

CoLA: Weakly-Supervised Temporal Action Localization with Snippet Contrastive Learning

Can Zhang¹, Meng Cao¹, Dongming Yang¹, Jie Chen^{1,2}, Yuexian Zou^{*1,2}

¹School of Electronic and Computer Engineering, Peking University ²Peng Cheng Laboratory

{zhangcan, mengcao, yangdongming, zouyx}@pku.edu.cn; chenjie@pcl.ac.cn

Abstract

Weakly-supervised temporal action localization (WS-TAL) aims to localize actions in untrimmed videos with only video-level labels. Most existing models follow the “localization by classification” procedure: locate temporal regions contributing most to the video-level classification. Generally, they process each snippet (or frame) individually and thus overlook the fruitful temporal context relation. Here arises the single snippet cheating issue: “hard” snippets are too vague to be classified. In this paper, we argue that learning by comparing helps identify these hard snippets and we propose to utilize snippet Contrastive learning to Localize Actions, **CoLA** for short. Specifically, we propose a Snippet Contrast (SniCo) Loss to refine the hard snippet representation in feature space, which guides the network to perceive precise temporal boundaries and avoid the temporal interval interruption. Besides, since it is infeasible to access frame-level annotations, we introduce a Hard Snippet Mining algorithm to locate the potential hard snippets. Substantial analyses verify that this mining strategy efficaciously captures the hard snippets and SniCo Loss leads to more informative feature representation. Extensive experiments show that CoLA achieves state-of-the-art results on THUMOS’14 and ActivityNet v1.2 datasets.

1. Introduction

Temporal action localization (TAL) aims at finding and classifying action intervals in untrimmed videos. It has been extensively studied in both industry and academia, due to its wide applications in surveillance analysis, video summarization and retrieval [38, 15, 23], etc. Traditionally, fully-supervised TAL is labor-demanding in its manual labeling procedure, thus weakly-supervised TAL (WS-TAL) which only needs video-level labels has gain popularity.

Most existing WS-TAL methods [39, 27, 30, 26, 14] employ the common attention mechanism or multiple instance learning formulation. Specifically, each input video is di-

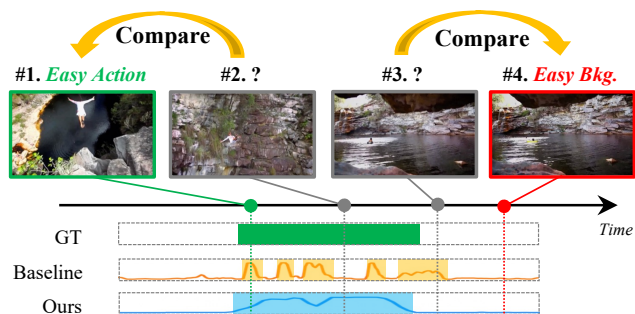


Figure 1. Which category do the two selected snippets (#2, #3) belong to? It is difficult to tell when evaluating independently and they are actually misclassified in *baseline* (We plot the one-dimensional T-CAS for *CliffDiving* and the thresholded results). By contrast, learning by comparing helps identify them: #2 snippet (person falling down) is inferred to be the action snippet by making a comparison with #1 “easy action” (different camera views of the *CliffDiving* action); The inference of #3 snippet is also rectified after the comparison with #4 “easy background” snippet.

vided into multiple fixed-size non-overlapping snippets and the snippet-wise classifications are performed over time to generate the Temporal Class Activation Map/Sequence (T-CAM/T-CAS)[27, 34]. The final localization results are generated by thresholding and merging the class activations. For illustration, we consider the naïve case where the whole process is optimized with a single video-level classification loss and we treat this pipeline as *baseline* in our paper.

In absence of frame-wise labels, WS-TAL suffers from the *single snippet cheating issue*: indistinguishable snippets are easily misclassified and hurt the localization performance. To illustrate it, we take *CliffDiving* in Figure 1 as an example. When evaluated individually, two selected snippets (#2, #3) seem ambiguous and are misclassified: 1) the #2 snippet is incorrectly categorized, thus breaking the time intervals; 2) the #3 snippet is misidentified as an action in *baseline*, resulting in inaccurately extended action interval boundaries. *How to address the single snippet cheating issue?* Let’s revisit the case in Figure 1. By comparing

snippets of interest with those “easy snippets” which can be classified effortlessly, action and background can be distinguished more easily. For example, the #2 snippet and the #1 easy action snippet are two different views of a man falling-down process in “CliffDiving”. The #3 snippet is similar to the #4 easy background snippet and can be easily classified as the background class. In light of this, we contend that *localizing actions by contextually comparing offers a powerful inductive bias that helps distinguish hard snippets*. Based on the above analysis, we propose an alternative, rather intuitive way to address the single snippet cheating issue – by conducting Contrastive learning on hard snippets to Localize Actions, CoLA for short. To this end, we introduce a new Snippet Contrast (SniCo) Loss to refine the feature representations of hard snippets under the guidance of those more discriminative easy snippets. Here these “cheating” snippets are named hard snippets due to their ambiguity.

This solution, however, faces one crucial challenge on how to identify reasonable snippets under our weakly-supervised setting. The selection of hard snippets is non-trivial as there is no specific attention distribution pattern for them. For example, in Figure 1 baseline, #3 hard snippet has a high response value while #2 remains low. Noticing that ambiguous hard snippets are commonly found around boundary areas of the action instances, we propose a boundary-aware Hard Snippet Mining algorithm – a simple yet effective importance sampling technique. Specifically, we first threshold T-CAS and then employ *dilation* and *erosion* operations temporally to mine the potential hard snippets. Since the hard snippets may either be action or background, we opt to distinguish them by their relative position. For easy snippets, they locate in the most discriminative parts, so snippets with *top-k/bottom-k* T-CAS scores are selected as easy action/background respectively. Moreover, we form two hard-easy contrastive pairs and conduct the feature refinement via the proposed SniCo Loss.

In a nutshell, the main contributions of this work are as follows: (1) Pioneeringly, we introduce the contrastive representation learning paradigm to WS-TAL and propose a SniCo Loss which effectively refines the feature representation of hard snippets. (2) A Hard Snippet Mining algorithm is proposed to locate potential hard snippets around boundaries, which serves as an efficient sampling strategy under our weakly-supervised setting. (3) Extensive experiments on THUMOS’14 and ActivityNet v1.2 datasets demonstrate the effectiveness of our proposed CoLA.

