_t$  belongs to the  $c$ -th class, and  $\theta = 4$  is a pre-set threshold. Besides, gradients are not calculated w.r.t.  $y_{t-1}(c)$ .

**Confidence Loss.** The confidence loss [Li *et al.*, 2021] is specifically designed for action segmentation under timestamp supervision:

$$\mathcal{L}_{conf} = \frac{1}{2(\tau_N - \tau_1)} \sum_{n=1}^N \left( \sum_{t=\tau_{n-1}}^{\tau_n+1} \delta_{c_{\tau_n}, t} \right), \quad (3)$$

$$\delta_{c_{\tau_n}, t} = \begin{cases} \max(0, \log y_t(c_{\tau_n}) - \log y_{t-1}(c_{\tau_n})) & \text{if } t \geq \tau_n, \\ \max(0, \log y_{t-1}(c_{\tau_n}) - \log y_t(c_{\tau_n})) & \text{if } t < \tau_n, \end{cases} \quad (4)$$

where  $2(\tau_N - \tau_1)$  is the number of frames that contributed to the loss. The loss encourages the model to predict higher probabilities for low-confident intervals surrounded by high-confident intervals while suppressing those high-confident regions that are far from the labeled timestamp.

**Clustering Loss.** To assist feature learning and clustering, we introduce a clustering loss commonly used in deep clustering, to distribute output features in the same action segment more compactly. IC aims to cluster similar features together and thus propagates pseudo-labels to frames within ambiguous intervals. As IC is similar to K-means clustering, we introduce the clustering loss [Yang *et al.*, 2017] commonly used in deep clustering into action segmentation for the first time. It can be formulated as:

$$\mathcal{L}_{clu} = \frac{1}{\sum_{n=1}^N (r_n - l_n)} \sum_{n=1}^N \left( \sum_{t=l_n}^{r_n-1} \|h_t - m_n\|_2^2 \right), \quad (5)$$

where  $m_n = \frac{1}{r_n - l_n} \sum_{t=l_n}^{r_n-1} h_t$  is the center vector of the  $n$ -th cluster. The loss encourages the high-level features of frames in the same cluster to be more compact so that the model is more discriminative for frames in different action segments.

The overall loss function during the training phase is a weighted sum of the above four losses:

$$\mathcal{L}^s = \mathcal{L}_{cls} + \lambda \mathcal{L}_{smo} + \beta \mathcal{L}_{conf} + \gamma \mathcal{L}_{clu}, \quad (6)$$

where  $\lambda, \beta$  and  $\gamma$  are hyper-parameters.

## 4 Experiments

### 4.1 Datasets and Evaluation Metrics

We empirically evaluate the performance of the proposed framework on the following three public datasets. The **50Salads** dataset [Stein and McKenna, 2013] contains 50 top-view videos with 17 action classes. On average, each video lasts for about 6.4 minutes and has about 20 action instances. The **GTEA** dataset [Fathi *et al.*, 2011] consists of 28 ego-centric videos with 11 action classes. On average, there are about 20 action instances in each video. The **Breakfast** dataset [Kuehne *et al.*, 2014] contains 1,712 third-person view videos with 48 action classes. On average, each video is composed of about 6 action instances. We perform four-fold cross-validation on the GTEA and Breakfast datasets, and five-fold cross-validation on the 50Salads dataset. For evaluation, we report the average of all the splits.

We also use the segmental F1 score at overlapping thresholds 10%, 25%, 50% (F1@10,25,50), segmental edit distance (Edit) measuring the difference between predicted segment paths and ground-truth segment paths, and frame-wise accuracy (Acc) as evaluation metrics.

### 4.2 Implementation Details

On all datasets, we represent each video as a sequence of visual features. We employ the I3D [Carreira and Zisserman, 2017] features provided in [Farha and Gall, 2019] as inputs. We utilize the same segmentation model following [Li *et al.*, 2021], which has two parallel stages for the first stage with kernel sizes 3 and 5 and passes the sum of both outputs to the next stages. We train the model for 70 epochs with Adam optimizer [Kingma and Ba, 2014]. Similar to [Li *et al.*, 2021], in the first 50 epochs, we only use the annotated timestamps to train to find a good initialization. Subsequently, we apply the IC algorithm to train for 20 epochs. We use the learning rate 0.0001 for Breakfast and 0.0005 for GTEA and 50Salads. We set  $\lambda, \beta, \gamma$  to 0.15, 0.075 and 0.15 respectively.

