GyF

This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

3C-Net: Category Count and Center Loss for Weakly-Supervised Action Localization

Sanath Narayan, Hisham Cholakkal, Fahad Shahbaz Khan, Ling Shao Inception Institute of Artificial Intelligence, UAE

firstname.lastname@inceptioniai.org

Abstract

Temporal action localization is a challenging computer vision problem with numerous real-world applications. Most existing methods require laborious frame-level supervision to train action localization models. In this work, we propose a framework, called 3C-Net, which only requires video-level supervision (weak supervision) in the form of action category labels and the corresponding count. We introduce a novel formulation to learn discriminative action features with enhanced localization capabilities. Our joint formulation has three terms: a classification term to ensure the separability of learned action features, an adapted multi-label center loss term to enhance the action feature discriminability and a counting loss term to delineate adjacent action sequences, leading to improved localization. Comprehensive experiments are performed on two challenging benchmarks: THUMOS14 and ActivityNet 1.2. Our approach sets a new state-of-the-art for weakly-supervised temporal action localization on both datasets. On the THU-*MOS14 dataset, the proposed method achieves an absolute* gain of 4.6% in terms of mean average precision (mAP), compared to the state-of-the-art [16]. Source code is available at https://github.com/naraysa/3c-net.

1. Introduction

Temporal action localization in untrimmed videos is a challenging problem due to intra-class variations, cluttered background, variations in video duration, and changes in viewpoints. In temporal action localization, the task is to find the start and end time (temporal boundaries or extent) of actions in a video. Most existing action localization approaches are based on strong supervision [15, 5, 21, 33, 23, 31], requiring manually annotated ground-truth temporal boundaries of actions during training. However, frame-level action boundary annotations are expensive compared to video-level action label annotations. Further, unlike object boundary annotations in images, manual annotations of temporal action boundaries are more subjective and prone to large variations [20, 18]. Here, we focus on learning to

temporally localize actions using only video-level supervision, commonly referred to as *weakly-supervised* learning.

Weakly-supervised temporal action localization has been investigated using different types of weak labels, e.g., action categories [25, 28, 14], movie scripts [12, 1] and sparse spatio-temporal points [13]. Recently, Paul et al. [16] proposed an action localization approach, demonstrating stateof-the-art results, using video-level category labels as the weak supervision. In their approach [16], a formulation based on co-activity similarity loss is introduced which distinguishes similar and dissimilar temporal segments (regions) in paired videos containing same action categories. This leads to improved action localization results. However, the formulation in [16] puts a constraint on the mini-batch, used for training, to mostly contain paired videos with actions belonging to the same category. In this work, we look into an alternative formulation that allows the mini-batch to contain diverse action samples during training.

We propose a framework, called 3C-Net, using a novel formulation to learn discriminative action features with enhanced localization capabilities using video-level supervision. As in [14, 16], our formulation contains a classification loss term that ensures the inter-class separability of learned features, for video-level action classification. However, this separability at the global video-level alone is insufficient for accurate action localization, which is generally a local temporal-context classification. This can be observed in Fig. 1, where the network trained with classification loss alone, denoted as 'CLS', localizes multiple instances of an action (central portion of the timeline) as a single instance. We therefore introduce two additional loss terms in our formulation that ensure both the discriminability of action categories at the global-level and separability of instances at the local-level.

The first additional term in our formulation is the center loss [30], introduced here for multi-label action classification. Originally designed for the face recognition problem [30], the objective of the center loss term is to reduce the intra-class variations in the feature representation of the training samples. This is achieved by learning the classspecific centers and penalizing the distance between the fea-



Figure 1. Predicted action proposals for a video clip containing *PoleVault* action category from THUMOS14 dataset. Sample frames from the video are shown in the top row. Frames containing actions have a *blue* border. GT indicates the ground-truth segments in the video containing the action. The network trained with classification loss term alone (CLS) inaccurately merges the four actions instances in the middle as a single instance. The network trained on classification and center loss terms (denoted as CLS + CL) improves the action localization but only partially delineates the merged action instances. The proposed 3C-Net framework, denoted as Ours (CLS + CL + CT), trained using a joint formulation of classification, center and counting loss terms, delineates the adjacent action instances in the middle. White regions in the timeline indicate background regions which do not contain actions of interest.

tures and their respective class centers. However, the standard center loss operates on training samples representing single-label instances. This prohibits its direct applicability in our multi-label action localization settings. We therefore propose to use a class-specific attention-based feature aggregation scheme to utilize multi-label action videos for training with center loss. As a result, a discriminative feature representation is obtained for improved localization. This improvement over 'CLS' can be observed in Fig. 1, where the network trained using the classification and center loss terms, denoted as 'CLS + CL', partially solves the incorrect grouping of multiple action instances.