2. Related Work

Fully-supervised Action Localization utilizes frame-level annotations to locate and classify the temporal intervals of action instances from long untrimmed videos. Most existing works may be classified into two cate-

gories: proposal-based (top-down) and frame-based methods (bottom-up). Proposal-based methods [35, 47, 40, 7, 5, 33, 19, 17, 44, 16] first generate action proposals and then classify them as well as conduct temporal boundary regression. On the contrary, frame-based methods [18, 2, 22, 46] directly predict frame-level action category and location followed by some post-processing techniques.

Weakly-Supervised Action Localization only requires video-level annotations and has drawn extensive attention. UntrimmedNets [39] address this problem by conducting the clip proposal classification first and then select relevant segments in a soft or hard manner. STPN [27] imposes a sparsity constraint to enforce the sparsity of the selected segments. Hide-and-seek [36] and MAAN [43] try to extend the discriminative regions via randomly hiding patches or suppressing the dominant response, respectively. Zhong *et al.* [48] introduce a progressive generation procedure to achieve similar ends. W-TALC [30] applies the deep metric learning to be complementary with the Multiple Instance Learning formulation.

Discussion. The single snippet cheating problem has not been fully studied though it is common in WS-TAL. Liu *et al.* [20] pinpoint the action completeness modeling problem and the action-context separation problem. They develop a parallel multi-branch classification architecture with the help of the generated hard negative data. In contrast, our CoLA unifies these two problems and settles them in a lighter way with the proposed SniCo Loss. DGAM [32] mentions the *action-context confusion issue*, *i.e.*, context snippets near action snippets tend to be misclassified, which can be considered as a sub-problem of our *single snippet cheating issue*. Besides, several background modeling works [28, 14, 32] can also be seen as one solution to this problem. Nguyen *et al.* [28] utilizes an attention mechanism to model both foreground and background frame appearances and guide the generation of the class activation map. BaS-Net [14] introduces an auxiliary class for background and applies an asymmetrical training strategy to suppress the background snippet activation. However, these methods have inherent drawbacks as background snippets are not necessarily motionless and it is difficult to include them into one specific class. By contrast, our CoLA is a more adaptive and explainable solution to tackle these issues.

Contrastive Representation Learning uses data internal patterns to learn an embedding space where associated signals are brought together while unassociated ones are distinguished via Noise Contrastive Estimation (NCE) [8]. CMC [37] presents a contrastive learning framework that maximize mutual information between different views of the same scene to achieve a view-invariant representation. SimCLR [6] selects the negative samples by using augmented views of other items in a minibatch. MoCo [9] uses a momentum updated memory bank of old negative repre-

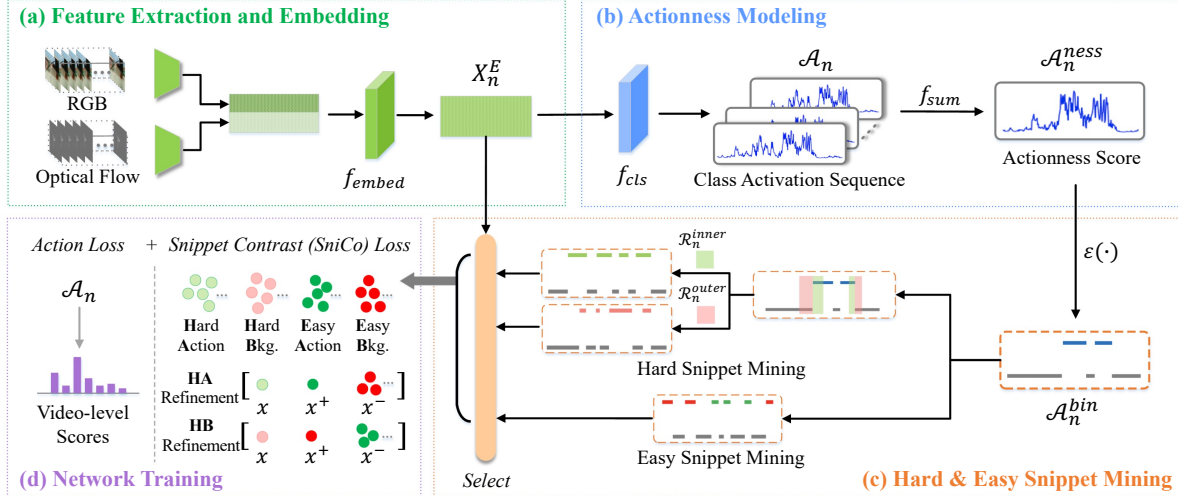


Figure 2. Illustration of the proposed CoLA, which consists of four parts: (a) Feature Extraction and Embedding to obtain the embedded feature X_n^E ; (b) Actionness Modeling to gather class-agnostic action likelihood \mathcal{A}_n^{ness} ; (c) Hard & Easy Snippet Mining to select hard and easy snippets. (d) Network Training driven by Action Loss and Snippet Contrast (SniCo) Loss.

representations to get rid of the batch size restriction and enable the consistent use of negative samples. To our best knowledge, we are the first to introduce the noise contrastive estimation to WS-TAL task. Experiment results show that CoLA refines the hard snippet representation, thus benefiting the action localization.

3. Method

Generally, CoLA (shown in Figure 2) follows the feature extraction (Section 3.1), actionness modeling (Section 3.2) and hard & easy snippet mining (Section 3.3) pipeline. The optimization loss terms and the inference process are detailed in Section 3.4 and Section 3.5, respectively.

