





A Generalized & Robust Framework For Timestamp Supervision in Temporal Action Segmentation

Rahul Rahaman¹, Dipika Singhania², Alexandre Thiery³, and Angela Yao⁴

rahul.rahaman@u.nus.edu, dipika16@comp.nus.edu.sg, a.h.thiery@nus.edu.sg,
ayao@comp.nus.edu.sg
National University of Singapore

Abstract. In temporal action segmentation, Timestamp Supervision requires only a handful of labelled frames per video sequence. For unlabelled frames, previous works rely on assigning hard labels, and performance rapidly collapses under subtle violations of the annotation assumptions. We propose a novel Expectation-Maximization (EM) based approach that leverages the label uncertainty of unlabelled frames and is robust enough to accommodate possible annotation errors. With accurate timestamp annotations, our proposed method produces SOTA results and even exceeds the fully-supervised setup in several metrics and datasets. When applied to timestamp annotations with missing action segments, our method presents stable performance. To further test our formulation’s robustness, we introduce the new challenging annotation setup of SkipTag Supervision. This setup relaxes constraints and requires annotations of any fixed number of random frames in a video, making it more flexible than Timestamp Supervision while remaining competitive.

Keywords: Timestamp Supervision, Temporal Action Segmentation

1 Introduction

Temporal action segmentation partitions and classifies a sequence of actions in long untrimmed video. In a fully-supervised setting, the exact action boundaries and actions of every frame are labelled for all training videos. **Timestamp Supervision (TSS)** [24,22] is a lightweight alternative in which the annotator labels one frame from each action in the video (Fig 1). TSS requires magnitudes fewer labels than the fully-supervised setting¹, but it has a strict constraint – there must be a timestamp for **every action** in the video sequence. To fulfill this constraint, annotators must watch the entire video carefully, so as not to miss any action segments. This affects the annotation efficiency and leads to only a 35% reduction in annotation time compared to full supervision. Furthermore, annotators may *mistakenly* skip actions, as observed in our user study on

¹ On Breakfast Actions [15], only 0.032% of the frames need labels!




Fig. 1: **Annotations levels.** Full-supervision requires labels for every frame in the video. Timestamp Supervision requires labels for a *single* frame from *each* action segment, although annotators may miss some action segments.

timestamp labelling. The annotators’ slow performance and mistakes highlight the impracticality of the constraints of Timestamp Supervision.

Existing works on TSS [24,22] are based on strong assumptions. For example, [24] directly parameterizes the position and width of action segments around the annotated timestamps, while [22] sets one action boundary with a handcrafted energy function between two annotated frames and assigns hard labels to in-between unlabelled frames. The performance of these methods collapses when annotation errors are introduced (see Fig. 4) as the errors violate the condition “*one* timestamp for *every* action segment”. Note that this sort of annotation error differs from the typical annotation error of mislabelling as it emerges from the difficulty of adhering to annotation constraints.

In this work, we first propose a novel, general model formulation for Timestamp Supervision, called **EM-TSS**. EM-TSS uses the first principles of an Expectation-Maximization (E-M) formulation to model frame-wise labels conditioned on action boundary locations. We restructure the EM formulation for the temporal segmentation task so that the maximization step reduces to a frame-wise cross-entropy minimization. Unlike the heuristics of [24,22], EM-TSS has theoretical groundings and assumes only an extremely weak prior distribution on segment length. EM-TSS also makes no additional assumptions per action segment and can leverage the label uncertainty provided by the E-M formulation to avoid assigning hard labels. EM-TSS surpasses previous works [24,22] for the majority of metrics by a significant margin. For several datasets and metrics, it even outperforms the fully-supervised setting with only a handful of labels. We attribute the success to the use of more accurate (albeit fewer) labels and verify with a user study and ablations on the position of the timestamps.

Next, we generalize our EM framework into **EM-Gen** to handle errors in Timestamp Supervision due to annotators missing action segments between timestamps. EM-Gen’s performance remains stable under such annotation error and can tolerate up to 20% missing segments, with a marginal drop in performance compared to the TSS methods. To further challenge the performance limits of EM-Gen, we introduce a new *SkipTag supervision* to allow for random frame annotations anywhere in the video. SkipTag supervision is far more flexible than timestamp supervision and is much less restrictive. Applying EM-Gen to SkipTag produces good results despite this weaker form of annotation.

Summarizing the **contributions** in this work,

1. We propose a novel E-M based method, EM-TSS, for action segmentation under Timestamp Supervision. EM-TSS surpasses SOTA timestamp methods by a large margin and is competitive with fully-supervised methods.
2. We generalize EM-TSS into EM-Gen, which, unlike competing works, remains robust under subtle timestamp annotation errors.
3. We push the limits of EM-Gen and apply it to the weaker SkipTag supervision, which enables the free annotation of a random selection of frames.
4. Through a user-study, we compare the annotation efforts for Timestamp and SkipTag supervision and investigate the ambiguity of boundary annotations.

2 Related Works

Temporal Action Segmentation assigns frame-wise action labels to long, untrimmed video; it is a 1-D temporal analogue of semantic segmentation. Recent approaches [21,30,39,12,4,38,34,11] have used pre-extracted I3D [32] snippet-level features as input and modelled the temporal structure using various models, such as HMMs [15], RNNs [33] or TCNs [18,38,21,4]. Temporal Convolution Networks (TCNs) [18,21,8] are shown to be the most effective, with the popular SOTA being MSTCNs [21]. We use I3D features as input and the MSTCN architecture as a backbone, similar to previous TSS work [22].

Fully-Supervised Methods [18,8,15] require action labels to be provided for all frames. To reduce the annotation costs, **weakly supervised** methods only require an ordered [7,29,3,19,35,26] or unordered list of actions [28,20,10,6] as supervision. **Unsupervised** approaches [31,17,37] do not require any annotation and are primarily focused on clustering segments.

