

AutoLoc: Weakly-supervised Temporal Action Localization in Untrimmed Videos

Zheng Shou¹, Hang Gao¹, Lei Zhang², Kazuyuki Miyazawa³, and Shih-Fu Chang¹

¹ Columbia University, New York, NY, USA

² Microsoft Research, Redmond, Washington, USA

³ Mitsubishi Electric, Japan

Abstract. Temporal Action Localization (TAL) in untrimmed video is important for many applications. But it is very expensive to annotate the segment-level ground truth (action class and temporal boundary). This raises the interest of addressing TAL with weak supervision, namely only video-level annotations are available during training). However, the state-of-the-art weakly-supervised TAL methods only focus on generating good Class Activation Sequence (CAS) over time but conduct simple thresholding on CAS to localize actions. In this paper, we first develop a novel weakly-supervised TAL framework called AutoLoc to directly predict the temporal boundary of each action instance. We propose a novel Outer-Inner-Contrastive (OIC) loss to automatically discover the needed segment-level supervision for training such a boundary predictor. Our method achieves dramatically improved performance: under the IoU threshold 0.5, our method improves mAP on THUMOS'14 from 13.7% to 21.2% and mAP on ActivityNet from 7.4% to 27.3%. It is also very encouraging to see that our weakly-supervised method achieves comparable results with some fully-supervised methods.

Keywords: Temporal Action Localization; Weak Supervision; Outer-Inner-Contrastive; Class Activation Sequence

1 Introduction

Impressive improvement has been made in the past two years to address **Temporal Action Localization (TAL)** in untrimmed videos [23,15,46,67,68,51,50,36,22,73,17,66,11,10,16,7,6,69,52]. These methods were proposed for the **fully-supervised** setting: the model training requires the full annotation of the ground truth temporal boundary (start time and end time) for each action instance. However, untrimmed videos are usually very long with substantial background content in time. Therefore, manually annotating temporal boundaries for a new large-scale dataset is very expensive and time-consuming [72], and thus might prohibit applying the fully-supervised methods to the new domains that lack enough training data with full annotations.

This motivates us to develop TAL methods that require significantly fewer ground truth annotations for training. As illustrated in Fig. 1, in this paper

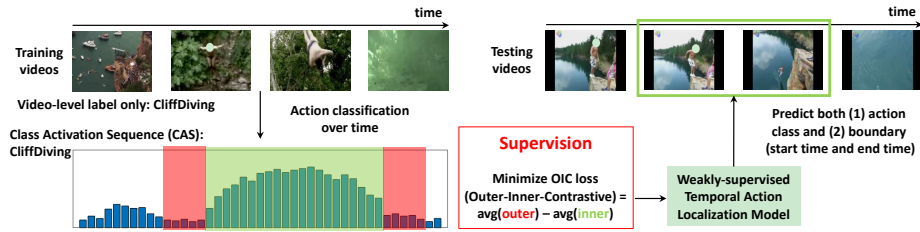


Fig. 1. We study the weakly-supervised temporal action localization problem: during training we only have videos with the video-level labels, but during testing we need to predict both (1) the action class and (2) the temporal boundary of each action instance. In order to obtain the segment-level supervision for training the action localization model to predict the boundary directly, we design a novel **Outer-Inner-Contrastive (OIC)** loss based on the action Class Activation Sequence. We denote the predicted action segment boundary as the *inner boundary*. The *outer boundary* is obtained by extending the inner boundary to include its surrounding area. A desirable boundary prediction should have high activations in the inner **green** area but low activations in the outer **red** area. Consequently, the OIC loss can be used to approximately determine the needed segment-level supervision for training the localization model

we focus on the following scenario: during training, we only have the video-level labels, which are much easier to collect, compared to the boundary annotations; during testing, we still aim to predict both (1) the action class and (2) the temporal boundary (i.e. start time and end time) of each action instance. We refer this scenario as the **weakly-supervised** setting that this paper works on.

Recently, a few methods have been proposed to tackle TAL in such a weakly-supervised setting. UntrimmedNet [64] and Hide-and-Seek [56] achieve the state-of-the-art performances and carry out the localization in a similar manner. Given a training video, several segments are randomly sampled and are fed into a network together to yield a video-level class prediction. During testing, the trained network is slid over time to produce the classification score sequence of being each action over time. The score sequence is similar to the Class Activation Map in [74] but just has one dimension, and thus we refer it as **Class Activation Sequence (CAS)**. Finally a simple thresholding method is applied on the CAS to localize each action instance in terms of the start time and the end time.

However, performing localization via thresholding in general may not be robust to noises in CAS: sometimes there are a few dips of low activations within an interval of high activations, using a large threshold might over-segment one action instance into several segments; but using a small threshold might include too many irrelevant backgrounds preceding and succeeding the action instance. One possible solution is improving the quality of CAS. Alternatively, instead of thresholding, many fully-supervised TAL methods detect action instances at the segment-level directly [51,6]. Some works further employ boundary regression models to learn to predict more accurate boundaries [36,17,66,16]. Thus, we de-

sign a framework called **AutoLoc** which can conduct direct boundary prediction via predicting the center location and the duration of each action instance.

But how to train the boundary prediction model without ground truth boundary annotations still remains unsolved. To address this challenge, we propose a novel **Outer-Inner-Contrastive (OIC)** loss to provide the needed segment-level supervision for training the boundary prediction model. Given the CAS of being the ground truth action, we denote the *inner boundary* as the boundary of a predicted action instance and we inflate the inner boundary slightly to obtain the *outer boundary*. As illustrated in Fig. 1, we propose an OIC loss as the average activation in the outer **red** area minus the average activation in the inner **green** area. By minimizing the OIC loss to find the area of high inner activations but low outer activations, we can make desirable localization of the salient interval on CAS, which is likely to be well-aligned with the ground truth segment. Equipped with the OIC loss, AutoLoc can automatically discover the segment-level supervision from the video-level annotations for training the boundary prediction model. In Sec. 5, we will experimentally compare with the state-of-the-art methods and also study several variants of our model.