3.1. Feature Extraction and Embedding

Assume that we are given a set of N untrimmed videos $\{V_n\}_{n=1}^N$ and their video-level labels $\{y_n\}_{n=1}^N$, where $y_n \in \mathbb{R}^C$ is a multi-hot vector, and C is the number of action categories. Following the common practice [27, 28, 14], for each input untrimmed video V_n , we divide it into multi-frame non-overlapping L_n snippets, *i.e.*, $V_n = \{S_{n,l}\}_{l=1}^{L_n}$. A fixed number of T snippets $\{S_{n,t}\}_{t=1}^T$ are sampled due to the variation of video length. Then the RGB features $X_n^R = \{x_t^R\}_{t=1}^T$ and optical flow features $X_n^O = \{x_t^O\}_{t=1}^T$ are extracted with pre-trained feature extractor (*e.g.*, I3D [4]), respectively. Here, $x_t^R \in \mathbb{R}^d$ and $x_t^O \in \mathbb{R}^d$, d is the feature dimension of each snippet. Afterwards, we apply an embedding function f_{embed} over the concatenation of X_n^R and X_n^O to obtain our extracted features $X_n^E \in \mathbb{R}^{T \times 2d}$. f_{embed} is implemented with a temporal convolution followed by the ReLU activation function.

3.2. Actionness Modeling

We introduce the concept *Actionness* referring to the likelihood of containing a general action instance for each snippet. Before we specify the Actionness Modeling process, let's revisit the commonly adopted Temporal Class Activation Sequence (T-CAS).

Given the embedded features X_n^E , we apply a classifier f_{cls} to obtain snippet-level T-CAS. Specifically, the classifier contains a temporal convolution followed by ReLU activation and Dropout. This can be formulated as follows for a video V_n :

$$\mathcal{A}_n = f_{cls}(X_n^E; \phi_{cls}), \quad (1)$$

where ϕ_{cls} represents the learnable parameters. The obtained $\mathcal{A}_n \in \mathbb{R}^{T \times C}$ represents the action classification results occurring at each temporal snippets.

Then, when it comes to modeling the actionness, one common way is to conduct the binary classification on each snippet, which yet will inevitably bring in extra overheads. Since the generated T-CAS $\mathcal{A}_n \in \mathbb{R}^{T \times C}$ in Eqn. 1 already contains snippet-level class-specific predictions, we simply sum T-CAS along the channel dimension (f_{sum}) followed by the Sigmoid function to obtain a class-agnostic aggregation and use it to represent the actionness $\mathcal{A}_n^{ness} \in \mathbb{R}^T$:

$$\mathcal{A}_n^{ness} = \text{Sigmoid}(f_{sum}(\mathcal{A}_n)). \quad (2)$$

3.3. Hard & Easy Snippet Mining

Recall that our aim is to use the easily spotted snippets as a priori to disambiguate controversial snippets. We systematically study the contrastive pair construction process for both hard and easy snippets.




Figure 3. Illustration of the Hard Snippet Mining algorithm. **Left:** Subtract the eroded sequences with different masks to get the inner regions (green color); **Right:** Subtract the dilated sequences with different masks to get the outer regions (pink color).

3.3.1 Hard Snippet Mining

Intuitively, for most snippets located inside the action or background intervals, they are far from the temporal borders with less noise interference and have the relatively trustworthy feature representation. For boundary-adjacent snippets, however, they are less reliable because they are in the transitional areas between action and background, thus leading to ambiguous detection.

Based on the above observations, we argue that boundary-adjacent snippets can serve as the potential hard snippets under the weak supervision setting. Therefore, we build a novel Hard Snippet Mining algorithm to exploit hard snippets from the border areas. Then these mined hard snippets are divided into **hard action** and **hard background** according to their locations.

Firstly, we threshold the actionness scores to generate a binary sequence (1 or 0 indicates the action or background location, respectively):

$$\mathcal{A}_n^{bin} = \varepsilon(\mathcal{A}_n^{ness} - \theta_b), \quad (3)$$

where $\varepsilon(\cdot)$ is the *Heaviside step function* and θ_b is the threshold value, i.e., \mathcal{A}_n^{bin} is 1 if $\mathcal{A}_n^{ness} \geq \theta_b$, 0 otherwise. Then, as shown in Figure 3, we apply two cascaded dilation or erosion operations to expand or narrow the temporal extent of action intervals. The differential areas with the diverse dilation or erosion degree are defined as the hard background or hard action regions:

$$\begin{aligned} \mathcal{R}_n^{inner} &= (\mathcal{A}_n^{bin}; m)^- - (\mathcal{A}_n^{bin}; \mathcal{M})^- \\ \mathcal{R}_n^{outer} &= (\mathcal{A}_n^{bin}; \mathcal{M})^+ - (\mathcal{A}_n^{bin}; m)^+, \end{aligned} \quad (4)$$

where $(\cdot; *)^+$ and $(\cdot; *)^-$ represent the binary dilation and erosion operations with mask $*$, respectively. The inner region \mathcal{R}_n^{inner} is defined as the different snippets between the eroded sequences with smaller mask m and larger mask \mathcal{M} , as shown in Figure 3 left part (in green color). Similarly, the outer region \mathcal{R}_n^{outer} is calculated as the difference between the dilated sequences with larger mask \mathcal{M} and smaller mask m , depicted in Figure 3 right part (in pink color). Empirically, we regard the inner regions \mathcal{R}_n^{inner} as hard action snippet sets since these regions are with $\mathcal{A}_n^{bin} = 1$. Similarly, the outer regions \mathcal{R}_n^{outer} are considered as hard

background snippet sets. Then the hard action snippets $X_n^{HA} \in \mathbb{R}^{k^{hard} \times 2d}$ are selected from \mathcal{R}_n^{inner} :

$$X_n^{HA} = \{X_{n;t}^E | t \in \mathcal{I}_n^{act}, \mathcal{I}_n^{act} \in \mathcal{I}_n^{inner}\}, \quad (5)$$

where \mathcal{I}_n^{inner} is the index set of snippets within \mathcal{R}_n^{inner} . \mathcal{I}_n^{act} is the subset of \mathcal{I}_n^{inner} with size k^{hard} (i.e., $|\mathcal{I}_n^{act}| = k^{hard}$), and $k^{hard} = \max(1, \lfloor \frac{T}{r^{hard}} \rfloor)$ is the hyper-parameter controlling the selected number of hard snippets, r^{hard} is the sampling ratio. Considering the case that $k^{hard} > |\mathcal{I}_n^{inner}|$, we adopt *sampling with replacement* mechanism to ensure the total k^{hard} snippets can be selected. Similarly, the hard background snippets $X_n^{HB} \in \mathbb{R}^{k^{hard} \times 2d}$ are selected from \mathcal{R}_n^{outer} :

$$X_n^{HB} = \{X_{n;t}^E | t \in \mathcal{I}_n^{bkg}, \mathcal{I}_n^{bkg} \in \mathcal{I}_n^{outer}\}, \quad (6)$$

where the notation definitions are similar to those in Eqn. 5 and we omit them for brevity.