### 4.3 Comparison with the State-of-the-art

We compare our method with recent state-of-the-art approaches, including four fully supervised methods (MS-TCN [Farha and Gall, 2019], MS-TCN++ [Li *et al.*, 2020], BCN [Wang *et al.*, 2020], ASRF [Ishikawa *et al.*, 2021]), two transcript-supervised methods (NN-Viterbi [Richard *et al.*, 2018b], CDFL [Li *et al.*, 2019]), two set-supervised methods (SCT [Fayyaz and Gall, 2020], ACV [Li and Todorovic, 2021]), and two timestamp-supervised methods (Li *et al.* [2021], GCN [Khan *et al.*, 2022]).

Supervision	Method	F1@{10,25,50}			Edit	Acc
Full	MS-TCN	76.3	74.0	64.5	67.9	80.7
	MS-TCN++	80.7	78.5	70.1	74.3	83.7
	BCN	82.3	81.3	74.0	74.3	84.4
	ASRF	84.9	83.5	77.3	79.3	84.5
Transcript	NN-Viterbi	-	-	-	-	49.4
	CDFL	-	-	-	-	54.7
Timestamp	Li <i>et al.</i>	73.9	70.9	60.1	66.8	75.6
	GCN	75.1	72.3	61.0	67.6	75.1
	Ours	<b>77.3</b>	<b>74.7</b>	<b>63.7</b>	<b>70.1</b>	<b>78.6</b>

Table 1: Comparison with the state-of-the-art on 50Salads.

Supervision	Method	F1@{10,25,50}			Edit	Acc
Full	MS-TCN	85.8	83.4	69.8	79.0	76.3
	MS-TCN++	88.8	85.7	76.0	83.5	80.1
	BCN	88.5	87.1	77.3	84.4	79.8
	ASRF	89.4	87.8	79.8	83.7	77.3
Timestamp	Li <i>et al.</i>	78.9	73.0	55.4	72.3	66.4
	GCN	81.5	77.5	60.8	75.6	66.1
	Ours	<b>83.7</b>	<b>79.8</b>	<b>65.4</b>	<b>77.2</b>	<b>70.1</b>

Table 2: Comparison with the state-of-the-art on GTEA.

Supervision	Method	F1@{10,25,50}			Edit	Acc
Full	MS-TCN	52.6	48.1	37.9	61.7	66.3
	MS-TCN++	64.1	58.6	45.9	65.6	67.6
	BCN	68.7	65.5	55.0	66.2	70.4
	ASRF	74.3	68.9	56.1	72.4	67.6
Transcript	NN-Viterbi	-	-	-	-	43.0
	CDFL	-	-	-	-	50.2
Set	SCT	-	-	-	-	30.4
	ACV	-	-	-	-	33.4
Timestamp	Li <i>et al.</i>	70.5	63.6	47.4	69.9	64.1
	GCN	67.9	61.0	45.3	67.0	61.4
	Ours	<b>71.2</b>	<b>64.6</b>	<b>48.9</b>	<b>71.6</b>	<b>65.7</b>

Table 3: Comparison with the state-of-the-art on Breakfast.

Results for three datasets are shown in Table 1, 2, 3. The best results are highlighted in bold. Under timestamp supervision, our method achieves state-of-the-art on all datasets. Besides, we find that our timestamp-supervised approach significantly outperforms all the set-supervised and transcript-supervised methods by a large margin. Furthermore, our method is comparable to or even better than some fully supervised models such as MS-TCN on the 50Salads and Breakfast datasets. In particular, on the largest Breakfast dataset, the performance is even comparable to one of the most advanced fully supervised methods, BCN.

### 4.4 Ablation Study

**Impact of Pseudo-Label Ensembling.** Given timestamp supervision and frame-wise visual features, we generate pseudo-labels with different methods and evaluate them with label accuracy and label rate in Table 4. We find that our temporal AGNES algorithm outperforms the other two. In




Figure 4: Visualization of pseudo-label sequences generated by several methods: (a) ground truth, (b) timestamp supervision, (c) energy function, (d) constrained k-medoids, (e) temporal AGNES, (f) pseudo-label ensembling. Frames in white intervals have no labels.