In summary, we make three novel contributions in this paper:

(1) To the best of our knowledge, AutoLoc is the first weakly-supervised TAL framework that can directly predict the temporal boundary of each action instance with only the video-level annotations available during training, specifically addressing the localization task at the segment level.

(2) To enable the training of such a parametric boundary prediction model, we design a novel OIC loss to automatically discover the segment-level supervision and we prove that the OIC loss is differentiable to the underlying boundary prediction model.

(3) We demonstrate the effectiveness of AutoLoc on two standard benchmarks. AutoLoc significantly outperforms the state-of-the-art weakly-supervised TAL methods and even achieves results comparable to some fully-supervised methods that use the boundary annotations during training. When the overlap IoU threshold is set to 0.5 during evaluation, our method improves mAP on THUMOS'14 from 13.7% to 21.2% (54.7% relative gain) and improves mAP on ActivityNet from 7.4% to 27.3% (268.9% relative gain).

2 Related Works

Video Action Analysis Detailed reviews can be found in recent surveys [65,42,2,9,3,31]. Researchers have developed quite a few deep networks for video action analysis such as 3D ConvNets [60,27,61], LSTM [13], two-stream network [55], I3D [8], etc. For example, Wang et al. proposed Temporal Segment Network [63], which employed the two-stream network to model the long-range temporal structure in video and served as an effective backbone network in various video analysis tasks such as recognition [63], localization [73], weakly-supervised learning [64].

Temporal Action Localization with Weak Supervision and Full Supervision Several large-scale video datasets have been created for TAL such as Charades [54,53], ActivityNet [21], THUMOS [29,19]. In order to obtain the ground truth temporal boundaries to provide full supervision for training the fully-supervised TAL models, substantial efforts are required for annotating each of such large-scale datasets. Therefore, it is useful and important to develop TAL models that can be trained with weak supervision only.

Video-level annotation is one kind of weak supervision that can be more easily collected and thus is quite interesting to the community. Sun et al. [58] was the first to consider TAL with only the video-level annotations available during training and the authors discovered the additional supervision from web images. Recently, Singh et al. designed Hide-and-Seek [56] to address the challenge that weakly-supervised detection methods usually focus on the most discriminative parts while neglect other relevant parts of the target instance. Wang et al. [64] proposed a framework called UntrimmedNet consisting of a classification module to perform action classification and a selection module to detect important temporal segments. These recent methods are effectively learning an action classification model during training in order to generate reasonably good Class Activation Sequence (CAS) over time. But in order to detect temporal boundaries, a simple thresholding is applied on the CAS during testing. Therefore, although these methods can excel at the video-level action recognition, the performance of temporal localization still has large room for improvement.

However, the fully-supervised TAL methods (boundary annotations available during training) have gone beyond the simple thresholding method. First, some researchers performed localization at segment-level: they first generated the candidate segments via sliding window or proposal methods, and then they classified each segment into certain actions [51,17,66,16,7]. Motivated by the success of single-shot object detection method [38,44,43], Lin et al. [36] removed the proposal stage and directly conducted TAL in a single-shot fashion to simultaneously predict temporal boundary and action class. Second, direct boundary prediction via anchor generation and boundary regression has been adapted from object detection [38,44,43,45,18] to fully-supervised TAL recently and proven to be quite effective in detecting more accurate boundaries [36,73,17,66,16]. This motivates us to generalize segment-level localization and direct boundary prediction to weakly-supervised TAL: we develop AutoLoc to first generate anchor segments and then regress their boundaries to obtain the predicted segments; in order to train the boundary regressors, we propose the OIC loss to provide the segment-level supervision.

Weakly-supervised Deep Learning Methods Other types of weak supervision for action detection have also been explored in the past. For instance, Huang et al. [25] and Richard et al. [47] both utilized the order of actions as the supervision used during training. Mettes et al. [39] worked on the spatio-temporal action detection using only the point-level supervision for training.

Weakly-supervised deep learning methods have been also widely studied in other vision tasks such as object detection [74,75,49,35,30,14,59,71,57,5,32,20], semantic segmentation [34,24,41,4], video captioning [48], visual relation detection [70], etc. As a counterpart of the weakly-supervised video TAL, the weakly-supervised image object detection has been significantly improved via combining Multiple Instance Learning (MIL) [12] and deep networks [49,30,59,5,32]: built upon Fast-RCNN [18], these methods first generated candidate proposals beforehand; then they employed deep networks to classify each proposal and the scores from all proposals were fused together to obtain one label prediction for the whole image to be compared with the image-level label. One of such MIL-based deep networks is ContextLocNet [32], which further inflated the prediction box to obtain its outer box to take into account the contextual information. Our work bypasses the costly proposal generation and predicts the boundaries from raw input videos in a single-shot fashion. Although we focus on video TAL in this paper, it would be also interesting to adapt our method for image object detection in the future.

3 Outer-Inner-Contrastive Loss

In this Section, we formulate how to compute the proposed OIC loss during the network forward pass of AutoLoc and prove that the OIC loss is differentiable to the underlying boundary prediction model during the backward pass. The whole pipeline and details of AutoLoc will be presented in Sec. 4.