Timestamp Supervision (TSS) has recently emerged as a new form of weak supervision that requires a single frame annotation for each action [13,16,24,22]. TSS thus requires segmentation methods to utilize and estimate a huge amount of unlabelled training frames. TSS is the preferred choice over other weak supervision methods due to its higher accuracy yet lower annotation cost [24].

Existing works [24,23] expand action segments' length around the annotated frames based on discriminative probability outputs. Li et al. [22] fix a boundary between two consecutive annotated frames based on similarity measures and assign the frames in between to either of the two actions. Our work emphasizes boundary localization. However, unlike solutions using hard boundaries [22] or distributional assumption [24,22], we formulate boundary estimation and annotate unlabelled frames in the training videos with the principles of the E-M algorithm. Our problem formulation and derivation of the steps for the E-M algorithm stem naturally from the model assumptions of the supervised setup.

3 Preliminaries

3.1 Temporal Action Segmentation Task

Consider a video V with duration T . For fully-supervised action segmentation, each frame $V[t] : t \leq T$ is labelled with the ground truth action label

$y[t] \in \{1, \dots, C\}$ of C possible action classes. The standard practice in action segmentation [21,39,34] is to use pre-trained frame-wise features, e.g., I3D [2] features. We denote these features as $\mathbf{f}[t] \in \mathbb{R}^d$. The TCN takes \mathbf{f} as input and generates frame-wise action class probabilities \mathbf{p} . For frame t , the vector $\mathbf{p}[t]$ is a C -dimensional vector and $p_{t,c} := \mathbf{p}[t, c]$ is the probability value assigned to action class c . A cross-entropy loss can be used to train the model.

$$\mathcal{L}_{\text{CE}} := - \sum_{t=1}^T \sum_{c=1}^C \omega_{t,c} \cdot \log p_{t,c} \quad (1)$$

Under full supervision, ω can be set as the ground truth labels, *i.e.* $\omega_{t,c} = \mathbb{I}[y[t] = c]$. To use the same loss under Timestamp Supervision, however, the weights $\omega_{t,c}$ must be estimated for the unlabelled frames where $y[t]$ is unknown. Previous works [24,22] do so by setting hard weights, *i.e.* $\omega_{t,c} \in \{0, 1\}$, through some rules or a ranking scheme.

3.2 E-M Algorithm


Expectation-Maximization (E-M) [5] is a classic iterative method for maximum likelihood estimation (MLE) with hidden variables and/or missing data. We consider the classical form with missing data. Formally, let X^{ob}, X^{ms} be the observed and missing parts of the data, respectively, and $\mathcal{P}(X^{ob}, X^{ms} | \Theta)$ be the data likelihood conditioned on some unknown parameter Θ , e.g., the weights of a neural network model. As X^{ob} is observed and X^{ms} is missing, directly maximizing \mathcal{P} as a function of Θ , and hence obtaining the MLE of Θ^* , is not possible. The E-M algorithm provides an iterative approach to obtain the MLE of Θ under this setting through the use of a \mathcal{Q} -function. In simple terms, the \mathcal{Q} -function serves as an expected version of the log-likelihood \mathcal{P} by replacing terms involving unobserved X^{ms} with the expected value based on our current estimate of the parameter. Formally, let $\Theta^{(m)}$ be the current estimate of Θ after the m^{th} iteration. At iteration $(m + 1)$, the *Expectation step* calculates the \mathcal{Q} -function:

$$\mathcal{Q}(\Theta, \Theta^{(m)}) = \mathbb{E}_{\mathbf{z}}[\log \mathcal{P}(X^{ob}, \mathbf{z} | \Theta)], \quad (2)$$

where $\mathbf{z} \sim f(X^{ms} | X^{ob}, \Theta^{(m)})$, *i.e.* the posterior distribution of the unobserved data X^{ms} given the observed data X^{ob} and the current estimate of the parameters $\Theta^{(m)}$. Note that the term in Eq. 2 is a function of the parameters Θ . In the *Maximization step*, the \mathcal{Q} function is then maximized with respect to the parameters Θ to obtain the new estimate

$$\Theta^{(m+1)} = \arg \max_{\Theta} \mathcal{Q}(\Theta, \Theta^{(m)}). \quad (3)$$

This process is repeated until convergence, *i.e.* $\Theta^{(m+1)}$ is desirably close to $\Theta^{(m)}$, or for a fixed number of iterations. The model (TCN) parameter learning happens during this M-step (Maximization), whereas the E-step (Expectation) iteratively updates the criteria to be optimized in the M-step.

Fig. 2: Depiction of **timestamp segment** \mathcal{S}_k .

4 Method

The assignment of hard labels to each frame may lead to some ambiguities at the action boundaries and deteriorate performance (see Sec. 5.4). To avoid the pitfalls of hard labels, we allow soft weights for ω in Eq. 1. We arrive at these soft weights using the E-M framework, as we show for Timestamp Supervision (Sec. 4.2), and extend it to the case of missing timestamps (Sec. 4.3).

4.1 Segment-Based Notation

While a video is typically treated as a sequence of frames (Sec. 3.1), it can also be regarded as a sequence of action segments. Each segment is a set of contiguous frames belonging to one action, represented by a single color in Fig. 1. A video with K action segments will have action classes c_1, \dots, c_K out of C action classes. K may vary for different videos; the same action may also occur multiple times in a video. As the adjacent action segments share boundaries, $c_k \neq c_{k+1}$ for all $k < K$. We denote s_k to be the starting frame for action segment k , with $s_1 \equiv 1$ and $s_k < s_{k+1} \leq T$. Frame s_k is also the *boundary* between action segments $k-1$ and k . In the fully-supervised setup, every frame is labelled; this is equivalent to labelling start frame s_k and action class c_k for each segment k .