Energy Func	Cons K-m	Temp AGNES	GTEA		50Salads		Breakfast	
			Acc	Label	Acc	Label	Acc	Label
✓			75.6	100	79.9	100	72.7	100
	✓		77.1	100	77.9	100	75.2	100
		✓	78.4	100	81.5	100	75.1	100
✓	✓	✓	84.0	76.9	90.0	70.7	84.2	68.1
✓*	✓*	✓*	<b>86.8</b>	60.4	<b>95.5</b>	51.2	<b>88.8</b>	50.9

Table 4: Comparison of the quality of pseudo-labels generated by different methods. *Energy Func*, *Cons K-m*, and *Temp AGNES* are abbreviations of the energy function, constrained k-medoids, and temporal AGNES, respectively. *Acc* and *Label* denote the label accuracy and label rate, respectively.

the penultimate row, we can obtain pseudo-labels with both a high label rate and high accuracy when we ensemble these three algorithms, i.e., PLE. In practice, we further divide the visual features into RGB features and optical flow features, and ensemble these two types of features separately to get the final results, which are listed in the last row. The final pseudo-labels achieve the highest accuracy and over 50% label rate. We visualize the generated pseudo-label sequences in Figure 4. The provided timestamp supervision signal is too sparse because the label rate is only 0.4%. We can generate full pseudo-label sequences by any clustering algorithm, but there are too many wrong labels. PLE makes a trade-off between the label rate and accuracy.

**Ablations about Iterative Clustering.** Iterative clustering (IC) is an important part of our framework for which we conduct detailed ablation studies. We use pseudo-label sequences updated by IC to train the model. We can also replace IC with other methods to generate pseudo-labels for training. In Table 5, we first generate pseudo-label sequences with and without ambiguous intervals by PLE and Temporal AGNES to train the model directly. The pseudo-label sequences depend only on the visual features and do not change during training. Therefore, these two methods have the worst performance. The methods in the last three rows of the table all iteratively update the pseudo-labels based on high-level features during training. Li et al. [2021] use the energy function to generate full pseudo-label sequences. The modified version

Method	F1@{10,25,50}			Edit	Acc
PLE	72.2	69.0	57.3	64.4	76.6
Temp AGNES	72.2	68.3	56.3	64.5	73.2
Li et al.	75.4	72.4	61.8	68.1	76.8
IC (Li et al.)	74.7	71.8	62.1	66.2	77.6
IC (Ours)	<b>76.0</b>	<b>73.2</b>	<b>63.7</b>	<b>68.5</b>	<b>78.5</b>

Table 5: Comparison of supervision generated by different methods. The first two and the middle two rows represent that pseudo-labels are fixed and updated during training, respectively. *IC (Li et al.)* refers to using the energy function to generate pseudo-labels for the ambiguous intervals in pseudo-label sequences during IC. *IC (Ours)* is IC we propose.

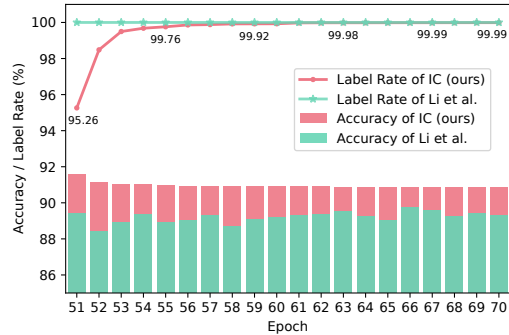


Figure 5: Comparison of the quality of the pseudo-labels generated by our proposed IC and the method of Li et al. during training. The method of Li et al. does not consider ambiguous intervals, and thus the label rate is always 100%.

of IC, called IC (Li et al.), uses the energy function to generate pseudo-labels for the frames within ambiguous intervals, which also yields full pseudo-label sequences. IC (Ours) generates pseudo-label sequences with ambiguous intervals to train the model and achieves the best results, which illustrates the importance of ambiguous intervals. As the training epoch increases, Figure 5 shows the variation of accuracy and label rate of pseudo-labels generated by energy function and IC. The accuracy of IC is always higher than that of Li et al., and the label rate of IC is increasing, indicating that the range of ambiguous intervals is decreasing.