The final term in our formulation is a counting loss term, which enhances the separability of action instances at the local-level. Count information has been previously exploited in the image domain for object delineation [8, 6]. In this work, the counting loss term incorporates information regarding the frequency of an action category in a video. The proposed loss term minimizes the distance between the predicted action count in a video and the ground-truth count. Consequently, the prediction scores sum up to a positive value within action instances and zero otherwise, leading to improved localization. This can be observed in Fig. 1, where the proposed 3C-Net trained using all the three loss terms, denoted as 'Ours (CLS + CL + CT)', delineates all four adjacent action instances, thereby leading to improved localization. Our counting term utilizes video-level action count and does not require user-intensive action location information (e.g. temporal boundaries).

1.1. Contributions

We introduce a weakly-supervised action localization framework, 3C-Net, with a novel formulation. Our formulation consists of a classification loss to ensure inter-class separability, a multi-label center loss to enhance the feature discriminability and a counting loss to improve the separability of adjacent action instances. The three loss terms in our formulation are jointly optimized in an end-to-end fashion. To the best of our knowledge, we are the first to propose a formulation containing center loss for multi-label action videos and counting loss to utilize video-level action count information for weakly-supervised action localization.

We perform comprehensive experiments on two benchmarks: THUMOS14 [9] and ActivityNet 1.2 [3]. Our joint formulation significantly improves the baseline containing only classification loss term. Further, our approach sets a new state-of-the-art on both datasets and achieves an absolute gain of 4.6% in terms of mAP, compared to the best existing weakly-supervised method on THUMOS14.

2. Related Work

Temporal action localization in untrimmed videos is a challenging problem that has gained significant attention in recent years. This is evident in popular challenges, such as THUMOS [9] and ActivityNet [3], where a separate track is dedicated to the problem of temporal action localization in untrimmed videos. Weakly-supervised action localization mitigates the need for temporal action boundary annotations and is therefore an active research problem. In the standard settings, only action category labels are available to train a localization model. Existing approaches have investigated different weak supervision strategies for action localization. The work of [25, 14, 28] use action category labels in videos for temporal localization, whereas [13] uses point-level supervision to spatio-temporally localize the actions. [17, 2] exploit the order of actions in a video as a weak supervision cue. The work of [12, 7] use video subtitles and movie scripts to obtain coarse temporal localization for training, while [1] utilizes actor-action pairs extracted from scripts for learning spatial actor-action localization. Recent work of [8] shows that object counting with image-level supervision is less expensive, in terms of annotation cost, compared to instance-level supervision (e.g., bounding-box). In this work, we propose to use action instance count as an additional cue for weakly-supervised action localization.

State-of-the-art weakly-supervised action localization methods utilize both appearance and motion features, typi-

cally extracted from backbone networks trained for the action recognition task. The work of [28] proposes a framework that consists of a classification and a selection module for classifying the actions and detecting the relevant temporal segments, respectively. The approach uses a two-stream Temporal Segment Network [29] as its backbone and emplovs a classification loss for training. In [14], a two-stream architecture is used to learn temporal class activation maps and a class-agnostic temporal attention. Their combination is then used to localize the human actions. Classification and sparsity-based losses are used to learn the activation maps and temporal attention, respectively. Recently, [16] proposed a framework to learn temporal localization from video-level labels, where a classification loss and a triplet loss for matching similar segments of an action category in paired videos is employed. In this work, we propose a joint formulation with explicit loss terms to ensure the separability of learned action features, enhance the feature discriminability and delineate adjacent action instances.

3. Method

In this section, we first describe the feature extraction scheme used in our approach. We then present our overall architecture followed by a detailed description of the different loss terms in the proposed formulation.

Feature Extraction: As in [14, 16], we use Inflated 3D (I3D) features extracted from the RGB and flow I3D deep networks [4], trained on the Kinetics dataset, to encode appearance and motion information, respectively. A video is divided into non-overlapping segments, each consisting of 16 frames. The input to the RGB and flow I3D networks are the color and the corresponding optical flow frames of a segment, respectively. A *D*-dimensional output I3D feature per segment, from each of the two networks, is used as input to the respective RGB and flow streams in our architecture.

3.1. Overall Architecture

Our overall 3C-Net architecture is shown in Fig. 2. In our approach, both appearance (RGB) and motion (flow) features are processed in parallel streams. The two streams are then fused at a later stage of the network. Both streams are structurally identical in design. Each stream in our network comprises of three fully-connected (FC) layers. Guided by the center loss [30], the first two FC layers learn to transform the I3D features into a discriminative intermediate feature representation. The final FC layer projects the intermediate features into the action category space under the guidance of the classification loss. The outputs of the final FC layer represent the sequence of classification scores for each action over time. This class-specific 1D representation, similar to the 2D class activation map in object detection [34], is called temporal class activation map (T-CAM), as in [14].

