

Learning Action Completeness from Points for Weakly-supervised Temporal Action Localization

Pilhyeon Lee¹ Hyeran Byun^{1,2*}

¹Department of Computer Science, Yonsei University

²Graduate school of AI, Yonsei University

{lph1114, hrbyun}@yonsei.ac.kr

Abstract

We tackle the problem of localizing temporal intervals of actions with only a single frame label for each action instance for training. Owing to label sparsity, existing work fails to learn action completeness, resulting in fragmentary action predictions. In this paper, we propose a novel framework, where dense pseudo-labels are generated to provide completeness guidance for the model. Concretely, we first select pseudo background points to supplement point-level action labels. Then, by taking the points as seeds, we search for the optimal sequence that is likely to contain complete action instances while agreeing with the seeds. To learn completeness from the obtained sequence, we introduce two novel losses that contrast action instances with background ones in terms of action score and feature similarity, respectively. Experimental results demonstrate that our completeness guidance indeed helps the model to locate complete action instances, leading to large performance gains especially under high IoU thresholds. Moreover, we demonstrate the superiority of our method over existing state-of-the-art methods on four benchmarks: THUMOS'14, GTEA, BEOID, and ActivityNet. Notably, our method even performs comparably to recent fully-supervised methods, at the 6× cheaper annotation cost. Our code is available at <https://github.com/Pilhyeon>.

1. Introduction

The goal of temporal action localization lies in locating starting and ending timestamps of action instances and classifying them. Thanks to the various applications [33, 45, 48], it has drawn much attention from researchers, leading to the rapid and remarkable progress in the fully-supervised setting (*i.e.*, frame-level labels) [27, 41, 43, 50]. Meanwhile, there appear attempts to reduce the prohibitively expensive cost of annotating individual frames by devising weakly-supervised models with video-level labels [7, 32, 46, 56].

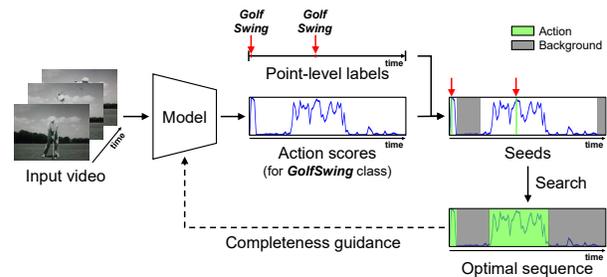


Figure 1: Simplified illustration of our idea. We use points as seeds to find the optimal sequence, which in turn provides completeness guidance to the model.

However, they fall largely behind the fully-supervised counterparts, mainly on account of their weak ability to distinguish action and background frames [17, 18, 38, 52].

To narrow the performance gap between them, another level of weak supervision has been proposed recently, namely the point-supervised setting. In this setting, only a single timestamp (point) with its action category is annotated for each action instance during training. In terms of the labeling cost, point-level labels require a negligible extra cost compared to video-level ones, while being 6× cheaper than frame-level ones (50s vs. 300s per 1-min video) [31].

Despite the affordable cost, it offers coarse locations as well as the total number of action instances, thus bringing a strong ability in spotting actions to the models. Consequently, point-supervised methods show comparable or even superior performances to fully-supervised counterparts under low intersection over union (IoU) thresholds. However, it has been revealed that they suffer from incomplete predictions, resulting in highly inferior performances in the case of high IoU thresholds. We conjecture that this problem is attributed to the sparse nature of point-level labels that induces the models to learn only a small part of actions rather than the full extent of action instances. In other words, they fail to learn *action completeness* from the point annotations. Although SF-Net [31] mines pseudo action and background points to alleviate the label sparsity, they are discontinuous and thus do not provide completeness cues.

*Corresponding author

In this paper, we aim to allow the model to learn action completeness under the point-supervised setting. To this end, we introduce a new framework, where dense pseudo-labels (*i.e.*, sequences) are generated based on the point annotations to provide completeness guidance to the model. The overall workflow is illustrated in Fig. 1.

Technically, we first select pseudo background points to augment point-level action labels. As aforementioned, such point annotations are discontinuous, so it is infeasible to learn completeness from them. To that end, we propose to search for the optimal sequence covering complete action instances among candidates consistent with the point labels. However, it is non-trivial to measure how complete the instances in each candidate sequence are, without full supervision. To realize it, we borrow the outer-inner-contrast concept [42] as a proxy for instance completeness. Intuitively, a complete action instance generally shows large score contrast, *i.e.*, much higher action scores for inner frames than those for surrounding frames. In contrast, a fragmentary instance probably has high action scores in its outer region (still within the action), leading to small score contrast. This can be generalized for background instances as well. Based on this property, we derive the score of an input sequence by aggregating the score contrast of action and background instances constituting the sequence. By maximizing the score, we can obtain the optimal sequence that is likely to be well-aligned with the ground-truth we do not have. In experiments, we present the accuracy of optimal sequences and the correlation between score contrast and completeness.

From the obtained sequence, the model is supposed to learn action completeness. To this end, we design score contrastive loss to maximize the agreement between the model outputs and the optimal sequence, by enlarging the completeness of the sequence. With the loss, the model is trained to discriminate each action (background) instance from its surroundings in terms of action scores. Moreover, we introduce feature contrastive loss to encourage feature discrepancy between action and background instances. Experiments validate that the proposed losses complementarily help the model to detect complete action instances, leading to large performance gains under high IoU thresholds.

To summarize, our contributions are three-fold.

- We introduce a new framework, where the dense optimal sequence is generated to provide completeness guidance to the model in the point-supervised setting.
- We propose two novel losses that facilitate the action completeness learning by contrasting action instances with background ones with respect to action score and feature similarity, respectively.
- Our model achieves a new state-of-the-art with a large gap on four benchmarks. Furthermore, it even performs favorably against fully-supervised approaches.

2. Related Work

Fully-supervised temporal action localization. In order to tackle temporal action localization, fully-supervised methods rely on precise temporal annotations, *i.e.*, **frame-level labels**. They mainly adopt the two-stage paradigm (proposal generation and classification), and can be roughly categorized into two groups regarding the way to generate proposals. The first group prepares a large number of proposals using the sliding window technique [4, 41, 43, 49, 53, 55, 62]. On the other hand, the second group first predicts the probability of each frame being a start (end) point of an action instance, and then uses the combinations of probable start and end points as proposals [21, 22, 23, 61]. Meanwhile, there are graph modeling methods taking snippets [1, 51] or proposals [57] as nodes. Different from fully-supervised methods that utilize expensive frame-level labels for action completeness learning, our method enables it with only point-level labels by introducing a novel framework.