3.1 Forward

As illustrated by the bottom-right part in Fig. 2, for each predicted segment ϕ , we can compute its OIC loss. Each predicted segment ϕ consists of the action/inner boundary $[x_1, x_2]$, the inflated outer boundary $[X_1, X_2]$, and the action class k . These boundaries are at the snippet-level granularity (for example, boundary $x = 1$ corresponds to the location of the 1-st snippet). In order to fetch the corresponding snippet-level activation on the CAS, we round each boundary of continuous value to its nearest integer (i.e. the location of the nearest snippet). We denote the class activation at the x -th snippet on the CAS of action k as $f_k(x)$. The OIC loss of the prediction ϕ is defined as the average activation $A_o(\phi)$ in the outer area minus the average activation $A_i(\phi)$ in the inner area:

$$\mathcal{L}_{\text{OIC}}(\phi) = A_o(\phi) - A_i(\phi) = \frac{\int_{X_1}^{X_2} f_k(u) du - \int_{x_1}^{x_2} f_k(u) du}{\underbrace{(X_2 - X_1 + 1) - (x_2 - x_1 + 1)}_{A_o(\phi)}} - \frac{\int_{x_1}^{x_2} f_k(u) du}{\underbrace{(x_2 - x_1 + 1)}_{A_i(\phi)}}. \quad (1)$$

During training, we set k to the ground truth action and we minimize $\mathcal{L}_{\text{OIC}}(\phi)$ to encourage high activations inside and penalize high activations outside.

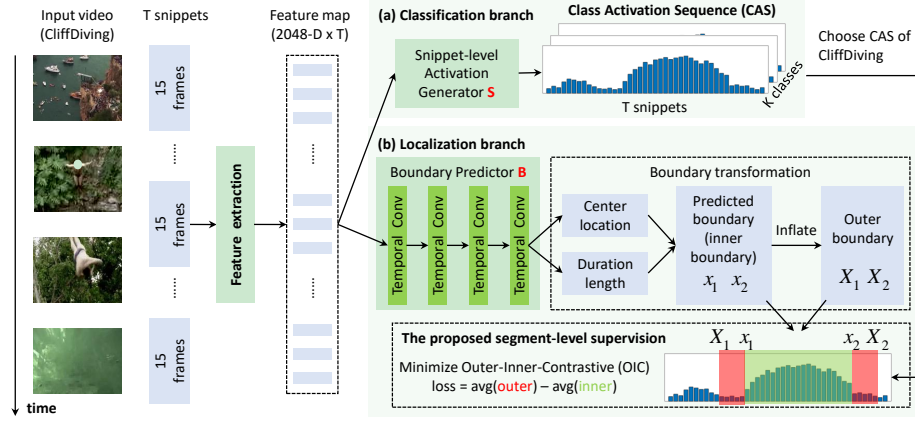


Fig. 2. The network architecture of AutoLoc. Given an input video during training, the video is chunked into T 15-frames-long snippets without overlap. We extract a feature vector of D dimension for each snippet. On top of the features, AutoLoc slides two separate branches over time: one classification branch for predicting the action scores of each snippet to obtain the **Class Activation Sequence (CAS)**; one localization branch for directly predicting the true action boundary which is denoted as the inner boundary and is inflated to obtain the outer boundary. Based on the CAS of the ground truth video-level action, we can compute the Outer-Inner-Contrastive loss (*the average activation in the outer red area minus the average activation in the inner green area*) to provide the needed segment-level supervision for training the boundary predictor

3.2 Backward

We prove that the OIC loss is differentiable to the inner and outer boundaries. Therefore, the supervision discovered by the OIC loss can be back-propagated to the underlying boundary prediction model. Detailed derivation can be found in the supplementary material. The gradients corresponding to the predicted segment ϕ w.r.t its inner boundary $[x_1, x_2]$ are as follows:

$$\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial x_1} = \frac{f_k(x_1) - A_o(\phi)}{\underbrace{(X_2 - X_1 + 1) - (x_2 - x_1 + 1)}_{\frac{\partial A_o(\phi)}{\partial x_1}}} - \frac{A_i(\phi) - f_k(x_1)}{\underbrace{(x_2 - x_1 + 1)}_{\frac{\partial A_i(\phi)}{\partial x_1}}}, \quad (2)$$

$$\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial x_2} = \frac{A_o(\phi) - f_k(x_2)}{\underbrace{(X_2 - X_1 + 1) - (x_2 - x_1 + 1)}_{\frac{\partial A_o(\phi)}{\partial x_2}}} - \frac{f_k(x_2) - A_i(\phi)}{\underbrace{(x_2 - x_1 + 1)}_{\frac{\partial A_i(\phi)}{\partial x_2}}}. \quad (3)$$

The gradients corresponding to the predicted segment ϕ w.r.t its outer boundary $[X_1, X_2]$ are as follows:

$$\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial X_1} = \frac{\partial A_o(\phi)}{\partial X_1} = \frac{A_o(\phi) - f_k(X_1)}{(X_2 - X_1 + 1) - (x_2 - x_1 + 1)}; \quad (4)$$

$$\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial X_2} = \frac{\partial A_o(\phi)}{\partial X_2} = \frac{f_k(X_2) - A_o(\phi)}{(X_2 - X_1 + 1) - (x_2 - x_1 + 1)}. \quad (5)$$

Note that these gradients indeed have the physical meanings about how to adjust the boundaries. For example, in Equation 2, $\frac{\partial A_i(\phi)}{\partial x_1}$ represents how much the average inner activation $A_i(\phi)$ is higher than the activation $f_k(x_1)$ at the inner left boundary x_1 . If the average inner activation is much higher than the activation at the inner left boundary x_1 , x_1 is likely to belong to the background and thus we would like to move x_1 in the positive (right) direction. Similarly, $\frac{\partial A_o(\phi)}{\partial x_1}$ represents how much the activation at the inner left boundary x_1 is higher than the average outer activation. $\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial x_1}$ is the adversarial outcome of $\frac{\partial A_o(\phi)}{\partial x_1}$ and $\frac{\partial A_i(\phi)}{\partial x_1}$. Consequently, $\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial x_1}$ indicates how the model wants to adjust the inner left boundary x_1 eventually: if $\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial x_1} < 0$, x_1 moves in the positive (right) direction; if $\frac{\partial \mathcal{L}_{\text{OIC}}(\phi)}{\partial x_1} > 0$, x_1 moves in the negative (left) direction.

4 AutoLoc

In this Section, we walk through the pipeline of AutoLoc as illustrated in Fig. 2. The training and testing pipelines are very similar in AutoLoc. So we only explicitly distinguish the training and testing pipelines when any difference appears.