Given a training video v_i , let $\mathbf{y}_i \in \mathbb{R}^{N_c}$ denote the

ground-truth multi-hot vector indicating the presence or absence of an action category in v_i , where $i \in [1, N]$. Here, N is the number of videos and N_c is the number of action classes in the dataset. Let $\mathbf{x}_i^a, \mathbf{x}_i^f \in \mathbb{R}^{s_i \times D}$ denote the intermediate features (outputs of the second FC layer) in the two streams, respectively. Here, s_i denotes the length (number of segments) of the video v_i . The output of the final FC layers represent the T-CAMs, denoted by $\mathbf{C}_i^a, \mathbf{C}_i^f \in \mathbb{R}^{s_i \times N_c}$, for the RGB and flow streams, respectively. The two T-CAMs (\mathbf{C}_i^a and \mathbf{C}_i^f) are weighted by learned class-specific parameters, $\mathbf{w}^a, \mathbf{w}^f \in \mathbb{R}^{N_c}$, and later combined by addition to result in the final T-CAM, $\mathbf{C}_i^F \in \mathbb{R}^{s_i \times N_c}$. The learning of the final T-CAM, \mathbf{C}_i^F is guided by the classification and counting loss terms. Consequently, our 3C-Net framework is trained using the overall loss formulation,

$$\mathcal{L} = \mathcal{L}_{cls} + \alpha \mathcal{L}_{center} + \beta \mathcal{L}_{count} \tag{1}$$

where \mathcal{L}_{cls} , \mathcal{L}_{center} and \mathcal{L}_{count} denote the classification loss, center loss and counting loss terms, respectively. The respective weights for the center loss and counting loss terms are denoted by α and β . Next, we describe the three loss terms utilized in the proposed formulation.

3.2. Classification Loss

The classification loss term is used in our formulation to ensure the inter-class separability of the features at the video-level and tackles the problem of multi-label action classification in the video. We utilize the cross-entropy classification loss as in [28, 16], to recognize different action categories in a video. The number of segments per video varies greatly in untrimmed videos. Hence, the *top-k* values per category (where $k = \lceil s_i/8 \rceil$, is proportional to the length, s_i , of the video) of a T-CAM¹ (\mathbf{C}_i^a) are selected, as in [16]. This results in a representation of size $k \times N_c$, for the video. Further, a temporal averaging is performed on this representation to obtain a class-specific encoding, $\mathbf{r}_i^a \in \mathbb{R}^{N_c}$, for the T-CAM, \mathbf{C}_i^a . Consequently, a probability mass function (pmf), $\mathbf{p}_i^a \in \mathbb{R}^{N_c}$, is computed using

$$\mathbf{p}_{i}^{a}(j) = \frac{\exp(\mathbf{r}_{i}^{a}(j))}{\sum_{l} \exp(\mathbf{r}_{i}^{a}(l))}$$
(2)

where $j \in [1, N_c]$ denotes the action category. As shown in Fig. 2, the 'Classification Module (CLS)' performs *top*k temporal pooling, averaging and category-wise softmax operations and outputs a predicted pmf, \mathbf{p}_i^a for an input, \mathbf{C}_i^a . The multi-hot encoded ground-truth action labels \mathbf{y}_i , are l_1 -normalized to generate a ground-truth pmf, \mathbf{q}_i . The classification loss is then represented as the cross-entropy between \mathbf{p}_i^a and \mathbf{q}_i . Let $\mathcal{L}_{cls}^a = -\mathbb{E}[\mathbf{q}_i^T log(\mathbf{p}_i^a)]$ denote the classification loss for the RGB stream, where \mathbf{p}_i^a is the pmf computed from \mathbf{C}_i^a . The loss for the flow stream T-CAM

¹For brevity, the loss computation is explained in detail for the RGB stream using the superscript a (denoting appearance) for the variables.



Figure 2. Our overall architecture (3C-Net) with different loss terms (classification, center and counting), and the associated modules. The architecture is based on a two-stream model (RGB and flow) with an associated backbone feature extractor in each stream. Both streams are structurally identical and consist of two fully-connected layers (FC). The outputs of the final FC layer in both streams are the temporal class activation maps (T-CAM), \mathbf{C}^a for RGB and \mathbf{C}^f for flow. The two T-CAMs are weighted by class-specific parameters (\mathbf{w}^a and \mathbf{w}^f) and combined in a late fusion manner. The resulting T-CAM, \mathbf{C}^F , is used for inference. The modules for the different loss terms do not have learnable parameters and are shown separately in the bottom row with sample inputs and corresponding outputs for clarity. Both center $(\mathcal{L}^a_{center}, \mathcal{L}^f_{center})$ and classification $(\mathcal{L}^a_{cls}, \mathcal{L}^f_{cls})$ losses are applied to each of the two streams (\mathbf{C}^a and \mathbf{C}^f) whereas the classification (\mathcal{L}^e_{cls}) and counting (\mathcal{L}_{count}) loss are applied to the fused representation (\mathbf{C}^F). Superscripts *a*, *f* and *F* denote appearance (RGB), flow and final, respectively. Color-coded arrows denote the association between the features in the network and the respective modules.