Weakly-supervised temporal action localization. To alleviate the cost issue of frame-level labels, many attempts have been made recently to solve the same task in the weakly-supervised setting, mainly using **video-level labels**. Untrimmednets [46] tackle it by selecting segments that contribute to video-level classification. STPN [37] puts a constraint that key frames should be sparse. In addition, there are background modeling approaches under the video-supervised setting [9, 17, 18, 38]. To learn reliable attention weights, DGAM [40] designs a generative modeling, while EM-MIL [30] adopts the Expectation-maximization strategy. Meanwhile, metric learning is utilized for action representation learning [10, 36, 39] or action-background separation [34]. There are also methods that explore sub-actions [11, 29] or exploit the complementarity of RGB and flow modalities [54, 58]. Besides, several methods leverage external information, *e.g.*, action count [36, 52], pose [60] or audio [16]. Moreover, some approaches aim to detect complete action instances by aggregating multiple predictions [25], erasing the most discriminative part [44, 63], or directly regressing the action intervals [28, 42].

Most recently, **point-level supervision** starts to be explored, which provides rich information at an affordable cost. Moltisanti *et al.* [35] first utilize the point-level labels for action localization. SF-Net [31] adopts the pseudo label mining strategy to acquire more labeled frames. Meanwhile, Ju *et al.* [13] perform boundary regression based on key frame prediction. However, they do not explicitly consider action completeness, and therefore produce predictions that cover only part of action instances. In contrast, we propose to learn action completeness from dense pseudo-labels by contrasting action instances with surrounding background ones. In Sec. 4, the efficacy of our method is clearly verified with notable performance boosts at high IoU thresholds.

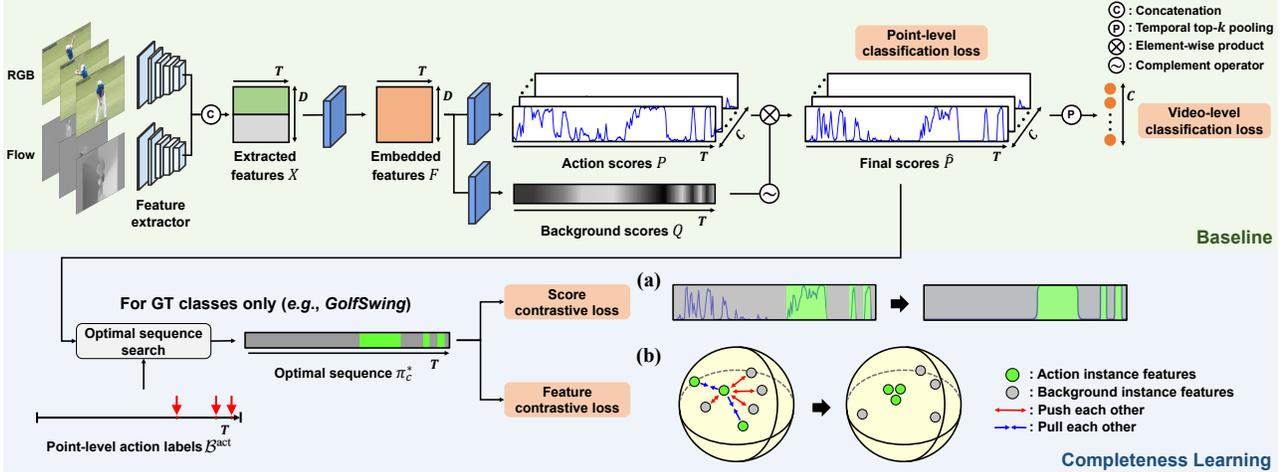


Figure 2: Overview of the proposed method. Besides the conventional objectives, *i.e.*, video-level and point-level classification losses, we propose to learn action completeness (the lower part). Based on the final action scores, the optimal sequence is selected among candidates consistent with the point-level labels. It in turn provides completeness guidance with two proposed losses that contrast action instances with background ones with respect to (a) action score and (b) feature similarity.

3. Method

In this section, we first describe the problem setting and detail the baseline setup. Afterward, the optimal sequence search is elaborated, followed by our action completeness learning strategy. Lastly, we explain the joint learning and the inference of our model. The overall architecture of our method is illustrated in Fig. 2.

Problem setting. Following [13, 31], we set up the problem of point-supervised temporal action localization. Given an input video, a single point and the category for each action instance is provided, *i.e.*, $\mathcal{B}^{\text{act}} = \{(t_i, \mathbf{y}_{t_i})\}_{i=1}^{M^{\text{act}}}$, where the i -th action instance is labeled at the t_i -th segment (frame) with its action label \mathbf{y}_{t_i} , and M^{act} is the total number of action instances in the input video. The points are sorted in temporal order (*i.e.*, $t_i < t_{i+1}$). The label \mathbf{y}_{t_i} is a binary vector with $y_{t_i}[c] = 1$ if the i -th action instance contains the c -th action class and otherwise 0 for C action classes. It is worth noting that the video-level label \mathbf{y}^{vid} can be readily acquired by aggregating the point-level ones, *i.e.*, $y^{\text{vid}}[c] = \mathbb{1} \left[\sum_{i=1}^{M^{\text{act}}} y_{t_i}[c] > 0 \right]$, where $\mathbb{1}[\cdot]$ is the indicator function.

3.1. Baseline Setup

Our baseline is shown in the upper part of Fig. 2. We first divide the input video into 16-frame segments, which are then fed to the pre-trained feature extractor. Following [17, 39], we exploit both of RGB and flow streams with early-fusion. The two-stream features are fused by concatenation, resulting in $X \in \mathbb{R}^{D \times T}$, where D and T denote the feature dimension and the number of segments, respectively.

The extracted features then go through a single 1D convolutional layer followed by ReLU activation, which pro-

duces the embedded features F . In practice, we set the dimension of the embedded features to the same as that of the extracted features X , *i.e.*, $F \in \mathbb{R}^{D \times T}$. Afterward, the embedded features are fed into a 1D convolutional layer with the sigmoid function, to predict the segment-level class scores $P \in \mathbb{R}^{C \times T}$, where C indicates the number of action classes. Meanwhile, we derive the class-agnostic background scores $Q \in \mathbb{R}^T$, to model background frames which do not belong to any action classes. Thereafter, we fuse the action scores with the complement of background probability to get the final scores \hat{P} , *i.e.*, $\hat{p}_t[c] = p_t[c](1 - q_t)$. This fusion strategy is similar to that of [18], although the out-of-distribution modeling is not incorporated in our model.