3.3.2 Easy Snippet Mining

In order to form contrastive pairs, we still need to mine the discriminative easy snippets. Based on the well-trained fully-supervised I3D features, we hypothesize that the video snippets with top- k and bottom- k actionness scores are exactly easy action ($X_n^{EA} \in \mathbb{R}^{k^{easy} \times 2d}$) and easy background snippets ($X_n^{EB} \in \mathbb{R}^{k^{easy} \times 2d}$), respectively. Therefore, we conduct easy snippet mining based on the actionness scores calculated in Eqn. 2. The specific process is as follows:

$$\begin{aligned} X_n^{EA} &= \{X_{n;t}^E | t \in \mathcal{S}_n^{act}, t \notin \mathcal{I}_n^{inner}, t \notin \mathcal{I}_n^{outer}, \\ &\quad \mathcal{S}_n^{act} = \mathcal{S}_n^{DESC}[:k^{easy}]\} \\ X_n^{EB} &= \{X_{n;t}^E | t \in \mathcal{S}_n^{bkg}, t \notin \mathcal{I}_n^{inner}, t \notin \mathcal{I}_n^{outer}, \\ &\quad \mathcal{S}_n^{bkg} = \mathcal{S}_n^{ASC}[:k^{easy}]\}, \end{aligned} \quad (7)$$

where \mathcal{S}_n^{DESC} and \mathcal{S}_n^{ASC} denotes the index of \mathcal{A}_n^{ness} sorting by DESC and ASC order respectively. $k^{easy} = \max(1, \lfloor \frac{T}{r^{easy}} \rfloor)$, r^{easy} is a hyper-parameter representing the selection ratio. Note that we remove the snippets in the hard snippet areas \mathcal{R}_n^{inner} and \mathcal{R}_n^{outer} to avoid conflict.

3.4. Network Training

Based on the mined hard and easy snippets, our CoLA introduces an additional Snippet Contrast (SniCo) Loss (\mathcal{L}_s) and achieves considerable improvement compared with the *baseline* model. The total loss can be represented as follows:

$$\mathcal{L}_{total} = \mathcal{L}_a + \lambda \mathcal{L}_s, \quad (8)$$

where \mathcal{L}_a and \mathcal{L}_s denote the Action Loss and the SniCo Loss, respectively. λ is the balance factor. We elaborate on these two terms as follows.

3.4.1 Action Loss

Action Loss (\mathcal{L}_a) is the classification loss between the predicted video category and the ground truth. To get the video-level predictions, we aggregate snippet-level class scores computed in Eqn. 1. Following [39, 30, 14], we take the *top-k mean* strategy: for each class c , we take k^{easy} terms with the largest class-specific T-CAS values and compute their means as $a_{n;c}$, namely the video-level class score for class c of video V_n . After obtaining $a_{n;c}$ for all the C classes, we apply a Softmax function on a_n along the class dimension to get the video-level class possibilities $p_n \in \mathbb{R}^C$. Action Loss (\mathcal{L}_a) is then calculated in the cross-entropy form:

$$\mathcal{L}_a = -\frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C \hat{y}_{n;c} \log(p_{n;c}), \quad (9)$$

where $\hat{y}_n \in \mathbb{R}^C$ is the normalized ground-truth.

3.4.2 Snippet Contrast (SniCo) Loss

Contrastive learning has been used on image or patch levels [1, 10]. For our application, given the extracted feature embedding X_n^E , the contrastive learning is applied in the snippet level. We name it *Snippet Contrast (SniCo) Loss* (\mathcal{L}_s), which aims to refine the snippet-level feature of hard snippets and obtain a more informative feature distribution. Considering that the hard snippets are classified as hard action and hard background, we form two contrastive pairs in \mathcal{L}_s accordingly, namely ‘‘HA refinement’’ and ‘‘HB refinement’’, where HA and HB are short for hard action and hard background respectively. ‘‘HA refinement’’ aims to transform the hard action snippet features by driving hard action and easy action snippets compactly in feature space and ‘‘HB refinement’’ is similar.

Formally, the query $\mathbf{x} \in \mathbb{R}^{1 \times 2d}$, positive $\mathbf{x}^+ \in \mathbb{R}^{1 \times 2d}$, and S negatives $\mathbf{x}^- \in \mathbb{R}^{S \times 2d}$ are selected from pre-mined snippets. As shown in Figure 2(d), for ‘‘HA refinement’’, $\mathbf{x} \sim X_n^{HA}$, $\mathbf{x}^+ \sim X_n^{EA}$, $\mathbf{x}^- \sim X_n^{EB}$; for ‘‘HB refinement’’, $\mathbf{x} \sim X_n^{HB}$, $\mathbf{x}^+ \sim X_n^{EB}$, $\mathbf{x}^- \sim X_n^{EA}$. We project them to a normalized unit sphere to prevent the space from collapsing or expanding. An $(S + 1)$ -way classification problem using the cross-entropy loss is set up to represent the probability of the positive example being selected over negatives. Following [9], we compute the distances between the query and other examples with a temperature scale $\tau = 0.07$:

$$\begin{aligned} & \ell(\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-) \\ &= -\log \left[\frac{\exp(\mathbf{x}^T \cdot \mathbf{x}^+ / \tau)}{\exp(\mathbf{x}^T \cdot \mathbf{x}^+ / \tau) + \sum_{s=1}^S \exp(\mathbf{x}^T \cdot \mathbf{x}_s^- / \tau)} \right], \end{aligned} \quad (10)$$

where \mathbf{x}^T is the transpose of \mathbf{x} and the proposed SniCo Loss is as follows:

$$\begin{aligned} \mathcal{L}_s = & \underbrace{\mathbb{E}_{\mathbf{x} \sim X_n^{HA}, \mathbf{x}^+ \sim X_n^{EA}, \mathbf{x}^- \sim X_n^{EB}} \ell(\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-)}_{\text{HA refinement}} \\ & + \underbrace{\mathbb{E}_{\mathbf{x} \sim X_n^{HB}, \mathbf{x}^+ \sim X_n^{EB}, \mathbf{x}^- \sim X_n^{EA}} \ell(\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-)}_{\text{HB refinement}}, \end{aligned} \quad (11)$$

where S represents the number of negative snippets and $\mathbf{x}_s^- \in \mathbb{R}^{2d}$ means the s -th negative. In this way, we maximize *mutual information* between the easy and hard snippets of the same category (action or background), which helps refine the feature representation and thereby alleviating the single snippet cheating issue.