 \mathbf{C}_{i}^{f} and the final T-CAM \mathbf{C}_{i}^{F} , are computed in a similar manner. The total classification loss, \mathcal{L}_{cls} , is then given by,

$$\mathcal{L}_{cls} = \mathcal{L}^a_{cls} + \mathcal{L}^f_{cls} + \mathcal{L}^F_{cls} \tag{3}$$

3.3. Center Loss for Multi-label Classification

We adapt and integrate the center loss term [30] in our overall formulation to cluster the features of different categories such that the same action category features are grouped together. The center loss learns the cluster centers of each action class and penalizes the distance between the features and the corresponding class centers. The objective of the classification loss, commonly employed in action localization, is to ensure the inter-class separability of learned features, whereas the center loss aims to enhance their discrminability through action-specific clustering and minimizing the intra-class variations. However, the standard center loss, originally proposed for face recognition [30], operates on training samples representing single-label instances. This hinders its usage in multi-label weakly-supervised action localization settings, where training samples (videos) contain multiple action categories. To counter this issue, we employ an attention-based per-class feature aggregation strategy to utilize videos with multiple action categories for training with the center loss. To the best of our knowledge, we are the first to introduce the center loss with multi-label training samples for weakly supervised action localization.

In the proposed 3C-Net framework, the center loss is ap-

plied on the features¹, \mathbf{x}_{i}^{a} (output of the penultimate FC layer as in Fig. 2). Typically, videos vary in length (s_{i}) and contain multiple action classes. Additionally, the action duration may be relatively short in untrimmed videos. Hence, aggregating category-specific features by considering only the high attention regions of those categories in the video is required. We perform the feature aggregation step on \mathbf{x}_{i}^{a} and compute a single feature $\mathbf{f}_{i}^{a}(j) \in \mathbb{R}^{D}$ if $\mathbf{y}_{i}(j) \neq 0$ (*i.e.*, if the action category j is present in video v_{i}). In the case of action categories which are not present in a video, the feature aggregation step is not performed, since these categories will not have a meaningful feature representation in that video. To this end, we first compute the attention, $\mathbf{a}_{i}^{a} \in \mathbb{R}^{s_{i} \times N_{c}}$, over time t, for a category j, using

$$\mathbf{a}_{i}^{a}(t,j) = \frac{\exp(\mathbf{C}_{i}^{a}(t,j))}{\sum_{l} \exp(\mathbf{C}_{i}^{a}(l,j))}$$
(4)

where \mathbf{C}_{i}^{a} represents the RGB stream T-CAM for video v_{i} . A threshold, $\tau_{j} = \text{median}(\mathbf{a}_{i}^{a}(j))$ is used to set the attention weights less than τ_{j} to 0 (*i.e.* $\mathbf{a}_{i}^{a}(t, j) = 0$, if $\mathbf{a}_{i}^{a}(t, j) < \tau_{j}$)). Here, s_{i} is the length of the video. This thresholding enables feature aggregation from category-specific high-attention regions of the video. The resulting aggregated features, $f_{i}^{a}(j)$, are then used with the center loss. The aggregated feature $f_{i}^{a}(j)$ is computed using

$$\mathbf{f}_{i}^{a}(j) = \frac{\sum_{t} \mathbf{a}_{i}^{a}(t, j) \mathbf{x}_{i}^{a}(t)}{\sum_{t} \mathbf{a}_{i}^{a}(t, j)}$$
(5)

As shown in Fig. 2, the 'Center Loss Module (CL)' implements Eq. 4 and 5 for each stream, using the outputs of the FC layers of the respective stream. Let $\mathbf{c}_j^a \in \mathbb{R}^D$ be the cluster center associated with the action category j. Following [30], the center loss and the update for center \mathbf{c}_j^a , used in our multi-label formulation, are given by,

$$\mathcal{L}_{center}^{a} = \frac{1}{N} \sum_{i} \sum_{j:\mathbf{y}_{i}(j)=1} ||\mathbf{f}_{i}^{a}(j) - \mathbf{c}_{j}^{a}||_{2}^{2}$$
(6)

$$\Delta \mathbf{c}_{j}^{a} = \frac{\sum_{i:\mathbf{y}_{i}(j)=1} (\mathbf{c}_{j}^{a} - \mathbf{f}_{i}^{a}(j))}{1 + \sum_{i} \mathbf{y}_{i}(j)}$$
(7)

For every category j, present in a mini-batch, the corresponding center, \mathbf{c}_j^a is updated using its $\Delta \mathbf{c}_j^a$ during training. The loss for the flow stream, \mathcal{L}_{center}^f , is also computed in a similar manner. The total center loss is then given by,

$$\mathcal{L}_{center} = \mathcal{L}_{center}^a + \mathcal{L}_{center}^f \tag{8}$$

3.4. Counting Loss

In this work, we propose to use auxiliary count information in addition to standard action category labels for weakly-supervised action localization. Here, count refers to the number of instances of an action category occurring in a video. As discussed earlier, integrating count information enhances the feature representation and delineation of temporally adjacent action instances in the video, leading to an improved temporal localization. In our 3C-Net framework, the counting loss is applied on the final T-CAM, C_i^F .