The segment-level action scores are then aggregated to build a single video-level class score. We use the temporal top- k pooling for aggregation as in [17, 39]. Formally, the video-level probability is calculated as follows.

$$\hat{p}^{\text{vid}}[c] = \frac{1}{k} \max_{S \subset \hat{P}[c, :]} \sum_{m \in S} m, \quad (1)$$

where $k = \lfloor \frac{T}{8} \rfloor$ and S denotes all possible subsets of $\hat{P}[c, :]$ containing k segments, *i.e.*, $|S| = k$.

Our baseline model includes two loss functions using video- and point-level labels respectively. As aforementioned, the video-level class label $y^{\text{vid}}[c]$ can be derived by accumulating the point-level labels. The video-level classification loss is then calculated with binary cross-entropy.

$$\mathcal{L}_{\text{video}} = - \sum_{c=1}^C \left(y^{\text{vid}}[c] \log \hat{p}^{\text{vid}}[c] + (1 - y^{\text{vid}}[c]) \log (1 - \hat{p}^{\text{vid}}[c]) \right). \quad (2)$$

The point-level classification loss is also computed by binary cross-entropy but involving the background term for effectively training Q . In addition, we adopt the focal loss [24] to facilitate the training process. Formally, the classification loss for action points is defined as follows.

$$\mathcal{L}_{\text{point}}^{\text{act}} = -\frac{1}{M^{\text{act}}} \sum_{\forall (t, \mathbf{y}_t) \in \mathcal{B}^{\text{act}}} \left(\sum_{c=1}^C \left(y_t[c] (1 - \hat{p}_t[c])^\beta \log \hat{p}_t[c] + (1 - y_t[c]) \hat{p}_t[c]^\beta \log (1 - \hat{p}_t[c]) \right) + q_t^\beta \log (1 - q_t) \right), \quad (3)$$

where M^{act} indicates the number of action instances in the video and β is the focusing parameter, which is set to 2 following the original paper [24].

Training only with action points would lead the network to always produce low background scores rather than learn to separate action and background. Therefore, we gather some pseudo background points to supplement action ones. Our principle for selection is that at least one background frame must be placed between two adjacent action instances to separate them. By the problem definition, two different action points are sampled from different instances, so we use the action points as surrogates for the corresponding instances. Concretely, between two adjacent action points, we find the segments whose background scores q_t are larger than the threshold γ . If no segment satisfies the condition in a section, we select one with the largest background score. Meanwhile, for the case where multiple background points are selected in a section, we mark all points between them as background, since it is trivial that no action exists there. In practice, this strategy is shown to be more effective than global mining [31] by collecting more hard points. Given the pseudo background point set, $\mathcal{B}^{\text{bkg}} = \{t_j\}_{j=1}^{M^{\text{bkg}}}$, the classification loss for background points is computed by:

$$\mathcal{L}_{\text{point}}^{\text{bkg}} = -\frac{1}{M^{\text{bkg}}} \sum_{\forall t \in \mathcal{B}^{\text{bkg}}} \left(\sum_{c=1}^C \hat{p}_t[c]^\beta \log (1 - \hat{p}_t[c]) + (1 - q_t)^\beta \log q_t \right), \quad (4)$$

where M^{bkg} denotes the number of the selected background points and β is the focusing factor, the same with (3). For pseudo background points, we penalize the final scores for all action classes, while encouraging the background scores.

The total point-level loss function is defined as the sum of the losses for action and pseudo background points.

$$\mathcal{L}_{\text{point}} = \mathcal{L}_{\text{point}}^{\text{act}} + \mathcal{L}_{\text{point}}^{\text{bkg}}. \quad (5)$$

3.2. Optimal Sequence Search

As discussed in Sec. 1, the point-level classification loss is insufficient to learn action completeness, as point labels cover only a small portion of action instances. Therefore, we propose to generate dense pseudo-labels that can offer

some hints about action completeness for the model. In detail, we consider all possible sequence candidates consistent with the action and pseudo background points. Among them, we find the optimal sequence that can provide good completeness guidance to the model. However, it is non-trivial without full supervision to measure how well a candidate sequence covers complete action instances. To enable it, we re-purpose the outer-inner-contrast concept [42] as a proxy for judging the completeness score of a sequence. Intuitively, the contrast between inner and outer scores is likely to be large for a complete action instance but small for a fragmentary one. Note that our purpose is different from the original paper [42]. It was originally designed for parametric boundary regression. In contrast, we utilize it as a scoring function to search for the optimal sequence, from which the model could learn action completeness.

Before detailing the scoring function, we present the formulation of candidate sequences. Due to the multi-label nature of temporal action localization, we consider class-specific sequences for each action class. Note that all segments belonging to other action classes are considered background for sequences of class c . Then, a sequence is defined as multiple action and background (including other actions) instances that alternate consecutively. Formally, a sequence of class c can be expressed as $\pi_c = \{(s_n^c, e_n^c, z_n^c)\}_{n=1}^{N_c}$, where s_n^c and e_n^c denote the start and end points of the n -th instance, respectively, while N_c is the total number of instances for class c . In addition, $z_n^c \in \{0, 1\}$ indicates the type of the instance, *i.e.*, $z_n^c = 1$ if n -th instance is of the c -th action class, otherwise 0 (background).

Given an input sequence, we compute its completeness score by averaging the contrast scores of individual action and background instances contained in the sequence. It would be noted that the contrast scores of background instances are included in the calculation, which proves to be effective for finding more accurate optimal sequences, as will be shown in Sec. 4.3. Formally, the completeness score of a sequence π_c for the c -th action class is computed by:

$$\mathcal{R}(\pi_c) = \frac{1}{N_c} \sum_{n=1}^{N_c} \left(\underbrace{\frac{1}{l_n^c} \sum_{t=s_n^c}^{e_n^c} u_n^c(t)}_{\text{Inner score}} - \underbrace{\frac{1}{\lceil \delta l_n^c \rceil + \lfloor \delta l_n^c \rfloor} \left(\sum_{t=s_n^c - \lfloor \delta l_n^c \rfloor}^{s_n^c - 1} u_n^c(t) + \sum_{t=e_n^c + 1}^{e_n^c + \lceil \delta l_n^c \rceil} u_n^c(t) \right)}_{\text{Outer score}} \right), \quad (6)$$

where $u_n^c(t) = \begin{cases} \hat{p}_t[c], & \text{if } z_n^c = 1. \\ 1 - \hat{p}_t[c], & \text{otherwise.} \end{cases}$

$l_n^c = e_n^c - s_n^c + 1$ is the temporal length of the n -th instance of π_c , δ is a hyper-parameter adjusting the outer range (set to 0.25), and N_c is the total number of action and background instances for class c . Then, the optimal sequence for