4.1 Input Data Preparation and Feature Extraction

Each input data sample fed into AutoLoc is one single untrimmed video. Following UntrimmedNet [64], for each input video, we first divide it into 15-frames-long snippets without overlap and extract feature for each snippet individually.

In particular, Temporal Segment Network (TSN) [63] is a state-of-the-art two-stream network for video analysis. UntrimmedNet [64] has been proven to be effective in training TSN classifier with only the video-level labels. Therefore, we first train an UntrimmedNet network (the soft version) in advance and then use the trained network as our backbone for feature extraction.

This backbone network consists of one spatial stream accepting RGB input and one temporal stream accepting Optical Flow input. For each stream, we employ the Inception network architecture with Batch Normalization [26] and extract the 1024-dimensional feature at the `global_pool` layer. Finally, for each snippet, we concatenate the extracted spatial feature and temporal feature into one feature vector of 2048 dimensions. For each input video of T snippets in total, we obtain a feature map of shape 2048 (channels) by T (snippets).

4.2 Classification Branch

The goal of the classification branch is to obtain the Class Activation Sequence (CAS). We build our Activation Generator **S** based on UntrimmedNet. On top of the `global_pool` layer, UntrimmedNet attaches one Fully Connected (FC) layer of K nodes to classify each snippet into K action categories and also attaches another Fully Connected (FC) layer of just 1 node to predict the attention score (importance) for each snippet. The corresponding scores from the spatial stream and the temporal stream are averaged to obtain the final score. For each video, we use these two FC layers in the UntrimmedNet that are trained beforehand to respectively extract a classification score sequence of shape K (actions) by T (snippets) and an attention score sequence of T dimensions. For each snippet, we set its classification scores of all classes to 0 when its attention score is lower than the threshold (7 is chosen via grid search on the THUMOS'14 training set and also works well on ActivityNet); then we regard such a gated classification score as the activation, which ranges within $[0, 1]$. Finally, for each video, we obtain its CAS of shape K (actions) by T (snippets).

4.3 Localization Branch

Overview. The goal of the localization branch is to learn a parametric model for predicting the segment boundary directly. Recent fully-supervised TAL methods [36,73,17,66,16] have shown the effectiveness of regressing anchors for direct boundary prediction: the anchor is a hypothesis of the possible segment; the predicted boundary is obtained by respectively regressing (1) the center location and (2) the temporal length of the anchor segment; multi-anchor mechanism is used to cover the possible segments of different temporal scales. Therefore, we design a localization network **B** to look at each temporal position on the feature map and output the needed two boundary regression values for each anchor. Then we regress the anchors using these regression values to obtain the predicted action boundaries (inner boundaries) and inflate the inner boundaries to obtain the outer boundaries. Finally, based on the CAS, we introduce an OIC layer equipped with the OIC loss to generate the final segment predictions.

Network Architecture of the Localization Network **B** Given an input video, its feature map of shape 2048 channels by T snippets is fed into **B**. **B** first stacks 3 same temporal convolutional layers, which slide convolutional filters over time. Each temporal convolutional layer has 128 filters, which all have kernel size 3 in time with stride 1 and padding 1. Each temporal convolutional layer is followed by one Batch Normalization layer and one ReLU layer.

Finally, **B** adds one more temporal convolutional layer `pred` to output the boundary regression values. Filters in `pred` have kernel size 3 in time with stride 1 and padding 1. Similar to YOLO [44,43], the boundary predicted by **B** is designed to be class-agnostic. This allows us to learn a generic boundary predictor, which may be used for generating action proposals for unseen actions in the future. Consequently, the total number of filters in `pred` is $2M$, where M is the

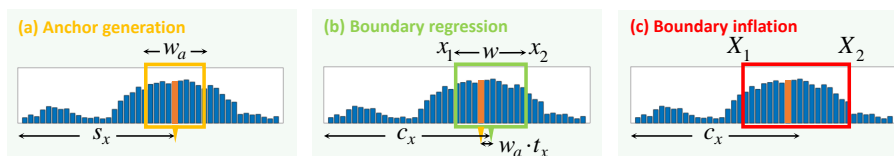


Fig. 3. Illustration of the boundary prediction procedure which consists of three steps sequentially: (1) **anchor generation** to obtain the boundary hypothesis; (2) **boundary regression** to obtain the predicted boundary of the action segment (denoted as the inner boundary); (3) **boundary inflation** to obtain the outer boundary. The score sequence is CAS and the orange score bar indicates the temporal position that the boundary predictor currently looks at

number of anchor scales. For each anchor, **B** predicts two boundary regression values: (1) t_x indicating how to shift the center location of the anchor and (2) t_w indicating how to scale the length of the anchor.

Details of the Boundary Transformation Since each temporal position on the feature map and each temporal position on the CAS both correspond to the same location of an input snippet, we make boundary predictions at the snippet-level granularity. We outline the boundary prediction procedure in Fig. 3.

Anchor generation. At the temporal position s_x on the feature map, we generate a hypothesized segment (anchor) of length w_a . In practice, we use multi-scale anchors. We determine their scales according to the time duration range of the ground truth segments in the training set.

Boundary regression. As aforementioned, for each anchor at the temporal position s_x , **B** predicts two boundary regression values t_x and t_w . We can obtain the predicted segment via regressing the center location $c_x = s_x + w_a \cdot t_x$ and the temporal length $w = w_a \cdot \exp(t_w)$. We denote the boundary of this predicted segment as the inner boundary, which can be computed by $x_1 = c_x - w/2$ and $x_2 = c_x + w/2$. Furthermore, we clip the predicted boundary x_1 and x_2 to fit into the range of the whole video. More details about clipping can be found in the supplementary material.

Boundary inflation. A ground truth segment usually exhibits relatively higher activations on CAS within the inner area $[x_1, x_2]$ compared to the contextual area preceding x_1 and succeeding x_2 . Therefore, we inflate the inner boundary by a ratio α to obtain the corresponding outer boundary $X_1 = x_1 - w \cdot \alpha$ and $X_2 = x_2 + w \cdot \alpha$. As discussed in the supplementary material, setting α to 0.25 is a good choice.