4.2 Timestamp Supervision

Timestamp Supervision provides one labelled frame or ‘*timestamp*’ per action segment, resulting in K labelled frames t_1, \dots, t_K and action labels c_1, \dots, c_K . There is exactly one timestamp, randomly positioned within the segment, for *every* action segment, *i.e.* $s_k \leq t_k < s_{k+1}$, or equivalently $t_{k-1} < s_k \leq t_k, \forall k$. As the action boundaries are unknown, we additionally define a *timestamp segment*; the k^{th} timestamp segment is denoted by \mathcal{S}_k and is bounded by frames $[t_{k-1}, t_k - 1]$ (see Fig.2). Within \mathcal{S}_k , the label $l := y[t_{k-1}] = c_{k-1}$ denotes the ‘*left action*’ class and $r := y[t_k] = c_k$ denotes the ‘*right action*’ class².

Timestamp Segment Likelihood: If s_k is known, all the frames $[t_{k-1}, s_k)$ can be labelled with action l and frames $[s_k, t_k)$ can be labelled with action r . Hence,

² We do not discuss the first and last timestamp segments $\mathcal{S}_1 := [1, t_1)$ and $\mathcal{S}_{K+1} := [t_K, T]$ as their action labels c_1, c_K are known.

the likelihood of segment \mathcal{S}_k given $s_k = j$, denoted as $\mathcal{P}_j(\mathcal{S}_k | \Theta)$, becomes

$$\mathcal{P}_j(\mathcal{S}_k | \Theta) := \mathcal{P}(\mathcal{S}_k | s_k = j, \Theta) = \prod_{i=t_{k-1}}^{j-1} p_{i,l}^{\Theta} \prod_{i=j}^{t_k-1} p_{i,r}^{\Theta}. \quad (4)$$

In Eq. 4, we add a super-script Θ to $p_{i,c}^{\Theta}$ to emphasize the dependence of the probability $p_{t,c}$ on the model parameters Θ of the temporal convolutional network. If s_k is known for each k , then maximizing the likelihood of Eq. 4 after taking the negative log and summing over all K segments simplifies to minimizing the fully-supervised frame-wise cross-entropy (CE) loss of Eq. 1. Note that Eq. 4 and fully-supervised Eq. 1 both assume independence of the frame-wise probability when conditioned on the network parameters and architecture.

In Timestamp Supervision, the boundary s_k is unknown, so we treat it as a random variable, with some prior probability of it being located at frame j , *i.e.* $\pi_k(j) := \pi(s_k = j)$. The joint likelihood of segment \mathcal{S}_k and boundary s_k then becomes $\mathcal{P}(\mathcal{S}_k, s_k = j | \Theta) = \mathcal{P}_j(\mathcal{S}_k | \Theta) \cdot \pi_k(j)$. We further elaborate on the prior distribution in subsection 4.4. As boundary s_k can be located at any of the frames $[t_{k-1}, t_k)$, the summarized joint likelihood can be written as

$$\mathcal{P}(\mathcal{S}_k, s_k | \Theta) = \prod_{j=t_{k-1}}^{t_k-1} [\mathcal{P}_j(\mathcal{S}_k | \Theta) \cdot \pi_k(j)]^{\mathbb{I}[s_k=j]}. \quad (5)$$

Q-Function & Posterior. Following Eq. 2, the \mathcal{Q} -function for timestamp segment \mathcal{S}_k is obtained by taking the expectation of the log of the likelihood in Eq. 5 with respect to $f(s_k | \mathcal{D}; \Theta^{(m)})$. This represents the posterior distribution of boundary s_k given observed data \mathcal{D} under the previous estimate of parameter $\Theta^{(m)}$. Here, observed data \mathcal{D} contains all the training information, such as input video features \mathbf{f} and the annotations $y[t_{k-1}], y[t_k]$. The \mathcal{Q} -function for \mathcal{S}_k , denoted by $\mathcal{Q}_k(\Theta, \Theta^{(m)})$, thus becomes

$$\begin{aligned} \mathcal{Q}_k(\Theta, \Theta^{(m)}) &= \sum_{j=t_{k-1}}^{t_k-1} \mathbb{E}_{s_k} \left[\mathbb{I}[s_k=j] \cdot \log \left(\mathcal{P}_j(\mathcal{S}_k | \Theta) \cdot \pi_k(j) \right) \right] \\ &= \Lambda + \sum_{j=t_{k-1}}^{t_k-1} \mathbb{P}^{(m)}[s_k=j | \mathcal{D}] \cdot \log \mathcal{P}_j(\mathcal{S}_k | \Theta), \end{aligned} \quad (6)$$

where $\mathbb{P}^{(m)}[s_k = j | \mathcal{D}]$ is the posterior probability that s_k is located at frame j given the data and under the model parameter $\Theta^{(m)}$. Λ consists of the terms (prior $\pi_k(j)$) that are constant with respect to the parameter Θ and hence can be ignored during the maximization of \mathcal{Q} with respect to Θ . The posterior probabilities can be obtained by

$$\mathbb{P}^{(m)}[s_k = j | \mathcal{D}] = \frac{\mathcal{P}_j(\mathcal{S}_k | \Theta^{(m)}) \cdot \pi_k(j)}{\sum_{i=t_{k-1}}^{t_k-1} \mathcal{P}_i(\mathcal{S}_k | \Theta^{(m)}) \cdot \pi_k(i)}, \quad (7)$$

where $\mathcal{P}_j(\mathcal{S}_k | \Theta^{(m)})$ is defined in Eq. 4. The final \mathcal{Q} -function for the whole video is calculated by $\mathcal{Q}(\Theta, \Theta^{(m)}) = \sum_k \mathcal{Q}_k(\Theta, \Theta^{(m)})$.