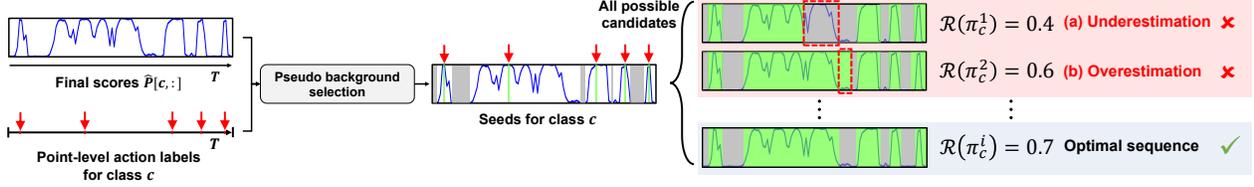


Figure 3: Optimal sequence search for class c . Given the final scores and the point-level labels, we select pseudo background points. Then, among all possible candidates, we search for the optimal sequence that maximizes the completeness score (6).

class c can be obtained by finding the sequence that maximizes the score, *i.e.*, $\pi_c^* = \arg \max_{\pi_c} \mathcal{R}(\pi_c)$ using (6). The optimal sequence search process is illustrated in Fig. 3. By evaluating the completeness score, our method can reject underestimation (Fig. 3a) and overestimation (Fig. 3b) cases. Consequently, we obtain the optimal sequence that is most likely to contain complete action instances.

However, the search space grows exponentially as T increases, leading to the exorbitant cost for optimal sequence search. To relieve the issue, we implement the search process with a greedy algorithm under a limited budget, which results in greatly saving the computational cost. Detailed algorithm and cost analysis are presented in Sec. B of the appendix. Note that optimal sequence search is performed only for the action classes contained in the video.

3.3. Action Completeness Learning

Given the class-specific optimal sequences $\{\pi_c^*\}_{c=1}^C$, our goal is to let the model learn action completeness. To this end, we design two losses that enable completeness learning by contrasting action instances from background ones. This helps in complete action predictions, as validated in Sec. 4.

Firstly, we propose *score contrastive loss* that encourages the model to separate action (background) instances from their surroundings in terms of final scores. It can be also interpreted as fitting the model outputs to the optimal sequences (Fig. 2a). Formally, the loss is computed by:

$$\mathcal{L}_{\text{score}} = \frac{1}{\sum_{c=1}^C y^{\text{vid}}[c]} \sum_{c=1}^C y^{\text{vid}}[c] (1 - \mathcal{R}(\pi_c^*))^\beta, \quad (7)$$

where we use β -squared term to focus on the instances that are largely inconsistent with the optimal sequence ($\beta = 2$).

Secondly, inspired by the recent success of contrastive learning [5, 8, 14], we design *feature contrastive loss*. Our intuition is that features from different instances but with the same action class should be closer to each other than any other background instances in the same video (Fig. 2b). We note that our loss differs from [5, 8, 14] in that they pull different views of an input image, whereas ours attracts different action instances in a given video. In addition, ours does not need negative sampling from different images, as background instances are obtained from the same video.

To extract the representative feature for each action (or background) instance, we modify the segment of interest (SOI) pooling [4] by replacing max-pooling with random sampling. In detail, we evenly divide each input instance into three intervals, from each of which a single segment is randomly sampled. Then, the embedded features of the sampled segments are averaged, producing the representative feature f_n^c for the n -th instance of the sequence π_c^* .

Taking the normalized instance features \bar{f}_n^c as inputs, we derive feature contrastive loss. The loss is computed only for the classes whose action counts are larger than 1, *i.e.*, at least two action instances exist in the video. Note that background instances do not attract each other. Given the optimal sequences $\{\pi_c^* = \{(s_n^c, e_n^c, z_n^c)\}_{n=1}^{N_c}\}_{c=1}^C$, the proposed feature contrastive loss is formulated as:

$$\mathcal{L}_{\text{feat}} = \frac{1}{\sum_{c=1}^C \mathbb{1} \left[\sum_{n=1}^{N_c} z_n^c > 1 \right]} \sum_{c=1}^C \mathbb{1} \left[\sum_{n=1}^{N_c} z_n^c > 1 \right] \ell_{\text{feat}}^c, \quad (8)$$

$$\ell_{\text{feat}}^c = - \frac{1}{\sum_{n=1}^{N_c} z_n^c} \sum_{n=1}^{N_c} z_n^c \log \frac{\sum_{\forall o \neq n} z_o^c \exp(\bar{f}_n^c \cdot \bar{f}_o^c / \tau)}{\sum_{\forall m \neq n} \exp(\bar{f}_n^c \cdot \bar{f}_m^c / \tau)},$$

where ℓ_{feat}^c is the partial loss for class c , τ denotes the temperature parameter, and $\mathbb{1}[\cdot]$ denotes the indicator function.

3.4. Joint Training and Inference

The overall training objective of our model is as follows.

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{video}} + \lambda_2 \mathcal{L}_{\text{point}} + \lambda_3 \mathcal{L}_{\text{score}} + \lambda_4 \mathcal{L}_{\text{feat}}, \quad (9)$$

where λ_* are weighting parameters for balancing the losses, which are determined empirically.

During the test time, we first threshold on the video score $\hat{\mathbf{p}}^{\text{vid}}$ with θ^{vid} to determine which action categories are to be localized. Then, only for the remaining classes, we threshold on the segment-level final scores $\hat{\mathbf{p}}_t$ with θ^{seg} to select candidate segments. Afterward, consecutive candidates are merged into a single proposal, which becomes a localization result. We set the confidence of each proposal to its outer-inner-contrast score, as in [17, 25]. To augment the proposal pool, we use multiple thresholds for θ^{seg} and perform non-maximum suppression (NMS) to remove overlapping proposals. Note that the optimal sequence search is not performed at test time, so does not affect the inference time.