3.5. Inference

Given an input video, we first predict its snippet-level class activations to form T-CAS and aggregate top- k^{easy} scores described in Sec. 3.4.1 to get the video-level predictions. Then the categories with scores larger than θ_v are selected for further localization. For each selected category, we threshold its corresponding T-CAS with θ_s to obtain candidate video snippets. Finally, continuous snippets are grouped into proposals and Non-Maximum Suppression (NMS) is applied to remove duplicated proposals.

4. Experiments

4.1. Datasets

We evaluate our CoLA on two popular action localization benchmark datasets including THUMOS’14 [11] and ActivityNet v1.2 [3]. We only use the video-level category labels for network training.

THUMOS’14 includes untrimmed videos with 20 categories. The video length varies greatly and each video may contain multiple action instances. By convention [14, 32], we use the 200 videos in validation set for training and the 213 videos in testing set for evaluation.

ActivityNet v1.2 is a popular large-scale benchmark for TAL with 100 categories. Following the common practice [39, 34], we train on the training set with 4,819 videos and test on the validation set with 2,383 videos.

4.2. Implementation Details

Evaluation Metrics. We follow the standard evaluation protocol by reporting mean Average Precision (mAP) values under different intersection over union (IoU) thresholds. The evaluation on both datasets are conducted using the benchmark code provided by ActivityNet¹.

¹<https://github.com/activitynet/ActivityNet/>

Feature Extractor. We use I3D [4] network pre-trained on Kinetics [4] for feature extraction. Note that the I3D feature extractor is not fine-tuned for fair comparison. TVL1 [31] algorithm is applied to extract optical flow stream from RGB stream in advance. Each video stream is divided into 16-frame non-overlapping snippets and the snippet-wise RGB and optical flow features are with 1024-dimension.

Training Details. The number of sampled snippets T for THUMOS’14 and ActivityNet v1.2 is set to 750 and 50, respectively. All hyper-parameters are determined by grid search: $r^{easy} = 8$, $r^{hard} = 32$, $S = k^{easy} = \max(1, \lfloor \frac{T}{r^{easy}} \rfloor)$. We set $\lambda = 5e - 3$ in Eqn. 8. θ_b in Eqn. 3 is set to 0.5 for both datasets. Dilation and erosion masks \mathcal{M} and m are set to 6 and 3 in our experiments. We utilize Adam optimizer with a learning rate of $1e - 4$. We train for total 6k epochs with a batch size of 16 for THUMOS’14 and for total 8k epochs with a batch size of 128 for ActivityNet v1.2.

Testing Details. We set θ_v to 0.2 and 0.1 for THUMOS’14 and ActivityNet v1.2, respectively. For proposal generation, we use multiple thresholds that θ_s is set as [0:0.25:0.025] for THUMOS’14 and [0:0.15:0.015] for ActivityNet v1.2, then Non-Maximum Suppression (NMS) is performed with IoU threshold 0.7.

4.3. Comparison with State-of-the-Arts

We compare our CoLA with the state-of-the-art fully-supervised and weakly-supervised TAL approaches on THUMOS’14 testing set. As shown in Table 1, CoLA achieves the impressive performance, *i.e.*, we consistently outperform previous weakly-supervised methods at all IoU thresholds. Specifically, our method achieves 32.2% mAP@0.5 and 40.9% mAP@AVG, bringing the state-of-the-art to a new level. Notably, even with a much lower level of supervision, our method is even comparable with several fully-supervised methods, following the latest fully-supervised approaches with the least gap.

We also conduct experiments on ActivityNet v1.2 validation set and the comparison results are summarized in Table 2. Again, our method shows significant improvements over state-of-the-art weakly-supervised TAL methods while maintaining competitive compared with other fully-supervised methods. The consistent superior results on both datasets signify the effectiveness of CoLA.

4.4. Ablation Studies

In this section, we conduct multiple ablation studies to provide more insights about our design intuition. By convention [28, 32, 14], all the ablation experiments are performed on the THUMOS’14 testing set.

Q1: How does the proposed SniCo Loss help? To evaluate the effectiveness of our SniCo Loss (\mathcal{L}_s), we con-

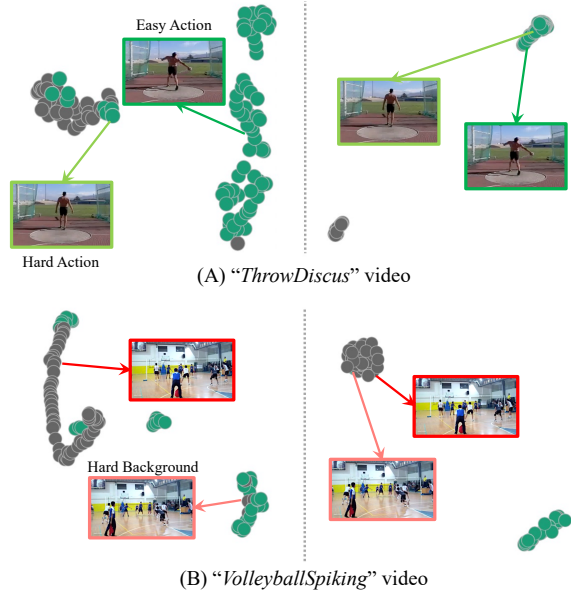


Figure 4. UMAP visualizations of feature embeddings X_n^E . Left: baseline; Right: CoLA. Green points represent action embeddings and gray points denote background embeddings. CoLA achieves a more separable feature distribution compared to baseline.

duct a comparison experiment with only the action loss \mathcal{L}_a as supervision, namely *baseline* in Table 3. The statistical results in Table 3 demonstrate that by introducing \mathcal{L}_s , the performance largely gains by 7.5% in mAP@0.5, partially because SniCo Loss effectively guides the network to achieve better feature distribution tailored for WS-TAL. To illustrate this, we randomly select 2 videos from THUMOS’14 testing set and calculate the feature embeddings X_n^E for baseline and CoLA, respectively. These embeddings are then projected to 2-dimensional space using UMAP [24], as shown in Figure 4. Notice that compared with baseline, SniCo Loss helps to separate the action and background snippets more precisely, especially for those ambiguous hard snippets. Overall, the above analyses strongly justify the significance of our proposed SniCo Loss.