\mathcal{Q} -function to Frame-Wise CE Loss. The formulation of the \mathcal{Q} -function in Eq. 6 requires estimating two terms for each frame j . The first is $\mathbb{P}^{(m)}[s_k = j | \mathcal{D}]$, which can be calculated only once per E-step as it does not depend on Θ . The second term $\mathcal{P}_j(\mathcal{S}_k | \Theta)$, however, must be calculated for every iteration of the M-step. Even after obtaining frame-wise probabilities $p_{j,c}^\Theta$ from the model, every calculation of $\log \mathcal{P}_j(\mathcal{S}_k | \Theta)$ for each frame j (as given in Eq. 4) takes linear time with respect to video length, i.e., of complexity $\mathcal{O}(T)$ per frame j . This may make each epoch of the parameter learning in the M-step computationally expensive. To reduce the computational complexity, we simplify Eq. 6 by substituting $\log \mathcal{P}_j(\mathcal{S}_k; \Theta)$ from Eq. 4 and rearranging the sums to get

$$\mathcal{Q}_k(\Theta, \Theta^{(m)}) = \sum_{j=t_{k-1}}^{t_k-1} \sum_{c=1}^C \omega_{j,c}^{(m)} \cdot \log p_{j,c}^\Theta, \quad (8)$$

where
$$\omega_{j,c}^{(m)} = \begin{cases} \sum_{i=j+1}^{t_k-1} \mathbb{P}^{(m)}[s_k = i | \mathcal{D}] & : c = l \\ 1 - \omega_{j,l}^{(m)} & : c = r \\ 0 & : \text{otherwise.} \end{cases}$$

This new form of $\mathcal{Q}_k(\Theta, \Theta^{(m)})$ is more intuitive than Eq. 6 and, as desired at the start of this section, reduces to a negative frame-wise CE loss (Eq. 1). Furthermore, our posterior weights $\omega^{(m)}$ are not hard labels, and the weights capture the label uncertainty of the unlabelled frames. This simplified formulation converges in very few iterations and adds negligible differences to the training time compared to a fully-supervised approach. The weights $(\omega_{j,l}^{(m)}, \omega_{j,r}^{(m)})$ are actually the posterior probabilities (given data and parameters) that frame j belongs to the left and right actions l and r , respectively. The weight $\omega_{j,l}^{(m)}$ acts as the current estimate of $\mathbb{I}[y[j] = l]$ (see Eq. 1) when $y[j]$ is unknown. Furthermore, regardless of the current parameter $\Theta^{(m)}$, the weights have the property that $\omega_{t_{k-1},l}^{(m)} = 1$ and $\omega_{j,l}^{(m)} \geq \omega_{j+1,l}^{(m)}$, $\omega_{j,r}^{(m)} \leq \omega_{j+1,r}^{(m)}$ for any $t_{k-1} \leq j < t_k$. This means the weight assigned to left action class l is 1 for frame t_{k-1} (as it should be), and it gradually decreases as j increases, with the converse being true for the action class r . These weights can be calculated during the E-step and be used in a CE loss without additional computation during the M-step.

Initialization and Overall Procedure. The E-M algorithm requires good initialization. Following [22], we initialize the TCN parameters with $\Theta^{(0)}$ derived from a ‘Naive’ baseline learned with a frame-wise cross-entropy loss applied *only* to the annotated frames t_k . From this initialization, we perform the first E-step to obtain the frame-wise weights $\omega_{j,c}^{(0)}$. Thereafter, we iterate between Maximization (maximize \mathcal{Q}) and Expectation (update $\omega^{(m)}$) steps, which we call **EM-TSS**, to find the best estimate Θ^* (see the pseudo-code in algorithm 1).

Algorithm 1 Our iterative procedure for EM-TSS

```

1: for  $epoch \leq N^{init}$  do ▷ Initialization:
2:   calculate  $\mathcal{L} := -\sum_{k=1}^K \log p_{t_k, c_k}^{\Theta}$ 
3:   minimize  $\mathcal{L}$ , update  $\Theta$ 
4: set  $\Theta^{(0)} = \Theta$ , set  $m = 0$ 
5: for  $m < M$  do
6:   calculate  $\omega^{(m)}$  using  $\Theta^{(m)}$  ▷ Expectation:
7:   for  $epoch \leq N^{max}$  do ▷ Maximization:
8:     maximize  $\mathcal{Q}(\Theta, \Theta^{(m)})$ , update  $\Theta$ 
9:   set  $\Theta^{(m+1)} = \Theta$ 
10: set final  $\Theta^* = \Theta^{(M)}$ 

```

4.3 Timestamp Supervision Under Missed Timestamps

Timestamp Supervision places a strong constraint that there is one timestamp for *every* action segment present in the video. In practice, as revealed by our user study (see *Supplementary*), annotators may miss action segments. These missed segments, as we show later in Sec 5, cause TSS methods to collapse rapidly.

As such, we target the impracticality of having perfectly accurate timestamps and develop a method that is more robust to such annotation errors. Specifically, we generalize our E-M framework **EM-TSS** into **EM-Gen** to accommodate possible missing segments between two consecutive annotated timestamps, as depicted in Fig. 3. The overall E-M framework remains the same as Timestamp supervision, but differs in the derivation of the associated \mathcal{Q} -function and weights ω . We consider the timestamp segment $[t_{k-1}, t_k]$ with the following cases .

Case 1 (C_1): Correct timestamp annotation. In this case, there will be exactly *one action boundary* s , where $t_{k-1} < s \leq t_k$ and $r \neq l$. This case is the same as a typical timestamp segment of TSS; given boundary location s , the segment likelihood is the same as Eq. 4 earlier,

$$\mathcal{P}(S|C_1, s, \Theta) = \prod_{i=t_{k-1}}^{s-1} p_{i,l}^{\Theta} \prod_{i=s}^{t_k-1} p_{i,r}^{\Theta}.$$




Fig. 3: **Correct vs. Missed Segments.** Timestamp segment \mathcal{S} between consecutive timestamps t_{k-1}, t_k with action labels l, r respectively. When correct, there is one unknown action boundary s ; when a segment is missed, there are actually two unknown boundaries, s_1 and s_2 , and an unknown action c in between.