Supervision	Method	mAP@IoU (%)							AVG	AVG
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	(0.1:0.5)	(0.3:0.7)
Frame-level (Full)	BMN [22]	-	-	56.0	47.4	38.8	29.7	20.5	-	38.5
	P-GCN [57]	69.5	67.8	63.6	57.8	49.1	-	-	61.6	-
	G-TAD [51]	-	-	54.5	47.6	40.2	30.8	23.4	-	39.3
	BC-GNN [1]	-	-	57.1	49.1	40.4	31.2	23.1	-	40.2
	Zhao <i>et al.</i> [61]	-	-	53.9	50.7	45.4	38.0	28.5	-	43.3
Video-level (Weak)	Lee <i>et al.</i> [18]	67.5	61.2	52.3	43.4	33.7	22.9	12.1	51.6	32.9
	CoLA [59]	66.2	59.5	51.5	41.9	32.2	22.0	13.1	50.3	32.1
	AUMN [29]	66.2	61.9	54.9	44.4	33.3	20.5	9.0	52.1	32.4
	TS-PCA [26]	67.6	61.1	53.4	43.4	34.3	24.7	13.7	52.0	33.9
	UGCT [54]	69.2	62.9	55.5	46.5	35.9	23.8	11.4	54.0	34.6
Point-level (Weak)	SF-Net [†] [31]	71.0	63.4	53.2	40.7	29.3	18.4	9.6	51.5	30.2
	Ju <i>et al.</i> [†] [13]	72.8	64.9	58.1	46.4	34.5	21.8	11.9	55.3	34.5
	Ours [†]	75.1	70.5	63.3	55.2	43.9	33.3	20.8	61.6	43.3
	Moltisanti <i>et al.</i> [‡] [35]	24.3	19.9	15.9	12.5	9.0	-	-	16.3	-
	SF-Net [‡] [31]	68.3	62.3	52.8	42.2	30.5	20.6	12.0	51.2	31.6
	Ju <i>et al.</i> [‡] [13]	72.3	64.7	58.2	47.1	35.9	23.0	12.8	55.6	35.4
	Ours [‡]	75.7	71.4	64.6	56.5	45.3	34.5	21.8	62.7	44.5

Table 1: State-of-the-art comparison on THUMOS’14. We also include the methods under video-level and frame-level supervision for reference. The average mAPs are computed under the IoU thresholds 0.1:0.5 and 0.3:0.7 with the step size of 0.1. While [†] indicates the use of manually annotated labels from [31], [‡] denotes the use of labels automatically generated in [35].

4. Experiments

4.1. Experimental Settings

Datasets. THUMOS’14 [12] is of 20 action classes with 200 and 213 untrimmed videos for validation and test, respectively. It is known to be challenging due to the diverse length and the frequent occurrence of action instances. Following the convention [37], we use the validation videos for training and test videos for test. GTEA [19] contains 28 videos of 7 fine-grained daily actions in the kitchen, among which 21 and 7 videos are utilized for training and test, respectively. BEOID [6] has 58 videos with a total of 30 action categories. We follow the data split provided by [31]. ActivityNet [2] is a large-scale dataset with two versions. The version 1.3 includes 10,024 training, 4,926 validation, and 5,044 test videos with 200 action classes. The version 1.2 consists of 4,819 training, 2,383 validation, and 2,480 test videos with 100 categories. We evaluate our model on the validation sets for both versions. It should be noted that our model takes only point-level annotations for training.

Evaluation metrics. Following the standard protocol of temporal action localization, we compute mean average precisions (mAPs) under several different levels of intersection over union (IoU) thresholds. We note that performances at small IoU thresholds demonstrate the ability in finding actions, while those under high IoU thresholds exhibit the completeness of action predictions.

Implementation details. We employ the two-stream I3D networks [3] pre-trained on Kinetics-400 [3] as our feature

extractor, which is not fine-tuned in our experiments for fair comparison. To obtain optical flow maps, we use TV-L1 algorithm [47]. Each video is split into 16-frame segments, which are taken as inputs by the feature extractor resulting in 1024-dim features for each modality (*i.e.*, $D = 2048$). We use the original number of segments as T without sampling. Our model is optimized by Adam [15] with the learning rate of 10^{-4} and the batch size of 16. Hyper-parameters are determined by grid search: $\gamma = 0.95$, $\tau = 0.1$. The video-level threshold θ^{vid} is set to 0.5, while the segment-level threshold θ^{seg} spans from 0 to 0.25 with a step size of 0.05. The NMS is performed with the threshold of 0.6.

4.2. Comparison with State-of-the-art Methods

In Table 1, we compare our method with state-of-the-art models under different levels of supervision on THUMOS’14. We note that fully-supervised models require far more expensive annotation costs compared to weakly-supervised counterparts. In the comparison, our model significantly outperforms the state-of-the-art point-supervised approaches. We also notice the large performance margins at high IoU thresholds, *e.g.*, $\sim 11\%$ in mAP@0.6 and $\sim 9\%$ in mAP@0.7. This confirms that the proposed method aids in locating the complete action instances. At the same time, our model largely surpasses the video-supervised methods with the comparable labeling cost. Further, our model even performs favorably against the fully-supervised methods in terms of average mAPs at the much lower annotation cost. It is, however, also shown that ours lags behind them at high IoU thresholds, due to the lack of boundary information.

Dataset	Method	mAP@IoU (%)				AVG
		0.1	0.3	0.5	0.7	
GTEA	SF-Net [31]	58.0	37.9	19.3	11.9	31.0
	SF-Net* [31]	52.9	37.6	21.7	13.7	31.1
	Ju <i>et al.</i> [13]	59.7	38.3	21.9	18.1	33.7
	Li <i>et al.</i> [20]	60.2	44.7	28.8	12.2	36.4
	Ours	63.9	55.7	33.9	20.8	43.5
BEOID	SF-Net [31]	62.9	40.6	16.7	3.5	30.9
	SF-Net* [31]	64.6	42.2	27.3	12.2	36.5
	Ju <i>et al.</i> [13]	63.2	46.8	20.9	5.8	34.9
	Li <i>et al.</i> [20]	71.5	40.3	20.3	5.5	34.4
	Ours	76.9	61.4	42.7	25.1	51.8

Table 2: State-of-the-art comparison on GTEA and BEOID. AVG denotes the average mAP at the thresholds 0.1:0.1:0.7. * denotes the reproduced results by official implementation.

Supervision	Method	mAP@IoU (%)			AVG
		0.5	0.75	0.95	
Frame-level	SSN [62]	41.3	27.0	6.1	26.6
Video-level	Lee <i>et al.</i> [18]	41.2	25.6	6.0	25.9
	AUMN [29]	42.0	25.0	5.6	25.5
	UGCT [54]	41.8	25.3	5.9	25.8
	CoLA [59]	42.7	25.7	5.8	26.1
Point-level	SF-Net [31]	37.8	-	-	22.8
	Ours	44.0	26.0	5.9	26.8

Table 3: State-of-the-art comparison on ActivityNet 1.2. AVG is the averaged mAP at the thresholds 0.5:0.05:0.95.

We provide the experimental results on GTEA and BEOID benchmarks in Table 2. On the both datasets, our method beats the existing state-of-the-art methods with a large gap. Notably, our method shows significant performance boosts under the high thresholds of 0.5 and 0.7, verifying the efficacy of the proposed completeness learning.

Table 3 and Table 4 summarize the results on ActivityNet. Our model shows the superior performances over all the existing weakly-supervised approaches on both versions. It can be also observed that the performance gains upon video-level labels are relatively small compared to THUMOS’14, which we conjecture is due to the far less frequent action instances (1.5 vs. 15 instances per video).