To compute the predicted count, first, the element-wise product of the category-specific temporal attention and the final T-CAM, C_i^F , is performed. The resulting attentionweighted T-CAM is equivalent to a density map [6] of the action category, and its summation yields the predicted count of that category. Let the attention for action category j be $a_i^F(j)$, which is computed using the final T-CAM, similar to Eq. 4. The predicted count for category j is given by,

$$\mathbf{m}_{i}(j) = \sum_{t} \mathbf{a}_{i}^{F}(t, j) \mathbf{C}_{i}^{F}(t, j)$$
(9)

where $\mathbf{m}_i(j)$ represents the sum of activation weighted by the temporal attention, over time for the j^{th} action category. As shown in Fig. 2, the 'Counting Module (CT)' implements Eq. 4 and 9 for the final T-CAM, \mathbf{C}_i^F . Temporal attention weighting ignores the background video segments not containing the action category j.

In the context of action localization, we observe that videos with a higher action count tend to have higher errors in count prediction during training. Training with absolute error results in an inferior T-CAM, since the mini-batch loss will be dominated by the count prediction error for the videos with a higher action count. To tackle this issue, we use a simple yet effective weighting strategy, where errors are inversely weighted depending on the action count in a video. A lower weight is assigned when the action count in a video is high and vice versa. The weighting penalizes the count error (*ce*) more at lower ground-truth count (GTC) compared to the same magnitude of *ce* at higher GTC. *E.g.*, *ce* = 1 at GTC of 5 is emphasized over *ce* = 1 at GTC of 100. To obtain a relative error for per-category count prediction, we divide the absolute error by the GTC of the categories present in the video. Absolute error is used for the action categories that are not present in a video to ensure that their predicted count is zero. The counting loss is then given by,

$$\mathcal{L}_{count}^{+} = \frac{1}{N} \sum_{i} \sum_{j:\mathbf{n}_{i}(j)>0} \frac{|\mathbf{m}_{i}(j) - \mathbf{n}_{i}(j)|}{\mathbf{n}_{i}(j)}$$
$$\mathcal{L}_{count}^{-} = \frac{1}{N} \sum_{i} \sum_{j:\mathbf{n}_{i}(j)=0} |\mathbf{m}_{i}(j)|$$
$$\mathcal{L}_{count} = \mathcal{L}_{count}^{+} + \lambda \mathcal{L}_{count}^{-}$$
(10)

where $\mathbf{n}_i \in \mathbb{R}^{N_c}$ is the ground-truth count label and λ is a hyper-parameter, typically set to 10^{-3} to compensate for the ratio of positive to negative instances for an action class.

To summarize, the loss terms in our overall formulation enhance the separability and discriminability of the learned features and improve the delineation of adjacent action instances. Consequently, a disrcriminative and improved T-CAM representation is obtained.

3.5. Classification and Localization using T-CAM

After training the 3C-Net, the CLS module (see Fig. 2 and Eq. 2) is used to compute the action-class scores (pmf) at the video-level using the final T-CAM, for the action classification task. Similar to [28, 16], we use the computed pmf without threshold, for evaluation. For the action localization task, detections are obtained using a similar approach used in [16]. Detections in a video are generated for the action categories with average top-k score above 0 (i.e. for categories in set $\{j : \mathbf{r}_i^F(j) > 0\}$, where \mathbf{r}_i^F is computed as in Sec. 3.2 using the final T-CAM). For a category j in the obtained set, continuous video segments between successive time instants when T-CAM goes above and below threshold η , correspond to a valid action detection. The resulting detections of an action category are non-overlapping. A weighted sum of the highest T-CAM value with in the detection and the category score for the video, corresponds to the score of a detection. The detection with the highest score that is overlapping (above IoU threshold) with the ground-truth is considered true-positive during evaluation.

4. Experiments

4.1. Experimental Setup

Datasets: The proposed 3C-Net is evaluated for temporal action localization on two challenging datasets containing

untrimmed videos with varying degree of activity duration. **THUMOS14** [9] dataset contains 1010 validation and 1574 test videos from 101 action categories. Out of these, 20 categories have temporal annotations in 200 validation and 213 test videos. The dataset is challenging, as it contains an average of 15 activity instances per video. Similar to [14, 16], we use the validation set for training and test set for evaluating our framework. **ActivityNet 1.2** [3] dataset has 4819 training, 2383 validation and 2480 testing videos from 100 activity categories. Note that the test set annotations for this dataset are withheld. There are an average of 1.5 activity instances per video. As in [22, 16], we use the training set to train and validation set to test our approach.

Count Labels: The ground-truth count labels for the videos in both datasets are generated using the available temporal action segments information. The total number of segments of an action category in a video is the ground-truth count video-label for the respective category. This was done to use the available annotations and avoid re-annotations. However, for a new dataset, action count can be independently annotated, without requiring action segment information.

Evaluation Metric: We follow the standard protocol, provided with the two datasets, for evaluation. The evaluation protocol is based on mean Average Precision (mAP) for different intersection over union (IoU) values for the action localization task. For the multi-label action classification task, we use the mAP computed from the predicted video-level scores for evaluation.

Implementation Details: We use an alternate mini-batch training approach to train the proposed 3C-Net framework. Since, the count labels are available at the video-level, all the segments of a video are required for count prediction. We use random temporal cropping of videos in alternate mini-batches to improve the generalization. Thus, the classification and center losses are used for every mini-batch training and the counting loss is applied *only* on the alternate mini-batches containing the full-length video features.