Q2: Is it necessary to consider both HA and HB refinements in SniCo Loss? To explore this, we conduct ablated experiments with two variants of SniCo Loss, each of which contains only one type of refinement in Eqn. 11, namely \mathcal{L}_s^{HA} and \mathcal{L}_s^{HB} , respectively. Table 3 shows that the performance drops dramatically with either kind of refinement removed, suggesting that both refinements contribute to the improved performance.

Q3: Are our mined hard snippets meaningful? How to evaluate the effectiveness of the mined hard snippets is nontrivial. As discussed in Sec. 3.3.1, indistinguishable frames usually exist within or near the action temporal intervals, so we define such temporal areas as *error-prone*

Table 1. Comparisons with state-of-the-art TAL methods on THUMOS’14 dataset. The mAP values at different IoU thresholds are reported. The AVG column shows the averaged mAP under the thresholds [0.1:0.7:0.1]. UNT is the abbreviation for UntrimmedNet feature.

Supervision (Feature)	Method	Publication	mAP@IoU (%)							
			0.1	0.2	0.3	0.4	0.5	0.6	0.7	AVG
Full (-)	R-C3D [40]	ICCV 2017	54.5	51.5	44.8	35.6	28.9	-	-	-
	SSN [47]	ICCV 2017	66.0	59.4	51.9	41.0	29.8	-	-	-
	TAL-Net [5]	CVPR 2018	59.8	57.1	53.2	48.5	42.8	33.8	20.8	45.1
	P-GCN [44]	ICCV 2019	69.5	67.8	63.6	57.8	49.1	-	-	-
	G-TAD [41]	CVPR 2020	-	-	66.4	60.4	51.6	37.6	22.9	-
Weak (-)	Hide-and-Seek [36]	ICCV 2017	36.4	27.8	19.5	12.7	6.8	-	-	-
	UntrimmedNet [39]	CVPR 2017	44.4	37.7	28.2	21.1	13.7	-	-	-
	Zhong <i>et al.</i> [48]	ACMMM 2018	45.8	39.0	31.1	22.5	15.9	-	-	-
Weak (UNT)	AutoLoc [34]	ECCV 2018	-	-	35.8	29.0	21.2	13.4	5.8	-
	CleanNet [21]	ICCV 2019	-	-	37.0	30.9	23.9	13.9	7.1	-
	BaS-Net [14]	AAAI 2020	-	-	42.8	34.7	25.1	17.1	9.3	-
Weak (I3D)	STPN [27]	CVPR 2018	52.0	44.7	35.5	25.8	16.9	9.9	4.3	27.0
	Liu <i>et al.</i> [20]	CVPR 2019	57.4	50.8	41.2	32.1	23.1	15.0	7.0	32.4
	Nguyen <i>et al.</i> [28]	ICCV 2019	60.4	56.0	46.6	37.5	26.8	17.6	9.0	36.3
	BaS-Net [14]	AAAI 2020	58.2	52.3	44.6	36.0	27.0	18.6	10.4	35.3
	DGAM [32]	CVPR 2020	60.0	54.2	46.8	38.2	28.8	19.8	11.4	37.0
	ActionBytes [12]	CVPR 2020	-	-	43.0	35.8	29.0	-	9.5	-
	A2CL-PT [25]	ECCV 2020	61.2	56.1	48.1	39.0	30.1	19.2	10.6	37.8
	TSCN [45]	ECCV 2020	63.4	57.6	47.8	37.7	28.7	19.4	10.2	37.8
	CoLA (Ours)	-	66.2	59.5	51.5	41.9	32.2	22.0	13.1	40.9

Table 2. Comparison results on ActivityNet v1.2 dataset. The AVG column shows the averaged mAP under the thresholds [0.5:0.95:0.05]. UNT and I3D are abbreviations for UntrimmedNet feature and I3D feature, respectively.

Sup.	Method	mAP@IoU (%)			
		0.5	0.75	0.95	AVG
Full	SSN [47]	41.3	27.0	6.1	26.6
Weak (UNT)	UntrimmedNet [39]	7.4	3.2	0.7	3.6
	AutoLoc [34]	27.3	15.1	3.3	16.0
	W-TALC [30]	37.0	12.7	1.5	18.0
Weak (I3D)	TSM [42]	28.3	17.0	3.5	17.1
	CleanNet [21]	37.1	20.3	5.0	21.6
	Liu <i>et al.</i> [20]	36.8	22.0	5.6	22.4
	BaS-Net [14]	38.5	24.2	5.6	24.3
	DGAM [32]	41.0	23.5	5.3	24.4
	TSCN [45]	37.6	23.7	5.7	23.6
	CoLA (Ours)	42.7	25.7	5.8	26.1

Table 3. Ablation analysis on loss terms on THUMOS’14.