Case 2 (C_2): Missed timestamp annotation. In this case, there will be two action boundaries s_1, s_2 , which define a middle action segment such that $t_{k-1} < s_1 < s_2 \leq t_k$. If the middle action label is c , where $c \neq l$ and $c \neq r$ (this allows for both $r \neq l$ and $r = l$), the segment likelihood is

$$\mathcal{P}(\mathcal{S}|C_2, s_1, s_2, c, \Theta) = \prod_{i=t_{k-1}}^{s_1-1} p_{i,l}^{\Theta} \prod_{i=s_1}^{s_2-1} p_{i,c}^{\Theta} \prod_{i=s_2}^{t_k-1} p_{i,r}^{\Theta}.$$

Similar to EM-TSS, to calculate the \mathcal{Q} -function, we now need to combine the log-likelihoods with the corresponding weights as the posterior probabilities. The \mathcal{Q} -function for the segment \mathcal{S} is calculated as

$$\begin{aligned} \mathcal{Q}(\Theta, \Theta^{(m)}) = & \sum_s \mathbb{P}^{(m)}[C_1, s|\mathcal{D}] \cdot \log \mathcal{P}(\mathcal{S}|C_1, s, \Theta) + \\ & \sum_{\tau} \mathbb{P}^{(m)}[C_2, \tau|\mathcal{D}] \cdot \log \mathcal{P}(\mathcal{S}|C_2, \tau, \Theta), \end{aligned} \quad (9)$$

where τ indexes all valid triplets (s_1, s_2, c) of Case 2. In the first term, $\mathbb{P}^{(m)}[C_2, s|\mathcal{D}]$ denotes the posterior probability that the segment belongs to Case 1 and the action boundary is located at s , given the data and under parameter $\Theta^{(m)}$. The second term contains an analogous posterior for Case 2 and some triplet τ . The probabilities of Eq. 9 can be estimated in the same way as Eq. 7 of EM-TSS. A few components in the \mathcal{Q} -function have high computational complexity, e.g., the second sum in Eq. 9 is of $\mathcal{O}(T^2)$ where T is the video duration. However, a few rearrangements similar to the EM-TSS change the \mathcal{Q} -function into the same form as in Eq. 8. We defer the exact forms to the *Supplementary*. The final posterior weights become

$$\begin{aligned} \omega_{j,l}^{(m)} &= \sum_{s>j} \mathbb{P}^{(m)}[C_1, s|\mathcal{D}] + \sum_{\tau:s_1>j} \mathbb{P}^{(m)}[C_2, \tau|\mathcal{D}] \\ \omega_{j,c}^{(m)} &= \sum_{s_1 \leq j} \sum_{s_2 > j} \mathbb{P}^{(m)}[C_2, s_1, s_2, c|\mathcal{D}] \quad \text{if } c \neq l, c \neq r, \end{aligned}$$

and $\omega_{j,r}^{(m)} = 1 - \omega_{j,l}^{(m)} - \sum_c \omega_{j,c}^{(m)}$. We calculate these weights during the E-step and simply apply the negative cross-entropy loss during the M-steps to perform our EM-Gen. The initialization and training are the same as algorithm 1.

4.4 Prior Distribution

We apply a prior on the position of the action boundaries s_k . Like previous works [35,16,27], we assume that the length of an action segment for action c is a random variable following a *Poisson* distribution $\text{Pois}(\mu_c)$ with a mean length $\mu_c > 0$. Using this assumption, one can arrive that $\pi(s_k = j)$ of Eq. 7 (probability of the k^{th} action boundary s_k is located at frame j), follows a binomial distribution, *i.e.* $s_k \sim \text{Bin}(T, p_k)$, with parameters T , the video length, and $p_k = \frac{\sum_{n \leq k-1} \mu_{c_n}}{\sum_{n \leq K} \mu_{c_n}}$, where K is the number of action segments in the video. For EM-Gen, the prior is less straightforward and we defer it to the *Supplementary*.

4.5 Loss Function

We apply $\mathcal{L}_{\text{EM}} := -\frac{1}{T}\mathcal{Q}(\Theta, \Theta^{(m)})$ during the M-step, *i.e.* we minimize the negative \mathcal{Q} -function normalized by the number of frames T as a loss for learning the TCN. We also add two auxiliary losses, similar to [22]: a *transition loss* $\mathcal{L}_{\text{TR}} = \frac{1}{TC} \sum_{t,c} \min(|\delta_{t,c}|, \varepsilon)$ to penalize small inter-frame differences (up to a threshold ε) to limit over-segmentation; and a *confidence loss* $\mathcal{L}_{\text{Conf}} = \frac{1}{T} \sum_{k=1}^K \sum_{t=t_{k-1}}^{t_k-1} (\delta_{t,y[t_{k-1}]}^+ + \delta_{t,y[t_k]}^-)$ to ensure peak confidence at the annotated timestamps. Here, $\delta_{t,c} := \log p_{t,c}^{\Theta} - \log p_{t-1,c}^{\Theta}$, $\delta^+ := \max(\delta, 0)$, $\delta^- := \max(-\delta, 0)$, C is the number of action classes, and K is the total number of annotated timestamps in the video. It is worth noting that our posterior weights from Eq. 8 already induce this property. The final loss function is

$$\mathcal{L} = \mathcal{L}_{\text{EM}} + \lambda_{\text{TR}} \cdot \mathcal{L}_{\text{TR}} + \lambda_{\text{Conf}} \cdot \mathcal{L}_{\text{Conf}}$$

We use the values of λ_{TR} , λ_{Conf} as suggested by [22].

4.6 SkipTag Supervision

Even though TSS requires annotation of very few frames, the constraint of exactly one frame for every action segment incurs significant amounts of annotation time (see Sec. 5.4). As the EM-Gen formulation tolerates missing frames between segments, we are motivated to introduce a more flexible form of frame-wise annotation. To this end, we propose the new SkipTag Supervision. SkipTag Supervision is less restrictive and allows annotators to simply label the actions of a random set of timestamps placed somewhat evenly. However, two consecutively annotated timestamps t_{k-1} and t_k now have an unknown number of action boundaries in between, including no boundaries. As an experimental setting, SkipTag annotations are more challenging than TSS not only because of removing the TSS annotation constraint but also because the number of annotated timestamps do not correspond to the (unknown) number of action segments. TSS implicitly gives the number of segments per video since the annotations match the exact number of action segments. Thus, Skiptag annotation is a weaker form of supervision than TSS annotation.