**Effect of ambiguous intervals.** To further explore the effect of ambiguous intervals, we remove the labels of ambiguous intervals from these ground-truth label sequences under full supervision. Then we use the resulting label sequences with a certain unlabeled rate to train the model. Since the range of ambiguous intervals is decreasing during IC, we can record multiple sets of ambiguous intervals, which correspond to multiple label sequences with different unlabeled rates. As shown in Figure 6, the frames within the ambiguous intervals generated by PLE take up 48.1% of the total frames. So we can use the label sequences with an unlabeled rate of 48.1% to train the model. Similarly, the unlabeled rates of 5.7%, 0.8% and 0.09% represent the ambiguous intervals when IC runs to three intermediate epochs. The un-

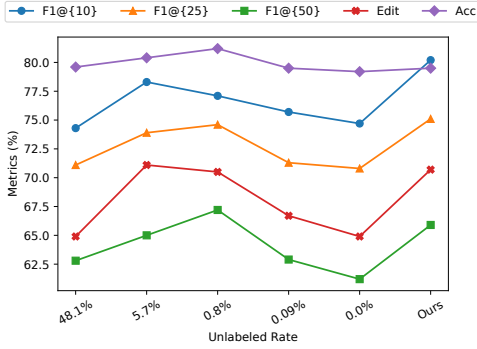


Figure 6: Comparison of the performance of training the model using ground-truth label sequences with different unlabeled rates, which can be obtained by removing the labels from the ground truth based on ambiguous intervals generated by the IC. For comparison, the last column shows the performance of our IC.

Dataset	$\mathcal{L}_{clu}$	F1@{10,25,50}			Edit	Acc
GTEA	✗	82.0	79.1	64.5	76.1	69.5
	✓	<b>83.7</b>	<b>79.8</b>	<b>65.4</b>	<b>77.2</b>	<b>70.1</b>
50Salads	✗	76.0	73.2	<b>63.7</b>	68.5	78.5
	✓	<b>77.3</b>	<b>74.7</b>	<b>63.7</b>	<b>70.1</b>	<b>78.6</b>

Table 6: Impact of the clustering loss  $\mathcal{L}_{clu}$  on two datasets.

labeled rate of 0% represents full supervision. The model performs best when the unlabeled rate is 0.8%, which indicates that about 0.8% labels in ground-truth label sequences are ambiguous and useless. We should train the model without these labels instead of using all the labels. Our IC does not use the information in the ground truth, but still achieves comparable results to the setting of unlabeled rate 0.8% and outperforms the fully supervised setting.

**Impact of the Clustering Loss.** To verify the effectiveness of our proposed clustering loss, we train the segmentation model with and without it. The metrics for both settings on two datasets are presented in Table 6. Additionally, we show the TSNE visualization of output features in Figure 7. We observe that the model’s output features become more discriminative with the help of clustering loss, leading to improvements in all metrics.

## 5 Conclusion

From the perspective of clustering, we propose pseudo-label ensembling and iterative clustering and introduce clustering loss to compose a framework for timestamp-supervised action segmentation. For the first time, we consider the ambiguity of action segment boundaries. Instead of directly generating pseudo-labels for all frames, our approach generates labels only for semantically explicit frames, which avoids the negative impact of inconsistent labels for ambiguous intervals at action segment boundaries. Experiments verify the effectiveness of our framework.




Figure 7: With or without the clustering loss, T-SNE visualization of all the learned output features in one video in the 50Salads dataset.

## A Temporal AGNES

Algorithm 2 summarizes the steps of our clustering algorithm, which we call the temporal AGNES algorithm. Function  $\text{dist}(\cdot, \cdot)$  calculates the Euclidean distance between visual features of two frames, and matrix  $D$  contains all the distances between any pair of frames. In contrast to the original AGNES algorithm, we make the following three changes. First, different timestamps should have distinct labels and cannot be merged. Therefore, we set the distance between different timestamp frame features to infinity to prevent them from being merged (step 9). Second, we only merge temporally adjacent clusters. Therefore, we only record and delete the boundaries of segments when merging two temporally contiguous clusters (step 3 and step 15). Third, since each action segment in the video contains a single timestamp, the algorithm stops when the number of clusters equals the number of timestamps (step 11). We utilize average linkage as the linkage criterion when merging clusters, i.e., the distance between a pair of clusters is determined by the average distance between frames (step 13 and step 14).