The OIC layer for Obtaining the Final Predictions Finally, we introduce an OIC layer which uses the OIC loss to measure how likely each segment contains actions and then removes the segments that are not likely to contain actions. During testing, this OIC layer outputs a set of predicted segments. During

training, this OIC layer further computes the total OIC loss and back-propagates the gradients to the underlying boundary prediction model.

Concretely, given an input video, the classification branch generates its CAS and the localization branch predicts the candidate class-agnostic segments. Note that since all temporal convolutional layers in \mathbf{B} slide over time with stride 1, the set of segments predicted at each temporal position on the feature map and the activations at each temporal position on the CAS are paired, corresponding to the same input snippet. Thus at the temporal position of each snippet, \mathbf{B} has predicted M class-agnostic anchor segments. Then for each action, we iteratively go through the following steps on the CAS to obtain the final class-specific segment predictions. Note that during training we consider only the ground truth actions while during testing we consider all actions. If a temporal position has the activation lower than 0.1 on the CAS, we discard all the predictions corresponding to this temporal position. For each of the remaining positions, among its M anchor segment predictions, we only keep the one with the lowest OIC loss which means selecting the anchor of the most likely scale. Finally, for all the kept segment predictions, we remove the segment predictions with the OIC loss higher than -0.3. We perform Non-Maximum Suppression (NMS) over all segment predictions with overlap IoU threshold 0.4. All these thresholds are chosen by grid search on the THUMOS’14 training set and also work well on ActivityNet. Alg. 1 in the supplementary material summarizes the above steps.

During training, the total loss is the summation of the OIC loss generated by each kept segment predictions. We can compute the gradients triggered by each kept segment prediction according to Sec. 3.2 and then accumulate them together to update the underlying boundary predictor \mathbf{B} . During testing, all the kept segment predictions are outputted as our final segment predictions. Each segment prediction consists of (1) the predicted action class, (2) the confidence score which is set to 1 minus its OIC loss, and (3) the start time and the end time obtained by converting the inner boundary $[x_1, x_2]$ from the snippet-level granularity (continuous value before rounding to its nearest integer) to time.

5 Experiments

In this section, we first introduce two standard benchmarks and the corresponding evaluation metrics. Note that during training, we only use the video-level labels; during testing, we use the ground truth segments with boundary annotations for evaluating the performance of temporal action localization. We compare our method with the state-of-the-art methods and then conduct some ablation studies to investigate different variants of our method.

5.1 Datasets and Evaluation

THUMOS’14 [29] The temporal action localization task in THUMOS’14 contains 20 actions. Its validation set has 200 untrimmed videos. Each video contains at least one action. We use these 200 videos in the validation set for training. The trained model is tested on the test set which contains 213 videos.

ActivityNet v1.2 [21] To facilitate comparisons, we follow Wang et al. [64] to use the ActivityNet release version 1.2 which covers 100 activity classes. The training set has 4,819 videos and the validation set has 2,383 videos. We train on the training set and test on the validation set.

Evaluation Metrics Given the testing videos, the system outputs a rank list of action segment predictions. Each prediction contains the action class, the starting time and the ending time, and the confidence score. We follow the conventions [29,1] to evaluate mean Average Precision (mAP). Each prediction is regarded as correct only when (1) the predicted class is correct and (2) its temporal overlap IoU with the ground truth segment exceeds the evaluation threshold. We do not allow duplicate detections for the same ground truth segment.

5.2 Implementation Details

We implement our AutoLoc using Caffe [28]. We use the stochastic gradient descent algorithm to train AutoLoc. Through the experimental studies, we find that the training process can converge quickly on both THUMOS'14 and ActivityNet datasets after 1 training epoch. Following Faster R-CNN [45], during each mini-batch, we process one whole untrimmed video. The learning rate is initially set to 0.001 and is reduced by one order of magnitude for every 200 iterations. We set the weight decay to 0.0005. We choose anchors of the snippet-level length 1, 2, 4, 8, 16, 32 for THUMOS'14 and 16, 32, 64, 128, 256, 512 for ActivityNet. We use CUDA 8.0 and cuDNN v5. We use one single NVIDIA GeForce GTX TITAN X GPU.

5.3 Comparisons with the State-of-the-art

The results on THUMOS'14 are shown in Table 1. Our method significantly outperforms the state-of-the-art weakly-supervised TAL methods that are trained with the video-level labels only. Regarding to the recent weakly-supervised TAL methods (i.e. Hide-and-Seek [56] and Wang et al. [64]), although they can generate reasonably good CAS, TAL is done by applying simple thresholding on the CAS which might not robust be to noises in CAS. Our method directly predicts the segment boundary with the contextual information taken into account. Our method can even achieve better or comparable results to some fully-supervised methods (e.g. S-CNN [51]) that are trained with the segment-level boundary annotations. The results of SSN [73] correspond to the model of the same backbone network architecture as ours.

The results on ActivityNet v1.2 are shown in Table 2 and our method can achieve substantial improvements again. Wang et al. [64] did not report temporal localization results on ActivityNet in their paper. But their trained models and source codes have been released online publicly and thus we can evaluate their results on ActivityNet as well.