Setting	Loss	mAP@0.5 (Δ)
CoLA (Ours)	$\mathcal{L}_a + \mathcal{L}_s$	32.2%
baseline	\mathcal{L}_a	24.7% (-7.5%)
CoLA w/o HB ref.	$\mathcal{L}_a + \mathcal{L}_s^{HA}$	29.7% (-2.5%)
CoLA w/o HA ref.	$\mathcal{L}_a + \mathcal{L}_s^{HB}$	30.4% (-1.8%)

regions. Specifically, given a ground-truth action instance with interval $[s, e]$ and duration $d = e - s$, we define its δ -scale error-prone regions as $[s - \delta \frac{d}{2}, e + \delta \frac{d}{2}]$, as illustrated in

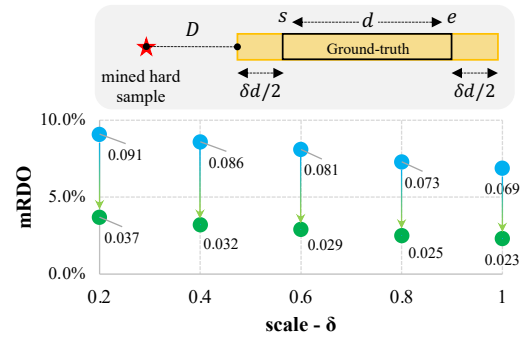


Figure 5. Effectiveness verification of Hard Snippet Mining algorithm. **Top:** Illustration of relative distance offsets (RDO) for a mined snippet. **Bottom:** The mean RDO (mRDO) vs. different scales δ at epoch 0 (blue) and epoch 2k (green).

Figure 5 top part. Then, to evaluate the positional relationship of our mined hard snippets with the error-prone areas, *relative distance offset* (RDO) is defined as follows: 1) if a mined hard snippet does not fall into any of the error-prone regions, $RDO = \frac{D}{T}$, where D is the nearest distance between this snippet and all error-prone regions, and T is the video length; 2) otherwise, $RDO = 0$. As shown in Figure 5 bottom part, the mean RDO values (mRDO) of all the videos are evaluated under different scales δ at two training snapshots (epoch 0 and epoch 2k). The mRDO consistently drops at all scales δ , indicating that our mined hard snippets

Table 4. Ablation analysis on the negative sample size S .

S	1	4	16	64	125 (k^{easy})
mAP@0.5	28.9	30.4	31.3	31.9	32.2

Table 5. Ablation analysis on the mask size \mathcal{M} and m .

$\mathcal{M}(m=3)$	4	5	6	7	8	9
mAP@0.5	30.9	31.8	32.2	32.0	31.8	32.1
$m(\mathcal{M}=6)$	0	1	2	3	4	5
mAP@0.5	30.3	31.7	32.0	32.2	32.0	31.9

are captured more precisely as the training goes on. Even under the most stringent condition ($\delta = 0.2$), the mRDO is only 3.7%, which suggests that most of our mined hard snippets locate in such error-prone areas and thus contribute to the network training.

Evaluation on the negative sample size S . Table 4 reports the experimental results evaluated with different negative sample sizes S . According to Eqn. 11, negative snippets are randomly chosen from the mined easy snippets, so $S \leq k^{easy}$. As shown, the mAP value is positively correlated with S , indicating that contrastive power increases by adding more negatives. This phenomenon is consistent with many self-supervised contrastive learning works [29, 9, 6] and a recent supervised one [13], which partially verifies the efficacy of our hard and easy snippet mining algorithm for weakly-supervised TAL task.

Evaluation on the mask size \mathcal{M} and m . We have defined two operation degrees (with larger \mathcal{M} and smaller m) for temporal interval erosion and dilation in Eqn. 4. Here we seek to evaluate the effect of different mask sizes. For simplification, we first fix $m = 3$ and vary \mathcal{M} from 4 to 9, then we fix $\mathcal{M} = 6$ and change m from 0 to 5. The results are shown in Table 5. The best result is achieved when setting $\mathcal{M} = 6$ and $m = 3$. Besides, it is quite evident that the performance remains stable across a wide range of \mathcal{M} and m , demonstrating the robustness of our proposed Hard Snippet Mining algorithm.

4.5. Qualitative Results

We visualize T-CAS results for two actions on THUMOS’14 in Figure 6. Our CoLA has a more informative T-CAS distribution compared to baseline, thus leading to more accurate localization. Figure 6-A depicts a typical case that all the frames in a video share the similar elements, *i.e.*, humans, billiard table and balls. By introducing SniCo Loss, our method can seek the subtle differences between action and hard background, thereby avoiding many false positives produced by single Action Loss (baseline). Figure 6-B demonstrates a “CliffDiving” action observed from different camera views. The baseline method fails to localize the complete interval and outputs short and sparse

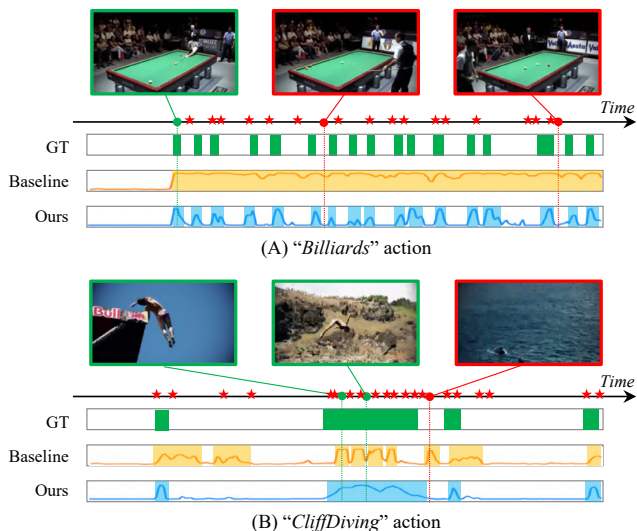


Figure 6. Qualitative comparisons with baseline on THUMOS’14. For baseline and CoLA, we visualize the one-dimensional T-CAS and the localized regions. For clarity, frames with green bounding boxes refer to ground-truth actions and those in red refer to ground-truth backgrounds. Red pentagrams along the time axis denote the mined hard snippet locations (computed at epoch 2k).

prediction results. Our method successfully identifies the entire “CliffDiving” action and suppress the false positive detections. We also visualize the mined hard snippet locations (computed at epoch 2k) on the time axis (marked as red pentagram). As expected, these snippets are misclassified in baseline and CoLA refines their representation to achieve better performance. This visualization also helps explain Q3 in Section 4.4. For more visualization results, please refer to our supplementary materials.

5. Conclusion

In this paper, we have proposed a novel framework (CoLA) to address the single snippet cheating issue in weakly-supervised action localization. We leverage the intuition that hard snippets frequently lay in the boundary regions of the action instances and propose a Hard Snippet Mining algorithm to localize them. Then we apply a SniCo Loss to refine the feature representation of the mined hard snippets with the help of easy snippets which locate in the most discriminative regions. Experiments conducted on two benchmarks including THUMOS’14 and ActivityNet v1.2 have validated the state-of-the-art performance of CoLA.