### Algorithm 2 Temporal AGNES clustering algorithm

**Input:** visual features:  $[x_1, x_2, \dots, x_T]$ , timestamp supervision:  $\{c_{\tau_1}, c_{\tau_2}, \dots, c_{\tau_N}\}$   
**Output:** pseudo-label sequence:  $[\hat{c}_1, \hat{c}_2, \dots, \hat{c}_T]$

- 1:  $q \leftarrow T$
- 2: **for**  $i = 1, 2, \dots, T$  **do**
- 3:  $b_i, d_i \leftarrow i, \text{dist}(x_i, x_{i+1})$
- 4: **for**  $j = 1, 2, \dots, T$  **do**
- 5:  $D[i, j] = \text{dist}(x_i, x_j)$
- 6: **end for**
- 7: **end for**
- 8: **for**  $n = 1, 2, \dots, N - 1$  **do**
- 9:  $D[\tau_n, \tau_{n+1}] \leftarrow \infty$
- 10: **end for**
- 11: **while**  $q > N$  **do**
- 12:  $i = \text{argmin}_j d_j$
- 13:  $d_{i-1} \leftarrow \frac{d_{i-1} * (b_i - b_{i-1}) * (b_{i+1} - b_i) + \sum_{j=b_{i-1}}^{b_i} \sum_{k=b_{i+1}}^{b_{i+2}} D[j, k]}{(b_i - b_{i-1}) * (b_{i+2} - b_i)}$
- 14:  $d_i \leftarrow \frac{d_i * (b_{i+2} - b_{i+1}) * (b_{i+3} - b_{i+2}) + \sum_{j=b_i}^{b_{i+1}} \sum_{k=b_{i+2}}^{b_{i+3}} D[j, k]}{(b_{i+3} - b_{i+2}) * (b_{i+2} - b_i)}$
- 15: **delete**  $b_{i+1}, d_{i+1}$
- 16:  $q \leftarrow q - 1$
- 17: **end while**
- 18: **for**  $n = 1, 2, \dots, N$  **do**
- 19:  $[\hat{c}_{b_n}, \hat{c}_{b_n+1}, \dots, \hat{c}_{b_{n+1}-1}] \leftarrow [c_{\tau_n}, c_{\tau_n}, \dots, c_{\tau_n}]$
- 20: **end for**

## Acknowledgments

This work was supported by the 2022 Special Research Assistant Grant project E3YD5901.