To accommodate SkipTag annotations, EM-Gen must also handle the case of no action boundaries between timestamps t_{k-1}, t_k . We defer the formulation related to this new case to the *Supplementary*. To limit complexity, we also do not consider beyond two action boundaries between consecutive timestamps t_{k-1}, t_k . Experimentally, this occurs very infrequently (see *Supplementary*). However, this simplification is our design choice and is not a limitation of the EM-Gen framework, *i.e.*, one can enumerate more cases in line with the same foundation.

5 Experiments

Datasets: We show our results on standard cross-validation splits of Breakfast [15] (1.7k, 3.6 mins/video, 48 actions), 50Salads [36] (50, 6.4 mins/video, 19

Table 1: **EM-TSS Posterior performance.** The obtained boundaries are highly accurate with high MoF and low boundary error (*‘Error %’*).

Method	50Salads		Breakfast	
	Error(%)	MoF	Error(%)	MoF
Uniform	1.23	76.3	4.69	69.2
EM-TSS	0.55	89.6	2.19	87.8

actions) and GTEA [9] (28, 1.5 mins/video, 11 actions). The standard evaluation criteria are the Mean-over-Frames (MoF), segment-wise edit distance (Edit), and $F1$ -scores with IoU thresholds of 0.25 and 0.50 ($F1@{25, 50}$) (see [21]).

Implementation Details: We compare using the same modified MS-TCN as [22] as our base network. We report results on additional base networks in the *Supplementary*. To initialize $\Theta^{(0)}$, we train the base TCN for $\{30, 50, 50\}$ epochs (N^{init}) for Breakfast, 50Salads, and GTEA, respectively. Subsequently, we apply $\{10, 20, 20\}$ E-M iterations (M), fixing each M-step to 5 epochs (N^{max}) for all datasets to update $\Theta^{(m)}$. We use the Adam optimizer [14] during initialization; for each M-step we use a learning rate $5e-4$ and a batch-size of 8.

Random Seeds & Variation in Annotations: All our results, including baselines, are reported for three seeds used for random selection of the annotated frames (timestamps) for both modes of supervision. Since [22] reports a single run, we report the results from that same set of timestamps in Table 3.

Baselines: We compare with two baselines: a *‘Naive’* baseline trained on *only* the annotated frames (see initialization in Sec. 4.2) and a *‘Uniform’* baseline (named according to [22], applicable only to TSS) that fixes the action boundary s_k to the mid-point between consecutive timestamps t_{k-1}, t_k . The action classes of unlabelled frames are then set accordingly and used during training.

5.1 Timestamp Supervision Results

Evaluation of Formulated Posterior. Table 1 compares the performance of our boundaries estimated from EM-TSS posterior probabilities against *‘Uniform’* (boundary at mid-point) for all training data. The action boundary locations s_k (unknown under Timestamp Supervision) are obtained by taking the expectation with respect to the posterior probabilities in Eq. 7. We evaluate *duration normalized* boundary error (Error (%)) as $|\hat{s}_k - s_k^{gt}|/T$, where \hat{s}_k and s_k^{gt} are the expected and ground-truth boundary locations, respectively, and T is the video duration. We also evaluate the frame-wise posterior weights $\omega_{j,c}^{(M)}$ (obtained after the final E-step) from Eq. 8 by regarding them as probability predictions and computing the frame-wise accuracy (MoF). Our boundary errors are less than half that of the uniform baseline; for 50Salads, our EM-TSS’s boundary error is less than 1%, while the MoF reaches 87% or higher for Breakfast and 50Salads.

Comparison to Baseline and Full Supervision. Table 2 compares our EM-TSS with the *‘Naive’* and *‘Uniform’* baselines (see section 5) and a fully-supervised setup. Our EM-TSS outperforms all baselines and even the fully-

Table 2: **Timestamp Supervision.** Our EM-TSS exceeds Fully-supervised (100% labels) on several metrics on 50Salads and GTEA. On Breakfast and 50Salads, EM-TSS has only 0.9% and 0.5% less MoF than Fully-supervised. We report the average results over three randomly sampled sets of annotations.

	50Salads				Breakfast				GTEA			
Method	$F1\{25,50\}$	Edit	MoF		$F1\{25,50\}$	Edit	MoF		$F1\{25,50\}$	Edit	MoF	
Naive	39.4	29.3	34.2	69.9	27.6	19.5	35.4	58.0	52.4	37.5	50.0	55.3
Uniform	58.2	42.3	60.4	63.4	56.3	36.4	68.1	51.0	72.5	50.9	73.1	56.5
EM-TSS	75.4	63.7	70.9	77.3	63.0	49.9	67.0	67.1	81.3	66.4	81.2	68.6
Fully-supervised	67.7	58.6	63.8	77.8	64.2	51.5	69.4	68.0	82.7	69.6	79.6	76.1

Table 3: **Comparison with SOTA:** In Timestamp Supervision (TSS), we exceed the previous SOTA by a considerable margin for almost all metrics. Our TSS results are reported for the same timestamp set as [22,24].

	50Salads				Breakfast				GTEA			
Method	$F1\{25,50\}$	Edit	MoF		$F1\{25,50\}$	Edit	MoF		$F1\{25,50\}$	Edit	MoF	
Plateau [24]	68.2	56.1	62.6	73.9	59.1	43.2	65.9	63.5	68.0	43.6	72.3	52.9
Li et al. [22]	70.9	60.1	66.8	75.6	63.6	47.4	69.9	64.1	73.0	55.4	72.3	66.4
Our EM	75.9	64.7	71.6	77.9	63.7	49.8	67.2	67.0	82.7	66.5	82.3	70.5
<i>(Inc. SOTA)</i>	+5.0	+4.6	+4.8	+2.3	+0.1	+2.4	-2.7	+2.9	+9.7	+11.1	+10	+4.1
Fully-supervised	67.7	58.6	63.8	77.8	64.2	51.5	69.4	68.0	82.7	69.6	79.6	76.1

supervised method in several metrics for the 50Salads and GTEA datasets. Most notable is the +7% difference in the $F1\{50,25\}$ and Edit scores in 50Salads.