4.3. Analysis

Effect of each component. In Table 5, we conduct ablation study to investigate the contribution of each component. The upper section reports the baseline performances, from which we observe a large score gain brought by the point-level supervision, especially under low IoU thresholds. It mainly comes from the background modeling [17, 18, 38] and the help of point annotations in spotting action instances. On the other hand, the lower section demonstrates the results of the proposed method, where completeness guidance is provided for the model. We observe the absolute average mAP gains of 4.7% and 1.7% from the pro-

Supervision	Method	mAP@IoU (%)			AVG
		0.5	0.75	0.95	
Frame-level	BMN [22]	50.1	34.8	8.3	33.9
	P-GCN [57]	48.3	33.2	3.3	31.1
	G-TAD [51]	50.4	34.6	9.0	34.1
	BC-GNN [1]	50.6	34.8	9.4	34.2
	Zhao <i>et al.</i> [61]	43.5	33.9	9.2	30.1
Video-level	Lee <i>et al.</i> [18]	37.0	23.9	5.7	23.7
	AUMN [29]	38.3	23.5	5.2	23.5
	TS-PCA [54]	37.4	23.5	5.9	23.7
Point-level	Ours	40.4	24.6	5.7	25.1

Table 4: State-of-the-art comparison on ActivityNet 1.3. AVG is the averaged mAP at the thresholds 0.5:0.05:0.95.

$\mathcal{L}_{\text{video}}$	$\mathcal{L}_{\text{point}}$	$\mathcal{L}_{\text{score}}$	$\mathcal{L}_{\text{feat}}$	mAP@IoU (%)				AVG
				0.1	0.3	0.5	0.7	
✓	✗	✗	✗	51.9	37.1	20.3	6.0	28.7
✓	✓	✗	✗	70.7	58.1	40.7	16.1	47.3
✓	✓	✓	✗	75.1	64.4	44.5	20.0	52.0
✓	✓	✗	✓	72.1	60.5	42.1	17.9	49.0
✓	✓	✓	✓	75.7	64.6	45.3	21.8	52.8

Table 5: Ablation study on THUMOS’14. AVG represents the average mAP at the IoU thresholds 0.1:0.1:0.7.

Scoring method	Sequence accuracy	mAP@IoU (%)				AVG
		0.1	0.3	0.5	0.7	
Baseline	N/A	70.7	58.1	40.7	16.1	47.3
(a) Inner scores	74.0	74.7	61.4	40.9	15.2	49.0
(b) Contrast-act	80.1	74.3	63.3	43.6	19.5	50.8
(c) Contrast-both	83.9	75.7	64.6	45.3	21.8	52.8

Table 6: Comparison of different scoring methods for optimal sequence search on THUMOS’14. AVG denotes the average mAP at the IoU thresholds 0.1:0.1:0.7.

posed contrastive losses regarding score and feature similarity, respectively. Moreover, with the two losses combined, the performance is further boosted to 52.8%. This clearly shows that the proposed two losses are complementary and beneficial for precise action localization. Notably, the scores at high IoU thresholds are largely improved, verifying the efficacy of our completeness learning.

Comparison of different scoring methods. In Table 6, we compare different sequence scoring methods regarding frame-level accuracy of optimal sequences in the training set as well as localization performances in the test set of THUMOS’14. Specifically, we investigate three variants: (a) inner scores and (b) score contrast of action instances, and (c) contrast of both action and background ones. As a result, compared to inner scores, the contrast methods generate more accurate optimal sequences and bring larger performance gains at high IoU thresholds. Moreover, we observe that incorporating background instances for score calculation helps to find highly accurate optimal sequences, thereby improving the localization performance at test time.

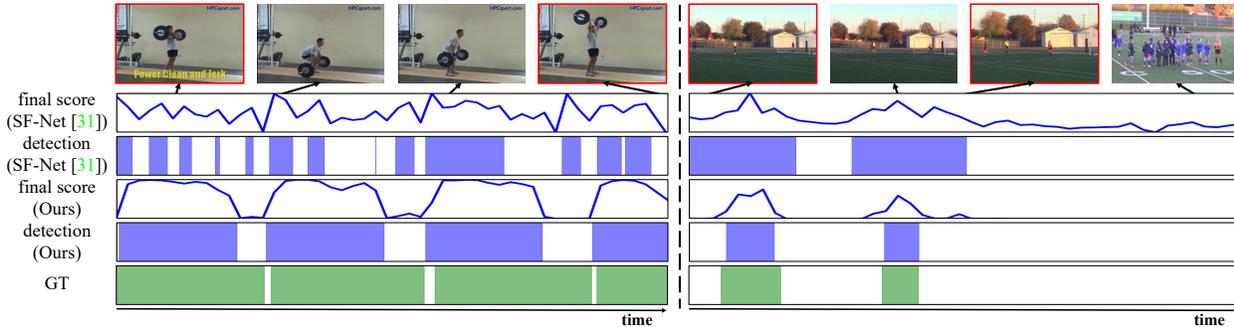


Figure 4: Qualitative comparison with SF-Net [31] on THUMOS’14. We provide two examples with different action classes: (1) *CleanAndJerk* and (2) *SoccerPenalty*. For each video, we present final scores and detection results from SF-Net and our model as well as ground truth action interval. The detection threshold is set to 0.2 for our method and set to the mean score for SF-Net following the original paper. The red boxes indicate the frames that are misclassified by SF-Net but detected by our method. Note that all of our detection results show high IoUs (> 0.6) with the ground-truths.

Method	Distribution	Sequence accuracy	mAP@IoU (%)			AVG
			0.3	0.5	0.7	
SF-Net [31]	Manual	N/A	53.3	28.8	9.7	40.6
	Uniform	N/A	52.0	30.2	11.8	40.5
	Gaussian	N/A	47.4	26.2	9.1	36.7
Ju <i>et al.</i> [13]	Manual	N/A	58.1	34.5	11.9	44.3
	Uniform	N/A	55.6	32.3	12.3	42.9
	Gaussian	N/A	58.2	35.9	12.8	44.8
Ours	Manual	83.7	63.3	43.9	20.8	51.7
	Uniform	76.6	60.4	42.6	20.2	49.3
	Gaussian	83.9	64.6	45.3	21.8	52.8

Table 7: Comparison of the point-level labels from different distributions on THUMOS’14. AVG denotes the average mAP at the IoU thresholds 0.1:0.1:0.7.