In our framework, a TV-L1 optical flow [32] is used to generate the optical flow frames of the video. The I3D features of size D = 1024 per segment of 16 video frames are obtained after spatio-temporal average pooling of *Mixed_5c* layers from the RGB and Flow I3D networks. These I3D features are then used as input to our framework. As in [14, 16], the backbone networks are not finetuned. Our 3C-Net is trained with a mini-batch size of 32 using the Adam [11] optimizer with 10^{-4} learning rate and 0.005 weight decay. The centers c_i are learned using the SGD optimizer with 0.1 learning rate. For both datasets, we set α in Eq. 1 to 10^{-3} since the center loss penalty is a squared error loss with a higher magnitude compared to other loss terms. We set β in Eq. 1 to 1 and 0.1 for the THUMOS14 and ActivityNet 1.2 datasets, respectively. η is set to $0.5[\min(\mathbf{C}_i^F(j)) + \max(\mathbf{C}_i^F(j))]$ for a j^{th} category

Approach	mAP @ IoU					
Approach	0.1	0.2	0.3	0.4	0.5	0.7
FV-DTF [15]+	36.6	33.6	27.0	20.8	14.4	-
S-CNN [23]+	47.7	43.5	36.3	28.7	19.0	5.3
CDC [21] ⁺	-	-	40.1	29.4	23.3	7.9
R-C3D [31]+	54.5	51.5	44.8	35.6	28.9	-
TAL-Net [5] ⁺	59.8	57.1	53.2	48.5	42.8	20.8
UntrimmedNets [28]	44.4	37.7	28.2	21.1	16.2	5.1
STPN [14]	52.0	44.7	35.5	25.8	16.9	4.3
Autoloc [22]	-	-	35.8	29.0	21.2	5.8
W-TALC [16]	53.7	48.5	39.2	29.9	22.0	7.3
Ours: CLS + CL	56.8	49.8	40.9	32.3	24.6	7.7
Ours: 3C-Net	59.1	53.5	44.2	34.1	26.6	8.1

Table 1. Action localization performance comparison (mAP) of our 3C-Net with state-of-the-art methods on THUMOS14 dataset. Superscript '+' for a method denotes that strong supervision is required for training. Our 3C-Net outperforms existing weaklysupervised methods and achieves an absolute gain of 4.6%, at IoU=0.5, compared to the best weakly-supervised result [16].

T-CAM in THUMOS14. Due to the nature of actions in ActivityNet 1.2, W-TALC [16] approach uses the Savitzky-Golay filter [19] for post-processing the T-CAMs. Here, we use a learnable temporal convolution filtering (kernel size=13, dilation=2) and set η to 0.

4.2. State-of-the-art comparison

Temporal Action Localization: Tab. 1 shows the comparison of our 3C-Net method with existing approaches in literature on the THUMOS14 dataset. Superscript '+' for a method in Tab. 1 denotes that frame-level labels (strong supervision) are required for training. Our approach is denoted as '3C-Net'. We report mAP scores at different IoU thresholds. Both UntrimmedNets [28] and Autoloc [22] use TSN [29] as the backbone, whereas STPN [14] and W-TALC [16] use I3D networks similar to our framework. The STPN approach obtains an mAP of 16.9 at IoU=0.5, while W-TALC achieves an mAP of 22.0. Our approach CLS + CL, without any count supervision, outperforms all existing weakly-supervised action localization approaches. With the integration of count supervision, our 3C-Net achieves an absolute gain of 4.6%, in terms of mAP at IoU=0.5, over W-TALC [16]. Further, a consistent improvement in performance is also obtained at other IoU thresholds.

Tab. 2 shows the state-of-the-art comparison on the ActivityNet 1.2 dataset. We follow the standard evaluation protocol [3] by reporting the mean mAP scores at different thresholds (0.5:0.05:0.95). Among the existing methods, the SSN approach [33] relies on frame-level annotations (strong supervision, denoted by superscript '+' in Tab. 2) for training and achieves a mean mAP score of 26.6. Our baseline approach, trained with the classification loss alone, achieves a mean mAP of 18.2. With only the center loss adaption, our approach achieves a mean mAP of 21.1 and surpasses all existing weakly-supervised methods.

Approach	mAP @ IoU			
Арргоасн	0.5	0.7	0.9	Avg*
SSN [33] ⁺	41.3	30.4	13.2	26.6
UntrimmedNets [28]	7.4	3.9	1.2	3.6
Autoloc [22]	27.3	17.5	6.8	16.0
W-TALC [16]	37.0	14.6	-	18.0
Ours: CLS + CL	35.4	22.9	8.5	21.1
Ours: 3C-Net	37.2	23.7	9.2	21.7

Table 2. Action localization performance comparison (mean mAP) of our 3C-Net with state-of-the-art methods on the ActivityNet 1.2 dataset. The mean mAP is denoted by Avg*. Note that SSN [33] requires frame-level labels (strong supervision) for training. Our 3C-Net outperforms all existing weakly-supervised methods and obtains an absolute gain of 3.7% in terms of mean mAP, compared to the state-of-the-art weakly-supervised W-TALC [16].