We discuss possible reasons behind EM-TSS exceeding full-supervision in the *Supplementary*, and in subsection 5.4, where the (negative) effect of boundary ambiguity is shown to be more prominent in 50Salads and GTEA than Breakfast.

SOTA comparison. In table 3, we compare our EM-TSS with [24,22]; we use the *same* set of annotated timestamps as [24,22]. Our method outperforms the previous SOTA for most metrics on all datasets, with margins of +11%, +10%, +5% in GTEA $F1@50$, GTEA Edit, and 50Salads $F1@25$, respectively.

5.2 Performance With Missed Segments

In figure 4, we show that existing TSS methods suffer from poor performance with missing segments, with a sharp drop in performance at 5% and at 20% missing segments. While EM-TSS fares better, it does not exhibit the robustness of EM-Gen, as EM-Gen explicitly models missing segments between two consecutive timestamps. We detail the numerical performance in the *Supplementary*.

5.3 SkipTag Supervision Results

For SkipTag annotation, we annotate K number of timestamps per video. We set K to be the average number of action segments, i.e., $K = \{19, 6, 32\}$ for 50Salads,




Fig. 4: **Performance under (% of) missing action segments.** Our EM-TSS performs significantly better than [22] under missing action segments (0% denotes no missed segment). Our EM-Gen is robust to missing segments, and with 20% missing segments, EM-Gen significantly outperforms TSS methods.

Table 4: **SkipTag Supervision.** Our EM-Gen provides a significant boost over ‘Naive’ and [22]’s baselines. It performs close to full supervision in 50Salads, and outperforms it in few metrics. Results are averaged over three sets of annotations. (*) indicates the results reported in [22].

	50Salads ($K = 19$)				Breakfast ($K = 6$)				GTEA ($K = 32$)			
Method	$F1\{25, 50\}$	Edit	MoF		$F1\{25, 50\}$	Edit	MoF		$F1\{25, 50\}$	Edit	MoF	
Naive	49.2	38.1	46.3	70.4	27.0	18.8	36.4	59.8	64.2	43.4	65.8	63.1
Li et al. [22]	54.3	40.1	54.6	70.7	-	-	-	*61.7	69.4	50.0	68.9	66.0
<i>EM-Gen</i>	68.1	54.9	64.3	74.4	57.3	45.2	59.9	64.1	76.7	57.9	73.5	69.8
Fully-supervised	67.7	58.6	63.8	77.8	64.2	51.5	69.4	68.0	82.7	69.6	79.6	76.1

Breakfast and GTEA, respectively. Having K random frames per video results in approximately the same number of annotated frames as Timestamp Supervision, but is much easier for the annotator. Table 4 compares the performance of previous methods to EM-Gen under SkipTag Supervision and shows that EM-Gen surpasses [22] across all metrics and datasets. Moreover, in 50Salads, EM-Gen is comparable to and sometimes surpasses *Fully-supervised* (i.e., 100% labels).

5.4 Additional Results

In *Supplementary*, we provide ablation related to **(1) different priors** (see subsection 4.4) (section S1.1), **(2) loss-functions** (section S1.2), **(3) TCN architectures** (section S1.4), **(4) effect of timestamp position** (section S1.3), other hyper-parameters, and **(5) qualitative analysis** (section S2).

Annotation Efficiency User Study. We perform a user study with 5 annotators labelling 50 videos of the Breakfast dataset with different types of annotations (Full, Timestamp, and SkipTag). The left sub-table of table 5 reports the duration-normalized annotation time, i.e., the annotation time required as a fraction of the video duration. Timestamp Supervision requires about two-thirds of the annotation time as full supervision, while SkipTag Supervision requires only one-third. Other details of the user study are given in the *Supplementary*.

Table 5: **(left)** Labelling cost on Breakfast. Time(%) denotes annotation cost as a fraction of video duration. **(right)** Choosing TSS timestamps near boundaries (start or end) lowers ‘*Naive*’ initialization performance significantly (sometimes $\geq -30\%$), indicating ambiguity in labelling the boundary frames.

Cost	Full-sup	Timestamp	SkipTag	Frame Type	F1@{25,50}	Edit	MoF
Frames (%)	100	0.03	0.03	Random	39.4	29.3	34.2 69.9
Time (%)	93.3	65.1	36.4	Mid	47.7	34.9	40.9 69.7
				Start	12.5	3.2	17.5 31.9
				End	12.2	2.7	21.4 29.1

Ambiguity in Boundary Frames. The right sub-table of table 5 compares the ‘*Naive*’ baseline on the 50Salads dataset when TSS timestamps are selected from different locations of the action segment. There is a $\geq 20\%$ difference in scores when trained with boundary frames (start/end frame) versus frames selected at (*Random*) or from the middle (*Mid*) of the action segment. This indicates ambiguity in boundary frame annotation (also see [25,1]). We further explain in the *Supplementary*, how EM-TSS exceeds fully-supervised results.

5.5 Training Complexity

The M-step of our E-M, transformed into the traditional cross-entropy loss, is extremely efficient and takes the same per-epoch time as fully-supervised. We only update our weights at the E-step once every (N^{max}) 5 epochs (unlike the updates at every iteration required in [22,24]). Our EM-TSS training time is one magnitude less than [22]. Our initialization time is exactly same as [22], whereas post-initialization, on one RTX 2080 GPU, 10 epochs (50Salads) of our EM-TSS require only 5 mins vs 50 mins for [22] (official code) under the same resources.

6 Conclusion

We first tackle the problem of timestamp supervision by introducing a novel E-M algorithm based method. Our proposed framework is general and robust enough to work with timestamp supervision, even with mistakenly missed action segments. We also propose SkipTag supervision, a novel and much weaker form of annotation than TSS. Our method seamlessly applies to all the scenarios, producing *SOTA* results, and in some cases surpassing fully-supervised results.