Table 1. Comparisons with the state-of-the-art methods in terms of temporal localization mAP (%) under different IoU thresholds on THUMOS’14 test set. Weak supervision means training with the video-level labels only. Full supervision indicates that the segment-level boundary annotations are used during training

Supervision	IoU threshold	0.3	0.4	0.5	0.6	0.7
Full	Karaman et al. [33]	0.5	0.3	0.2	0.2	0.1
Full	Wang et al. [62]	14.6	12.1	8.5	4.7	1.5
Full	Heilbron et al. [23]	-	-	13.5	-	-
Full	Escorcia et al. [15]	-	-	13.9	-	-
Full	Oneata et al. [40]	28.8	21.8	15.0	8.5	3.2
Full	Richard and Gall [46]	30.0	23.2	15.2	-	-
Full	Yeung et al. [67]	36.0	26.4	17.1	-	-
Full	Yuan et al. [68]	33.6	26.1	18.8	-	-
Full	Yuan et al. [69]	36.5	27.8	17.8	-	-
Full	S-CNN [51]	36.3	28.7	19.0	10.3	5.3
Full	SST [7]	37.8	-	23.0	-	-
Full	CDC [50]	40.1	29.4	23.3	13.1	7.9
Full	Dai et al. [10]	-	33.3	25.6	15.9	9.0
Full	SSAD [36]	43.0	35.0	24.6	-	-
Full	TURN TAP [17]	44.1	34.9	25.6	-	-
Full	R-C3D [66]	44.7	35.6	28.9	-	-
Full	SS-TAD [6]	45.7	-	29.2	-	9.6
Full	Gao et al. [16]	50.1	41.3	31.0	19.1	9.9
Full	SSN [73]	51.9	41.0	29.8	19.6	10.7
Weak	Sun et al. [58]	8.5	5.2	4.4	-	-
Weak	Hide-and-Seek [56]	19.5	12.7	6.8	-	-
Weak	Wang et al. [64]	28.2	21.1	13.7	-	-
Weak	Ours - AutoLoc	35.8	29.0	21.2	13.4	5.8

5.4 Discussions

In this Section, we address several questions quantitatively to analyze our model.

Q1: How Effective is the Proposed OIC Loss? In order to evaluate the effectiveness of the proposed OIC loss, we enumerate all candidate segments at the snippet-level granularity (for example, a segment starting at the location of the 2-nd snippet and ending at the location of the 6-th snippet). We leverage the OIC loss to measure how likely each segment contains actions and then select the most likely ones. Concretely, for each segment, we compute its OIC loss of being each action. Then we follow Sec. 4.3 to remove segments with high OIC loss and remove duplicate predictions via NMS. We denote this approach as **OIC Selection**. As shown in Table 3, although not as good as AutoLoc, **OIC Selection** still significantly improves the state-of-the-art results [64]. Because the OIC loss explicitly favors the segment which has high activations inside and low activations outside, and also such a segment of low OIC loss is usually

Table 2. Comparisons with the state-of-the-art methods in terms of temporal localization mAP (%) under different IoU thresholds on ActivityNet v1.2 validation set. Weak supervision means training with the video-level labels only. Full supervision indicates that the segment-level boundary annotations are used during training

Supervision	IoU threshold	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	Avg
Full	SSN [73]	41.3	38.8	35.9	32.9	30.4	27.0	22.2	18.2	13.2	6.1	26.6
Weak	Wang et al. [64]	7.4	6.1	5.2	4.5	3.9	3.2	2.5	1.8	1.2	0.7	3.6
Weak	Ours - AutoLoc	27.3	24.9	22.5	19.9	17.5	15.1	13.0	10.0	6.8	3.3	16.0



Fig. 4. Illustration of the LoG filter and our OIC loss, which in effect is a step function

well aligned to the ground truth segment. This confirms the effectiveness of the proposed OIC loss.

Q2: How Important is Looking into the Contrast Between the Inner Area and the Outer Area? The core idea of the OIC loss is encouraging high activations in the inner area while penalizing high activations in the outer area. We consider another variant that can also discover the segment-level supervision but does not model the contrast between inner and outer. Specifically, we change the OIC loss in AutoLoc to **Inner Only Loss**, which only encourages high activations inside the segment but does not look into the contextual area. The detailed formulation of how to compute the loss and gradients in the **Inner Only Loss** can be found in the supplementary material. As shown in Table 3, the performances drop a lot. Consequently, when designing the loss for training the boundary predictor, it is very important and effective to take into account the contrast between the inner area and the outer area.

Notably, the idea of looking into the contrast between inner and outer is related to the usage of **Laplacian of Gaussian (LoG)** filter for blob detection [37]. The operation of computing the OIC loss is effectively convolving the CAS with a step function as shown in Fig. 4, which can be regarded as a variant of the LoG filter for the sake of easing the network training. As proven in the supplementary material, the integral of the LoG filter and the integral of the step function are both zero on the range $(-\text{Inf}, +\text{Inf})$. Further, we approach the scale selection in blob detection by the multi-anchor mechanism and the boundary regression method. Despite the simplicity of the OIC loss, it turns out to be quite effective in practice for localizing likely action segments.

Q3: What is the Advantage of Learning a Model on the Training Videos Compared to Directly Optimizing the Boundaries on the Test-

Table 3. Temporal localization mAP (%) under different IoU thresholds on ActivityNet v1.2 validation set. All approaches are trained with weak supervision

IoU threshold	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	Avg
Wang et al. [64]	7.4	6.1	5.2	4.5	3.9	3.2	2.5	1.8	1.2	0.7	3.6
Ours - AutoLoc	27.3	24.9	22.5	19.9	17.5	15.1	13.0	10.0	6.8	3.3	16.0
Q1: OIC Selection	15.8	13.7	11.9	10.3	8.8	7.5	6.4	5.1	3.6	2.2	8.5
Q2: Inner Only Loss	4.6	3.7	2.7	1.9	1.3	0.9	0.5	0.2	0.1	0.0	1.6
Q3: Direct Optimization	21.8	19.6	17.8	15.8	13.8	11.7	9.8	7.8	5.5	2.7	12.6