Comparison of different label distributions. In Table 7, we explore different label distributions. “Manual” indicates the use of human annotations from [31], whereas the others denote the simulated labels from the corresponding distributions. It is shown that our method significantly outperforms the existing methods regardless of the distribution choice, showing its robustness. We also observe that our method performs slightly worse in “Uniform” compared to the other distributions. We conjecture this is because less discriminative points have more chances to be annotated. Their neighbors are likely to have lower confidence, probably leading to sub-optimal sequences by the greedy algorithm. Indeed, the optimal sequence accuracy is shown to be the lowest in the uniform distribution, which supports our claim.

4.4. Qualitative Comparison

We present qualitative comparisons with SF-Net [31] in Fig. 4. It can be clearly noticed that our method locates the action instances more precisely. Specifically, in the left example, SF-Net produces fragmentary predictions with false negatives, whereas our method detects the complete action instances without splitting them. In the right sample, while

SF-Net overestimates the action instances with false positives, our method produces precise detection results by contrasting action frames from background ones well. The red boxes highlight the false negatives and false positives of SF-Net in the left and right examples, respectively. We note that all the predictions of our model in both examples have high IoUs larger than 0.6 with the corresponding ground-truth instances, validating the effectiveness of our completeness learning. Comparisons on other benchmarks and more visualization results can be found in Sec. C of the appendix.

5. Conclusion

In this paper, we presented a new framework for point-supervised temporal action localization, where dense sequences provide completeness guidance to the model. Concretely, we find the optimal sequence consistent with point labels based on the completeness score, which is efficiently implemented with a greedy algorithm. To learn completeness from the obtained sequence, we introduced two novel losses which encourage contrast between action and background instances regarding action score and feature similarity, respectively. Experiments validated that the optimal sequences are accurate and the proposed losses indeed help to detect complete action instances. Moreover, our model achieves a new state-of-the-art with a large gap on four benchmarks. Notably, it even outperforms fully-supervised methods on average despite the lower supervision level.

Acknowledgements

This project was partly supported by the National Research Foundation of Korea grant funded by the Korea government (MSIT) (No. 2019R1A2C2003760) and the Institute for Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (No. 2020-0-01361: Artificial Intelligence Graduate School Program (YONSEI UNIVERSITY)).

References

- [1] Yueran Bai, Yingying Wang, Yunhai Tong, Yang Yang, Qiyue Liu, and Junhui Liu. Boundary content graph neural network for temporal action proposal generation. In *ECCV*, pages 121–137, 2020. 2, 6, 7
- [2] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015. 6
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. 6
- [4] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018. 2, 5
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020. 5
- [6] Dima Damen, Teesid Leelasawassuk, Osian Haines, Andrew Calway, and Walterio W Mayol-Cuevas. You-do, i-learn: Discovering task relevant objects and their modes of interaction from multi-user egocentric video. In *BMVC*, volume 2, page 3, 2014. 6
- [7] Basura Fernando, Cheston Tan Yin Chet, and Hakan Bilen. Weakly supervised gaussian networks for action detection. In *WACV*, pages 526–535, 2020. 1
- [8] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 5
- [9] Ashraf Islam, Chengjiang Long, and Richard Radke. A hybrid attention mechanism for weakly-supervised temporal action localization. In *AAAI*, pages 1637–1645, 2021. 2
- [10] Ashraf Islam and Richard J. Radke. Weakly supervised temporal action localization using deep metric learning. In *WACV*, pages 536–545, 2020. 2
- [11] Mihir Jain, Amir Ghodrati, and Cees G. M. Snoek. Actionbytes: Learning from trimmed videos to localize actions. In *CVPR*, 2020. 2
- [12] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://csrc.ucf.edu/THUMOS14/>, 2014. 6
- [13] Chen Ju, Peisen Zhao, Ya Zhang, Yanfeng Wang, and Qi Tian. Point-level temporal action localization: Bridging fully-supervised proposals to weakly-supervised losses. *arXiv preprint arXiv:2012.08236*, 2020. 2, 3, 6, 7, 8
- [14] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020. 5
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [16] Jun-Tae Lee, Mihir Jain, Hyungwoo Park, and Sungchack Yun. Cross-attentional audio-visual fusion for weakly-supervised action localization. In *ICLR*, 2021. 2
- [17] Pilhyeon Lee, Youngjung Uh, and Hyeran Byun. Background suppression network for weakly-supervised temporal action localization. In *AAAI*, pages 11320–11327, 2020. 1, 2, 3, 5, 7
- [18] Pilhyeon Lee, Jinglu Wang, Yan Lu, and Hyeran Byun. Weakly-supervised temporal action localization by uncertainty modeling. In *AAAI*, pages 1854–1862, 2021. 1, 2, 3, 6, 7
- [19] Peng Lei and Sinisa Todorovic. Temporal deformable residual networks for action segmentation in videos. In *CVPR*, pages 6742–6751, 2018. 6
- [20] Zhe Li, Yazan Abu Farha, and Juergen Gall. Temporal action segmentation from timestamp supervision. In *CVPR*, pages 8365–8374, 2021. 7
- [21] Chuming Lin, Jian Li, Yabiao Wang, Ying Tai, Donghao Luo, Zhipeng Cui, Chengjie Wang, Jilin Li, Feiyue Huang, and Rongrong Ji. Fast learning of temporal action proposal via dense boundary generator. In *AAAI*, pages 11499–11506, 2020. 2
- [22] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, pages 3888–3897, 2019. 2, 6, 7
- [23] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *ECCV*, pages 3–19, 2018. 2
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 4
- [25] Daochang Liu, Tingting Jiang, and Yizhou Wang. Completeness modeling and context separation for weakly supervised temporal action localization. In *CVPR*, pages 1298–1307, 2019. 2, 5
- [26] Yuan Liu, Jingyuan Chen, Zhenfang Chen, Bing Deng, Jianqiang Huang, and Hanwang Zhang. The blessings of unlabeled background in untrimmed videos. In *CVPR*, pages 6176–6185, 2021. 6
- [27] Yuan Liu, Lin Ma, Yifeng Zhang, Wei Liu, and Shih-Fu Chang. Multi-granularity generator for temporal action proposal. In *CVPR*, pages 3604–3613, 2019. 1
- [28] Ziyi Liu, Le Wang, Qilin Zhang, Zhanning Gao, Zhenxing Niu, Nanning Zheng, and Gang Hua. Weakly supervised temporal action localization through contrast based evaluation networks. In *ICCV*, pages 3899–3908, 2019. 2
- [29] Wang Luo, Tianzhu Zhang, Wenfei Yang, Jingen Liu, Tao Mei, Feng Wu, and Yongdong Zhang. Action unit memory network for weakly supervised temporal action localization. In *CVPR*, pages 9969–9979, 2021. 2, 6, 7
- [30] Zhekun Luo, Devin Guillory, Baifeng Shi, Wei Ke, Fang Wan, Trevor Darrell, and Huijuan Xu. Weakly-supervised action localization with expectation-maximization multi-instance learning. In *ECCV*, 2020. 2
- [31] Fan Ma, Linchao Zhu, Yi Yang, Shengxin Zha, Gourab Kundu, Matt Feiszli, and Zheng Shou. Sf-net: Single-frame supervision for temporal action localization. In *ECCV*, 2020. 1, 2, 3, 4, 6, 7, 8