Approach	THUMOS14	ActivityNet 1.2
iDT+FV [27]	63.1	66.5
Objects + Motion [10]	71.6	-
Two Stream [24]	66.1	71.9
C3D [26]	-	74.1
TSN [29]	67.7	88.8
UntrimmedNets [28]	82.2	87.7
W-TALC [16]	85.6	93.2
Ours: 3C-Net	86.9	92.4

Table 3. Action classification performance comparison (mAP) of our 3C-Net with state-of-the-art methods on the THUMOS14 and ActivityNet 1.2 datasets. On THUMOS14, our 3C-Net achieves superior classification result, compared to existing methods.

With the integration of count supervision, the performance further improves to 21.7 and outperforms the state-of-theart weakly-supervised approach [16] by 3.7%, in terms of mean mAP. The relatively lower margin of improvement using count labels, compared to THUMOS14, is likely due to fewer multi-instance videos in training and noisy annotations in this dataset.

Action Classification: We also evaluate our method for action classification. Tab. 3 shows the comparison on THU-MOS14 and ActivityNet 1.2 datasets. Our 3C-Net achieves a superior classification performance of 86.9, in terms of mAP, compared to existing methods on the THUMOS14 dataset and is comparable to W-TALC on ActivityNet 1.2.

4.3. Baseline Comparison and Ablation Study

Baseline comparison: Tab. 4 shows the action localization performance comparison on THUMOS14 (at IoU=0.5). We also show the impact of progressively integrating one contribution at a time in our 3C-Net framework. The baseline (CLS) trained using classification loss alone obtains a mAP score of 19.1. The integration of our multi-label center loss term (CLS + CL) significantly improves the performance by obtaining a mAP score of 24.6. The action localization performance is further improved to 26.6 mAP, by the integration of our counting loss term (CLS + CL + CT).

Ablation study: Fig. 3 shows the results with respect to

Baseline: CLS	CLS + CL	3C-Net: CLS + CL + CT		
19.1	24.6	26.6		

Table 4. Baseline action localization performance comparison (mAP) on THUMOS14 at IoU=0.5. Our 3C-Net achieves an absolute gain of 7.5% in terms of mAP, compared to the baseline.



Figure 3. Ablation study with respect to difference design choices and different loss terms in our action localization framework on the THUMOS14 dataset. See text for details.

different design choices and impact of different loss terms in our action localization framework on the THUMOS14 dataset. All the experiments are conducted independently and show the deviation in performance relative to the proposed 3C-Net framework. The localization performance of our final proposed 3C-Net framework is shown as yellow bar. First, we show the impact of removing the classification loss in both the streams and retaining it only for the final T-CAM (\mathbf{C}_{i}^{F}). This results (orange bar) in a drop of 2.5% mAP. Next, we observe that retaining center loss term only in the flow stream results in a drop of 2.1% mAP (purple bar). Retaining the center loss term only in the RGB stream results in a drop of 1.9% mAP (green bar). Afterwards, we observe that removing the negative category counting loss in Eq. 10 results in a drop of 1.5% mAP (blue bar). Further, replacing the relative error for counting loss with absolute error deteriorates the results by 1.2% mAP (red bar). These results show that both our design choices and different loss terms contribute in the overall performance of our approach.

4.4. Qualitative Analysis

We now present the qualitative analysis of our 3C-Net approach. Fig. 4 shows the qualitative temporal action localization results of our 3C-Net on example videos from the THUMOS14 and ActivityNet 1.2 datasets. For each video, example frames are shown in the top row. GT denotes the ground-truth segments. The category-specific confidence scores over time are indicated by T-CAM. Detection denotes the action segments predicted using the T-CAM. The top two videos are from THUMOS14. The multiple instances of *HighJump* action (first video) are accurately localized by our 3C-Net. The second video contains visually similar multiple actions (*Shotput* and *ThrowDiscus*) and has overlapping ground-truth annotations. In this case, 3C-Net mostly localizes the two actions accurately.

The bottom two examples from the ActivityNet 1.2



Figure 4. Qualitative temporal action localization results of our 3C-Net approach on example videos from the THUMOS14 and ActivityNet 1.2 datasets. For each video, we show the example frames in the top row, ground-truth segments indicating the action instances as GT and the class-specific confidence scores over time as T-CAM (for brevity, only the thresholded T-CAM is shown). Action segments predicted using the T-CAM are denoted as Detection. Examples show different scenarios: multiple instances of same action (first video), visually-similar multiple action categories (second video) and long duration activities (third and fourth video). Our approach achieves promising localization performance on these variety of actions.

dataset contain long duration activities from Playing Violin and Parallel Bars categories. Observing the T-CAM progression in both videos, we see that the proposed framework detects the action instances reasonably well. For Playing *Violin* video, prediction with respect to the second instance is correctly detected, while the first instance is partially detected. This is due to imprecise annotation of the first instance which has some segments without the playing activity. In Parallel Bars video, a single action instance is annotated. However, the video contains an activity instance followed by background segments without any action and ends with the replay of the first action instance. This progression of activity-background-activity has been clearly identified by our approach as observed in the T-CAM. These results suggest the effectiveness of our approach for the problem of temporal action localization. We observe common failure reasons to be extreme scale change, visually similar actions confusion and temporally quantized segments for I3D feature generation. Few failure instances in Fig. 4 are: detections having minimal overlap with the GT (first two detected instances of *ThrowDiscus*), false detections (third and fourth detected instances of *ThrowDiscus*) and multiple detections (first two detected instances of *Parallel Bars*).