Acknowledgment

This research is supported by the National Research Foundation, Singapore under its NRF Fellowship for AI (NRF-NRFFAI1-2019-0001). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

1. Alwassel, H., Heilbron, F.C., Escorcía, V., Ghanem, B.: Diagnosing error in temporal action detectors. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 256–272 (2018)
2. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
3. Chang, C.Y., Huang, D.A., Sui, Y., Fei-Fei, L., Niebles, J.C.: 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. pp. 3546–3555 (2019)
4. Chen, M.H., Li, B., Bao, Y., AlRegib, G., Kira, Z.: Action segmentation with joint self-supervised temporal domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9454–9463 (2020)
5. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* **39**(1), 1–38 (1977), <http://www.jstor.org/stable/2984875>
6. Ding, G., Yao, A.: Temporal action segmentation with high-level complex activity labels. arXiv preprint arXiv:2108.06706 (2021)
7. Ding, L., Xu, C.: Weakly-supervised action segmentation with iterative soft boundary assignment. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6508–6516 (2018)
8. Farha, Y.A., Gall, J.: Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3575–3584 (2019)
9. Fathi, A., Ren, X., Rehg, J.M.: Learning to recognize objects in egocentric activities. In: CVPR 2011. pp. 3281–3288. IEEE (2011)
10. Fayyaz, M., Gall, J.: Sct: Set constrained temporal transformer for set supervised action segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
11. Gao, S.H., Han, Q., Li, Z.Y., Peng, P., Wang, L., Cheng, M.M.: Global2local: Efficient structure search for video action segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16805–16814 (2021)
12. Ishikawa, Y., Kasai, S., Aoki, Y., Kataoka, H.: Alleviating over-segmentation errors by detecting action boundaries. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2322–2331 (2021)
13. Ju, C., Zhao, P., Chen, S., Zhang, Y., Wang, Y., Tian, Q.: Divide and conquer for single-frame temporal action localization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13455–13464 (2021)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv:1412.6980 (2014)
15. Kuehne, H., Arslan, A., Serre, T.: The language of actions: Recovering the syntax and semantics of goal-directed human activities. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 780–787 (2014)
16. Kuehne, H., Richard, A., Gall, J.: A hybrid rnn-hmm approach for weakly supervised temporal action segmentation. *IEEE transactions on pattern analysis and machine intelligence* **42**(4), 765–779 (2018)

17. Kukleva, A., Kuehne, H., Sener, F., Gall, J.: Unsupervised learning of action classes with continuous temporal embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12066–12074 (2019)
18. Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 156–165 (2017)
19. Li, J., Lei, P., Todorovic, S.: Weakly supervised energy-based learning for action segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6243–6251 (2019)
20. Li, J., Todorovic, S.: Set-constrained viterbi for set-supervised action segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
21. Li, S.J., AbuFarha, Y., Liu, Y., Cheng, M.M., Gall, J.: Ms-tcn++: Multi-stage temporal convolutional network for action segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
22. Li, Z., Abu Farha, Y., Gall, J.: Temporal action segmentation from timestamp supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8365–8374 (2021)
23. Ma, F., Zhu, L., Yang, Y., Zha, S., Kundu, G., Feiszli, M., Shou, Z.: Sf-net: Single-frame supervision for temporal action localization. In: European conference on computer vision. pp. 420–437. Springer (2020)
24. Moltisanti, D., Fidler, S., Damen, D.: Action recognition from single timestamp supervision in untrimmed videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9915–9924 (2019)
25. Moltisanti, D., Wray, M., Mayol-Cuevas, W., Damen, D.: Trespassing the boundaries: Labeling temporal bounds for object interactions in egocentric video. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2886–2894 (2017)
26. Rashid, M., Kjellstrom, H., Lee, Y.J.: Action graphs: Weakly-supervised action localization with graph convolution networks. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (March 2020)
27. Richard, A., Gall, J.: Temporal action detection using a statistical language model. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3131–3140 (2016)
28. Richard, A., Kuehne, H., Gall, J.: Action sets: Weakly supervised action segmentation without ordering constraints. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 5987–5996 (2018)
29. Richard, A., Kuehne, H., Iqbal, A., Gall, J.: Neuralnetwork-viterbi: A framework for weakly supervised video learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7386–7395 (2018)
30. Sener, F., Singhania, D., Yao, A.: Temporal aggregate representations for long-range video understanding. In: European Conference on Computer Vision. pp. 154–171. Springer (2020)
31. Sener, F., Yao, A.: Unsupervised learning and segmentation of complex activities from video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8368–8376 (2018)
32. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: Advances in neural information processing systems. pp. 568–576 (2014)

33. Singh, B., Marks, T.K., Jones, M., Tuzel, O., Shao, M.: A multi-stream bi-directional recurrent neural network for fine-grained action detection. pp. 1961–1970 (2016)
34. Singhania, D., Rahaman, R., Yao, A.: Coarse to fine multi-resolution temporal convolutional network. arXiv preprint arXiv:2105.10859 (2021)
35. Souri, Y., Fayyaz, M., Minciullo, L., Francesca, G., Gall, J.: Fast weakly supervised action segmentation using mutual consistency. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
36. Stein, S., McKenna, S.J.: 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*. pp. 729–738. ACM (2013)
37. VidalMata, R.G., Scheirer, W.J., Kukleva, A., Cox, D., Kuehne, H.: Joint visual-temporal embedding for unsupervised learning of actions in untrimmed sequences. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 1238–1247 (2021)
38. Wang, D., Yuan, Y., Wang, Q.: Gated forward refinement network for action segmentation. *Neurocomputing* **407**, 63–71 (2020)
39. Wang, Z., Gao, Z., Wang, L., Li, Z., Wu, G.: Boundary-aware cascade networks for temporal action segmentation. In: *European Conference on Computer Vision*. pp. 34–51. Springer (2020)