ing Videos? AutoLoc trains a model on the training videos and then applies the trained model to perform inference on the testing videos. Alternatively, without training the boundary predictor \mathbf{B} on the training videos, we can directly train/optimize \mathbf{B} from scratch on each testing video individually: we follow the testing pipeline as described in Sec. 4.3 while we also conduct the back-propagation to update \mathbf{B} to iteratively find likely segments on each testing video. We refer this approach as **Direct Optimization**. As shown in Table 3, its performance is not bad, which confirms the effectiveness of the OIC loss again. But it is still not as good as AutoLoc. Because **Direct Optimization** optimizes the predicted boundaries according to the testing video’s CAS, which may not be very accurate. Eventually **Direct Optimization** overfits such an inaccurate CAS and thus results into imperfect boundary predictions. In AutoLoc, \mathbf{B} has been trained on multiple training videos and thus is robust to the noises in CAS. Consequently, AutoLoc may still predicts good boundary even when the testing video’s CAS is not perfect. Furthermore, **Direct Optimization** requires optimizing the boundary predictions on the testing video until convergence and thus its testing speed is much slower than AutoLoc. For example, on ActivityNet, **Direct Optimization** converges after 25 training iterations (25 forward passes and 25 backward passes). However, AutoLoc directly applies the trained model to do inference on the testing video and thus requires only one forward pass during testing.

6 Conclusion and Future Works

In this paper, we have presented a novel weakly-supervised TAL framework to directly predict temporal boundary in a single-shot fashion and proposed a novel OIC loss to provide the needed segment-level supervision. In the future, it would be interesting to extend AutoLoc for object detection in image.

7 Acknowledgment

We appreciate the support from Mitsubishi Electric for this project. This project adapted Zheng’s internship project at MSR from image domain to video domain.

References

1. Activitynet challenge 2016. <http://activity-net.org/challenges/2016/> (2016)
2. Aggarwal, J.K., Ryoo, M.S.: Human activity analysis: A review. In: ACM Computing Surveys (2011)
3. Asadi-Aghbolaghi, M., Clapés, A., Bellantonio, M., Escalante, H.J., Ponce-López, V., Baró, X., Guyon, I., Kasaei, S., Escalera, S.: A survey on deep learning based approaches for action and gesture recognition in image sequences. In: FG (2017)
4. Bearman, A., Russakovsky, O., Ferrari, V., Fei-Fei, L.: Whats the point: Semantic segmentation with point supervision. In: ECCV (2016)
5. Bilen, H., Vedaldi, A.: Weakly supervised deep detection networks. In: CVPR (2016)
6. Buch, S., Escorcia, V., Ghanem, B., Fei-Fei, L., Niebles, J.C.: End-to-end, single-stream temporal action detection in untrimmed videos. In: BMVC (2017)
7. Buch, S., Escorcia, V., Shen, C., Ghanem, B., Niebles, J.C.: Sst: Single-stream temporal action proposals. In: CVPR (2017)
8. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: CVPR (2017)
9. Cheng, G., Wan, Y., Saudagar, A.N., Namuduri, K., Buckles, B.P.: Advances in human action recognition: A survey (2015), <http://arxiv.org/abs/1501.05964>
10. Dai, X., Singh, B., Zhang, G., Davis, L.S., Chen, Y.Q.: Temporal context network for activity localization in videos. In: ICCV (2017)
11. Dave, A., Russakovsky, O., Ramanan, D.: Predictive-corrective networks for action detection. In: CVPR (2017)
12. Dietterich, T.G., Lathrop, R.H., Lozano-Pérez, T.: Solving the multiple instance problem with axis-parallel rectangles (1997)
13. Donahue, J., Hendricks, L.A., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T.: Long-term recurrent convolutional networks for visual recognition and description. In: CVPR (2015)
14. Durand, T., Mordan, T., Thome, N., Cord, M.: Wildcat: Weakly supervised learning of deep convnets for image classification, pointwise localization and segmentation. In: CVPR (2017)
15. Escorcia, V., Heilbron, F.C., Niebles, J.C., Ghanem, B.: Daps: Deep action proposals for action understanding. In: ECCV (2016)
16. Gao, J., Yang, Z., Nevatia, R.: Cascaded boundary regression for temporal action detection. In: BMVC (2017)
17. Gao, J., Yang, Z., Sun, C., Chen, K., Nevatia, R.: Turn tap: Temporal unit regression network for temporal action proposals. In: ICCV (2017)
18. Girshick, R.: Fast r-cnn. In: ICCV (2015)
19. Gorban, A., Idrees, H., Jiang, Y.G., Zamir, A.R., Laptev, I., Shah, M., Sukthankar, R.: THUMOS challenge: Action recognition with a large number of classes. <http://www.thumos.info/> (2015)
20. Gudi, A., van Rosmalen, N., Loog, M., van Gemert, J.: Object-extent pooling for weakly supervised single-shot localization. In: BMVC (2017)
21. Heilbron, F.C., Escorcia, V., Ghanem, B., Niebles, J.C.: Activitynet: A large-scale video benchmark for human activity understanding. In: CVPR (2015)
22. Heilbron, F.C., Barrios, W., Escorcia, V., Ghanem, B.: Scc: Semantic context cascade for efficient action detection. In: CVPR (2017)
23. Heilbron, F.C., Niebles, J.C., Ghanem, B.: Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In: CVPR (2016)