- [32] Junwei Ma, Satya Krishna Gorti, Maksims Volkovs, and Guangwei Yu. Weakly supervised action selection learning in video. In *CVPR*, pages 7587–7596, 2021. 1
- [33] Yu-Fei Ma, Xian-Sheng Hua, Lie Lu, and Hong-Jiang Zhang. A generic framework of user attention model and its application in video summarization. *IEEE Transactions on Multimedia*, 7:907–919, 2005. 1
- [34] Kyle Min and Jason J. Corso. Adversarial background-aware loss for weakly-supervised temporal activity localization. In *ECCV*, 2020. 2
- [35] Davide Moltisanti, Sanja Fidler, and Dima Damen. Action recognition from single timestamp supervision in untrimmed videos. In *CVPR*, 2019. 2, 6
- [36] Sanath Narayan, Hisham Cholakkal, Fahad Shabaz Khan, and Ling Shao. 3c-net: Category count and center loss for weakly-supervised action localization. In *ICCV*, pages 8678–8686, 2019. 2
- [37] Phuc Nguyen, Ting Liu, Gautam Prasad, and Bohyung Han. Weakly supervised action localization by sparse temporal pooling network. In *CVPR*, pages 6752–6761, 2018. 2, 6
- [38] Phuc Xuan Nguyen, Deva Ramanan, and Charles C. Fowlkes. Weakly-supervised action localization with background modeling. In *ICCV*, pages 5501–5510, 2019. 1, 2, 7
- [39] Sujoy Paul, Sourya Roy, and Amit K Roy-Chowdhury. W-talc: Weakly-supervised temporal activity localization and classification. In *ECCV*, pages 563–579, 2018. 2, 3
- [40] Baifeng Shi, Qi Dai, Yadong Mu, and Jingdong Wang. Weakly-supervised action localization by generative attention modeling. In *CVPR*, 2020. 2
- [41] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-deconvolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, pages 5734–5743, 2017. 1, 2
- [42] Zheng Shou, Hang Gao, Lei Zhang, Kazuyuki Miyazawa, and Shih-Fu Chang. Autoloc: Weakly-supervised temporal action localization in untrimmed videos. In *ECCV*, pages 154–171, 2018. 2, 4
- [43] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, pages 1049–1058, 2016. 1, 2
- [44] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *ICCV*, pages 3544–3553, 2017. 2
- [45] Sarvesh Vishwakarma and Anupam Agrawal. A survey on activity recognition and behavior understanding in video surveillance. *The Visual Computer*, 29:983–1009, 2012. 1
- [46] Limin Wang, Yuanjun Xiong, Dahua Lin, and Luc Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *CVPR*, pages 4325–4334, 2017. 1, 2
- [47] Andreas Wedel, Thomas Pock, Christopher Zach, Horst Bischof, and Daniel Cremers. An improved algorithm for tv-l1 optical flow. In *Statistical and geometrical approaches to visual motion analysis*, pages 23–45. Springer, 2009. 6
- [48] Bo Xiong, Yannis Kalantidis, Deepti Ghadiyaram, and Kristen Grauman. Less is more: Learning highlight detection from video duration. In *CVPR*, pages 1258–1267, 2019. 1
- [49] Yuanjun Xiong, Yue Zhao, Limin Wang, Dahua Lin, and Xiaou Tang. A pursuit of temporal accuracy in general activity detection. *arXiv preprint arXiv:1703.02716*, 2017. 2
- [50] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *ICCV*, pages 5783–5792, 2017. 1
- [51] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *CVPR*, pages 10156–10165, 2020. 2, 6, 7
- [52] Yunlu Xu, Chengwei Zhang, Zhanzhan Cheng, Jianwen Xie, Yi Niu, Shiliang Pu, and Fei Wu. Segregated temporal assembly recurrent networks for weakly supervised multiple action detection. In *AAAI*, 2019. 1, 2
- [53] Ke Yang, Peng Qiao, Dongsheng Li, Shaohe Lv, and Yong Dou. Exploring temporal preservation networks for precise temporal action localization. In *AAAI*, 2018. 2
- [54] Wenfei Yang, Tianzhu Zhang, Xiaoyuan Yu, Tian Qi, Yongdong Zhang, and Feng Wu. Uncertainty guided collaborative training for weakly supervised temporal action detection. In *CVPR*, pages 53–63, 2021. 2, 6, 7
- [55] Jun Yuan, Bingbing Ni, Xiaokang Yang, and Ashraf A Kasim. Temporal action localization with pyramid of score distribution features. In *CVPR*, pages 3093–3102, 2016. 2
- [56] Yuan Yuan, Yueming Lyu, Xi Shen, Ivor Wai-Hung Tsang, and Dit-Yan Yeung. Marginalized average attentional network for weakly-supervised learning. In *ICLR*, 2019. 1
- [57] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *ICCV*, pages 7094–7103, 2019. 2, 6, 7
- [58] Yuanhao Zhai, Le Wang, Wei Tang, Qilin Zhang, Junsong Yuan, and Gang Hua. Two-stream consensus network for weakly-supervised temporal action localization. In *ECCV*, 2020. 2
- [59] Can Zhang, Meng Cao, Dongming Yang, Jie Chen, and Yuexian Zou. Cola: Weakly-supervised temporal action localization with snippet contrastive learning. In *CVPR*, pages 16010–16019, 2021. 6, 7
- [60] Xiaoyu Zhang, Haichao Shi, Changsheng Li, and Peng Li. Multi-instance multi-label action recognition and localization based on spatio-temporal pre-trimming for untrimmed videos. In *AAAI*, 2020. 2
- [61] Peisen Zhao, Lingxi Xie, Chen Ju, Ya Zhang, Yanfeng Wang, and Qi Tian. Bottom-up temporal action localization with mutual regularization. In *ECCV*, pages 539–555, 2020. 2, 6, 7
- [62] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, pages 2914–2923, 2017. 2, 7
- [63] Jia-Xing Zhong, Nannan Li, Weijie Kong, Tao Zhang, Thomas H Li, and Ge Li. Step-by-step erosion, one-by-one collection: a weakly supervised temporal action detector. In *ACM MM*, pages 35–44, 2018. 2