5. Conclusion

We proposed a novel formulation with classification loss, center loss and counting loss terms for weakly-supervised action localization. We first proposed to use a class-specific attention-based feature aggregation strategy to utilize multilabel videos for training with center loss. We further introduced a counting loss term to leverage video-level action count information. To the best of our knowledge, we are the first to propose a formulation with multi-label center loss and action counting loss terms for weakly-supervised action localization. Experiments on two challenging datasets clearly demonstrate the effectiveness of our approach for both action localization and classification.

References

- Piotr Bojanowski, Francis Bach, Ivan Laptev, Jean Ponce, Cordelia Schmid, and Josef Sivic. Finding actors and actions in movies. In *ICCV*, 2013. 1, 2
- [2] Piotr Bojanowski, Rémi Lajugie, Francis Bach, Ivan Laptev, Jean Ponce, Cordelia Schmid, and Josef Sivic. Weakly supervised action labeling in videos under ordering constraints. In ECCV, 2014. 2
- [3] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 2, 6
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 3
- [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 *CVPR*, 2018. 1, 6
- [6] Hisham Cholakkal, Guolei Sun, Fahad Shahbaz Khan, and Ling Shao. Object counting and instance segmentation with image-level supervision. In *CVPR*, 2019. 2, 5
- [7] Olivier Duchenne, Ivan Laptev, Josef Sivic, Francis Bach, and Jean Ponce. Automatic annotation of human actions in video. In *ICCV*, 2009. 2
- [8] Mingfei Gao, Ang Li, Ruichi Yu, Vlad I Morariu, and Larry S Davis. C-wsl: Count-guided weakly supervised localization. In ECCV, 2018. 2
- [9] 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. *CVIU*, 2017. 2, 6
- [10] Mihir Jain, Jan C Van Gemert, and Cees GM Snoek. What do 15,000 object categories tell us about classifying and localizing actions? In *CVPR*, 2015. 7
- [11] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6
- [12] Ivan Laptev, Marcin Marszałek, Cordelia Schmid, and Benjamin Rozenfeld. Learning realistic human actions from movies. In *CVPR*, 2008. 1, 2
- [13] Pascal Mettes, Jan C Van Gemert, and Cees GM Snoek. Spot on: Action localization from pointly-supervised proposals. In *ECCV*, 2016. 1, 2
- [14] Phuc Nguyen, Ting Liu, Gautam Prasad, and Bohyung Han. Weakly supervised action localization by sparse temporal pooling network. In *CVPR*, 2018. 1, 2, 3, 6
- [15] Dan Oneata, Jakob Verbeek, and Cordelia Schmid. The lear submission at thumos 2014. 2013. 1, 6
- [16] Sujoy Paul, Sourya Roy, and Amit K Roy-Chowdhury. Wtalc: Weakly-supervised temporal activity localization and classification. In *ECCV*, 2018. 1, 3, 5, 6, 7
- [17] Alexander Richard, Hilde Kuehne, and Juergen Gall. Weakly supervised action learning with rnn based fine-to-coarse modeling. In *CVPR*, 2017. 2
- [18] Scott Satkin and Martial Hebert. Modeling the temporal extent of actions. In ECCV, 2010. 1

- [19] Abraham Savitzky and Marcel JE Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8):1627–1639, 1964. 6
- [20] Konrad Schindler and Luc Van Gool. Action snippets: How many frames does human action recognition require? In *CVPR*, 2008. 1
- [21] 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, 2017. 1, 6
- [22] Zheng Shou, Hang Gao, Lei Zhang, Kazuyuki Miyazawa, and Shih-Fu Chang. Autoloc: weakly-supervised temporal action localization in untrimmed videos. In *ECCV*, 2018. 6, 7
- [23] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, 2016. 1, 6
- [24] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014. 7
- [25] 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*, 2017. 1, 2
- [26] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 7
- [27] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *ICCV*, 2013. 7
- [28] Limin Wang, Yuanjun Xiong, Dahua Lin, and Luc Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *CVPR*, 2017. 1, 2, 3, 5, 6, 7
- [29] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 3, 6, 7
- [30] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In ECCV, 2016. 1, 3, 4, 5
- [31] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *ICCV*, 2017. 1, 6
- [32] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-1 1 optical flow. In *Joint pattern recognition symposium*, 2007. 6
- [33] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, 2017. 1, 6, 7
- [34] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 3