24. Hong, S., Yeo, D., Kwak, S., Lee, H., Han, B.: Weakly supervised semantic segmentation using web-crawled videos. In: CVPR (2017)
25. Huang, D.A., Fei-Fei, L., Niebles, J.C.: Connectionist temporal modeling for weakly supervised action labeling. In: ECCV (2016)
26. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
27. Ji, S., Xu, W., Yang, M., Yu, K.: 3d convolutional neural networks for human action recognition. In: TPMAI (2013)
28. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. In: ACM MM (2014)
29. Jiang, Y.G., Liu, J., Zamir, A.R., Toderici, G., Laptev, I., Shah, M., Sukthankar, R.: THUMOS challenge: Action recognition with a large number of classes. <http://csrcv.ucf.edu/THUMOS14/> (2014)
30. Jie, Z., Wei, Y., Jin, X., Feng, J., Liu, W.: Deep self-taught learning for weakly supervised object localization. In: CVPR (2017)
31. Kang, S.M., Wildes, R.P.: Review of action recognition and detection methods. arXiv preprint arXiv:1610.06906 (2016)
32. Kantorov, V., Oquab, M., Cho, M., Laptev, I.: Contextlocnet: Context-aware deep network models for weakly supervised localization. In: ECCV (2016)
33. Karaman, S., Seidenari, L., Bimbo, A.D.: Fast saliency based pooling of fisher encoded dense trajectories. In: ECCV THUMOS Workshop (2014)
34. Khoreva, A., Benenson, R., Hosang, J., Hein, M., Schiele, B.: Simple does it: Weakly supervised instance and semantic segmentation. In: CVPR (2017)
35. Kim, D., Yoo, D., Kweon, I.S., et al.: Two-phase learning for weakly supervised object localization. In: ICCV (2017)
36. Lin, T., Zhao, X., Shou, Z.: Single shot temporal action detection. In: ACM MM (2017)
37. Lindeberg, T.: Feature detection with automatic scale selection. IJCV (1998)
38. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: ECCV (2016)
39. Mettes, P., van Gemert, J.C., Snoek, C.G.: Spot on: Action localization from pointly-supervised proposals. In: ECCV (2016)
40. Oneata, D., Verbeek, J., Schmid, C.: The lear submission at thumos 2014. In: ECCV THUMOS Workshop (2014)
41. Papandreou, G., Chen, L.C., Murphy, K.P., Yuille, A.L.: Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In: CVPR (2015)
42. Poppe, R.: A survey on vision-based human action recognition. In: Image and vision computing (2010)
43. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR (2016)
44. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: CVPR (2017)
45. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS (2015)
46. Richard, A., Gall, J.: Temporal action detection using a statistical language model. In: CVPR (2016)
47. Richard, A., Kuehne, H., Gall, J.: Weakly supervised action learning with rnn based fine-to-coarse modeling. In: CVPR (2017)
48. Shen, Z., Li, J., Su, Z., Li, M., Chen, Y., Jiang, Y.G., Xue, X.: Weakly supervised dense video captioning. In: CVPR (2017)

49. Shi, M., Caesar, H., Ferrari, V.: Weakly supervised object localization using things and stuff transfer. In: ICCV (2017)
50. Shou, Z., Chan, J., Zareian, A., Miyazawa, K., Chang, S.F.: Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. In: CVPR (2017)
51. Shou, Z., Wang, D., Chang, S.F.: Temporal action localization in untrimmed videos via multi-stage cnns. In: CVPR (2016)
52. Sigurdsson, G.A., Divvala, S., Farhadi, A., Gupta, A.: Asynchronous temporal fields for action recognition. In: CVPR (2017)
53. Sigurdsson, G.A., Russakovsky, O., Farhadi, A., Laptev, I., Gupta, A.: Much ado about time: Exhaustive annotation of temporal data. In: HCOMP (2016)
54. Sigurdsson, G.A., Varol, G., Wang, X., Farhadi, A., Laptev, I., Gupta, A.: Hollywood in homes: Crowdsourcing data collection for activity understanding. In: ECCV (2016)
55. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
56. Singh, K.K., Lee, Y.J.: Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In: ICCV (2017)
57. Sun, C., Paluri, M., Collobert, R., Nevatia, R., Bourdev, L.: Pronet: Learning to propose object-specific boxes for cascaded neural networks. In: CVPR (2016)
58. Sun, C., Shetty, S., Sukthankar, R., Nevatia, R.: Temporal localization of fine-grained actions in videos by domain transfer from web images. In: ACM MM (2015)
59. Tang, P., Wang, X., Bai, X., Liu, W.: Multiple instance detection network with online instance classifier refinement. In: CVPR (2017)
60. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: ICCV (2015)
61. Tran, D., Ray, J., Shou, Z., Chang, S.F., Paluri, M.: Convnet architecture search for spatiotemporal feature learning. arXiv preprint arXiv:1708.05038 (2017)
62. Wang, L., Qiao, Y., Tang, X.: Action recognition and detection by combining motion and appearance features. In: ECCV THUMOS Workshop (2014)
63. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Gool, L.V.: Temporal segment networks: Towards good practices for deep action recognition. In: ECCV (2016)
64. Wang, L., Xiong, Y., Lin, D., Van Gool, L.: Untrimmednets for weakly supervised action recognition and detection. In: CVPR (2017)
65. Weinland, D., Ronfard, R., Boyer, E.: A survey of vision-based methods for action representation, segmentation and recognition. In: Computer Vision and Image Understanding (2011)
66. Xu, H., Das, A., Saenko, K.: R-c3d: Region convolutional 3d network for temporal activity detection. In: ICCV (2017)
67. Yeung, S., Russakovsky, O., Mori, G., Fei-Fei, L.: End-to-end learning of action detection from frame glimpses in videos. In: CVPR (2016)
68. Yuan, J., Ni, B., Yang, X., Kassim, A.: Temporal action localization with pyramid of score distribution features. In: CVPR (2016)
69. Yuan, Z., Stroud, J.C., Lu, T., Deng, J.: Temporal action localization by structured maximal sums. In: CVPR (2017)
70. Zhang, H., Kyaw, Z., Yu, J., Chang, S.F.: Ppr-fcn: weakly supervised visual relation detection via parallel pairwise r-fcn. In: ICCV (2017)
71. Zhang, J., Lin, Z., Brandt, J., Shen, X., Sclaroff, S.: Top-down neural attention by excitation backprop. In: ECCV (2016)

72. Zhao, H., Yan, Z., Wang, H., Torresani, L., Torralba, A.: Slac: A sparsely labeled dataset for action classification and localization. arXiv preprint arXiv:1712.09374 (2017)
73. Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., Lin, D.: Temporal action detection with structured segment networks. In: ICCV (2017)
74. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)
75. Zhu, Y., Zhou, Y., Ye, Q., Qiu, Q., Jiao, J.: Soft proposal networks for weakly supervised object localization. In: ICCV (2017)