Acknowledgements. This paper was partially supported by the IER foundation (No. HT-JD-CXY-201904) and Shenzhen Municipal Development and Reform Commission (Disciplinary Development Program for Data Science and Intelligent Computing). Special acknowledgements are given to Aoto-PKUSZ Joint Lab for its support.

References

- [1] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15535–15545, 2019. 5
- [2] Shyamal Buch, Victor Escorcia, Bernard Ghanem, Li Fei-Fei, and Juan Carlos Niebles. End-to-end, single-stream temporal action detection in untrimmed videos. 2019. 2
- [3] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. 5
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 3, 6
- [5] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1130–1139, 2018. 2, 7
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 2, 8
- [7] Xiyang Dai, Bharat Singh, Guyue Zhang, Larry S Davis, and Yan Qiu Chen. Temporal context network for activity localization in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5793–5802, 2017. 2
- [8] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304, 2010. 2
- [9] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 2, 5, 8
- [10] Olivier J Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019. 5
- [11] Haroon Idrees, Amir R Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding*, 155:1–23, 2017. 5
- [12] Mihir Jain, Amir Ghodrati, and Cees G. M. Snoek. Action-bytes: Learning from trimmed videos to localize actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 7
- [13] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020. 8
- [14] 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, 6, 7
- [15] Yong Jae Lee, Joydeep Ghosh, and Kristen Grauman. Discovering important people and objects for egocentric video summarization. In *2012 IEEE conference on computer vision and pattern recognition*, pages 1346–1353. IEEE, 2012. 1
- [16] 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
- [17] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3889–3898, 2019. 2
- [18] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 988–996, 2017. 2
- [19] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018. 2
- [20] Daochang Liu, Tingting Jiang, and Yizhou Wang. Completeness modeling and context separation for weakly supervised temporal action localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1298–1307, 2019. 2, 7
- [21] 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 *Proceedings of the IEEE International Conference on Computer Vision*, pages 3899–3908, 2019. 7
- [22] Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. Gaussian temporal awareness networks for action localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 344–353, 2019. 2
- [23] 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(5):907–919, 2005. 1
- [24] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *ArXiv e-prints*, Feb. 2018. 6
- [25] Kyle Min and Jason J Corso. Adversarial background-aware loss for weakly-supervised temporal activity localization. In *European Conference on Computer Vision*, pages 283–299. Springer, 2020. 7
- [26] Sanath Narayan, Hisham Cholakkal, Fahad Shahbaz Khan, and Ling Shao. 3c-net: Category count and center loss for

- weakly-supervised action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8679–8687, 2019. 1
- [27] Phuc Nguyen, Ting Liu, Gautam Prasad, and Bohyung Han. Weakly supervised action localization by sparse temporal pooling network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6752–6761, 2018. 1, 2, 3, 7
- [28] Phuc Xuan Nguyen, Deva Ramanan, and Charless C Fowlkes. Weakly-supervised action localization with background modeling. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5502–5511, 2019. 2, 3, 6, 7
- [29] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 8
- [30] Sujoy Paul, Sourya Roy, and Amit K Roy-Chowdhury. Wtalc: Weakly-supervised temporal activity localization and classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 563–579, 2018. 1, 2, 5, 7
- [31] Javier Sánchez Pérez, Enric Meinhardt-Llopis, and Gabriele Facciolo. Tv-11 optical flow estimation. *Image Processing On Line*, 2013:137–150, 2013. 6
- [32] Baifeng Shi, Qi Dai, Yadong Mu, and Jingdong Wang. Weakly-supervised action localization by generative attention modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1009–1019, 2020. 2, 5, 6, 7
- [33] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5734–5743, 2017. 2
- [34] Zheng Shou, Hang Gao, Lei Zhang, Kazuyuki Miyazawa, and Shih-Fu Chang. Autoloc: Weakly-supervised temporal action localization in untrimmed videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 154–171, 2018. 1, 5, 7
- [35] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1049–1058, 2016. 2
- [36] Krishna Kumar Singh and Yong Jae Lee. Hide-and-peek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *2017 IEEE international conference on computer vision (ICCV)*, pages 3544–3553. IEEE, 2017. 2, 7
- [37] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 2
- [38] Sarvesh Vishwakarma and Anupam Agrawal. A survey on activity recognition and behavior understanding in video surveillance. *The Visual Computer*, 29(10):983–1009, 2013. 1
- [39] Limin Wang, Yuanjun Xiong, Dahua Lin, and Luc Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4325–4334, 2017. 1, 2, 5, 7
- [40] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *Proceedings of the IEEE international conference on computer vision*, pages 5783–5792, 2017. 2, 7
- [41] Mengmeng Xu, Chen Zhao, David S Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10156–10165, 2020. 7
- [42] Tan Yu, Zhou Ren, Yuncheng Li, Enxu Yan, Ning Xu, and Junsong Yuan. Temporal structure mining for weakly supervised action detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5522–5531, 2019. 7
- [43] Yuan Yuan, Yueming Lyu, Xi Shen, Ivor W Tsang, and Dit-Yan Yeung. Marginalized average attentional network for weakly-supervised learning. *arXiv preprint arXiv:1905.08586*, 2019. 2
- [44] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7094–7103, 2019. 2, 7
- [45] Yuanhao Zhai, Le Wang, Wei Tang, Qilin Zhang, Junsong Yuan, and Gang Hua. Two-stream consensus network for weakly-supervised temporal action localization. In *European Conference on Computer Vision*, pages 37–54. Springer, 2020. 7
- [46] Peisen Zhao, Lingxi Xie, Chen Ju, Ya Zhang, Yanfeng Wang, and Qi Tian. Bottom-up temporal action localization with mutual regularization. In *European Conference on Computer Vision*, pages 539–555. Springer, 2020. 2
- [47] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2914–2923, 2017. 2, 7
- [48] Jia-Xing Zhong, Nannan Li, Weijie Kong, Tao Zhang, Thomas H Li, and Ge Li. Step-by-step erasing, one-by-one collection: a weakly supervised temporal action detector. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 35–44, 2018. 2, 7