## References

- [Behrmann *et al.*, 2022] Nadine Behrmann, S. Alireza Golestaneh, Zico Kolter, Juergen Gall, and Mehdi Noroozi. Unified fully and timestamp supervised temporal action segmentation via sequence to sequence translation. In *ECCV*, 2022.
- [Carreira and Zisserman, 2017] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [Chang *et al.*, 2019] Chien-Yi Chang, De-An Huang, Yanan Sui, Li Fei-Fei, and Juan Carlos Nieves. D3tw: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3546–3555, 2019.
- [Dietterich, 2002] Thomas G Dietterich. Ensemble learning. *The handbook of brain theory and neural networks*, 2(1):110–125, 2002.
- [Du *et al.*, 2022] Dazhao Du, Bing Su, Yu Li, Zhongang Qi, Lingyu Si, and Ying Shan. Do we really need temporal convolutions in action segmentation? *arXiv preprint arXiv:2205.13425*, 2022.
- [Farha and Gall, 2019] Yazan Abu Farha and Jurgen Gall. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3575–3584, 2019.
- [Fathi *et al.*, 2011] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In *CVPR 2011*, pages 3281–3288. IEEE, 2011.
- [Fayyaz and Gall, 2020] Mohsen Fayyaz and Jurgen Gall. Sct: Set constrained temporal transformer for set supervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 501–510, 2020.
- [Han and Kamber, 2006] Jiawei Han and Micheline Kamber. Data mining: concepts and techniques, 2nd. *University of Illinois at Urbana Champaign: Morgan Kaufmann*, 2006.
- [Huang *et al.*, 2016] De-An Huang, Li Fei-Fei, and Juan Carlos Nieves. Connectionist temporal modeling for weakly supervised action labeling. In *European Conference on Computer Vision*, pages 137–153. Springer, 2016.
- [Isken *et al.*, 2019] Ahmet Isken, Giorgos Toliass, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5070–5079, 2019.
- [Ishikawa *et al.*, 2021] Yuchi Ishikawa, Seito Kasai, Yoshimitsu Aoki, and Hirokatsu Kataoka. Alleviating over-segmentation errors by detecting action boundaries. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2322–2331, 2021.
- [Khan *et al.*, 2022] Hamza Khan, Sanjay Haresh, Awais Ahmed, Shakeeb Siddiqui, Andrey Konin, M Zeeshan Zia, and Quoc-Huy Tran. Timestamp-supervised action segmentation with graph convolutional networks. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2014.
- [Kuehne *et al.*, 2014] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*, pages 780–787, 2014.
- [Kuehne *et al.*, 2016] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-to-end generative framework for video segmentation and recognition. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–8. IEEE, 2016.
- [Kuehne *et al.*, 2017] Hilde Kuehne, Alexander Richard, and Juergen Gall. Weakly supervised learning of actions from transcripts. *Computer Vision and Image Understanding*, 163:78–89, 2017.
- [Lea *et al.*, 2017] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *CVPR*, pages 156–165, 2017.
- [Lee and others, 2013] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, page 896, 2013.
- [Li and Todorovic, 2021] Jun Li and Sinisa Todorovic. Anchor-constrained viterbi for set-supervised action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9806–9815, 2021.
- [Li *et al.*, 2019] Jun Li, Peng Lei, and Sinisa Todorovic. Weakly supervised energy-based learning for action segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6243–6251, 2019.
- [Li *et al.*, 2020] Shi-Jie Li, Yazan AbuFarha, 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*, 2020.
- [Li *et al.*, 2021] Zhe Li, Yazan Abu Farha, and Jurgen Gall. Temporal action segmentation from timestamp supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8365–8374, 2021.



- [Qiang *et al.*, 2023] Wenwen Qiang, Jiangmeng Li, Bing Su, Jianlong Fu, Hui Xiong, and Wen Jirong. Meta attention-generation network for cross-granularity few-shot learning. *International Journal of Computer Vision*, 2023.
- [Rahaman *et al.*, 2022] Rahul Rahaman, Dipika Singhania, Alexandre Thiery, and Angela Yao. A generalized and robust framework for timestamp supervision in temporal action segmentation. In *ECCV*, pages 279–296. Springer, 2022.
- [Richard *et al.*, 2018a] Alexander Richard, Hilde Kuehne, and Juergen Gall. Action sets: Weakly supervised action segmentation without ordering constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5987–5996, 2018.
- [Richard *et al.*, 2018b] Alexander Richard, Hilde Kuehne, Ahsan Iqbal, and Juergen Gall. Neuralnetwork-viterbi: A framework for weakly supervised video learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7386–7395, 2018.
- [Singh *et al.*, 2016] Bharat Singh, Tim K Marks, Michael Jones, Oncel Tuzel, and Ming Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1961–1970, 2016.
- [Stein and McKenna, 2013] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 729–738, 2013.
- [Tang *et al.*, 2012] Kevin Tang, Li Fei-Fei, and Daphne Koller. Learning latent temporal structure for complex event detection. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1250–1257. IEEE, 2012.
- [Wang *et al.*, 2020] Zhenzhi Wang, Ziteng Gao, Limin Wang, Zhifeng Li, and Gangshan Wu. Boundary-aware cascade networks for temporal action segmentation. In *ECCV*, pages 34–51. Springer, 2020.
- [Yang *et al.*, 2017] Bo Yang, Xiao Fu, Nicholas D Sidiropoulos, and Mingyi Hong. Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In *international conference on machine learning*, pages 3861–3870. PMLR, 2017.
- [Yi *et al.*, 2021] Fangqiu Yi, Hongyu Wen, and Tingting Jiang. Asformer: Transformer for action segmentation. In *The British Machine Vision Conference*, 2021.