

Temporal Action Proposal Generation with Background Constraint

Haosen Yang^{1,2*}, Wenhao Wu^{2*}, Lining Wang¹, Sheng Jin¹,
Boyang Xia^{2,3}, Hongxun Yao^{1†}, Hujie Huang¹

¹ Harbin Institute of Technology

² Department of Computer Vision Technology (VIS), Baidu Inc.

³ University of Chinese Academy of Sciences

{haosen.yang.6, liningwang98, jsh.hit.doc}@gmail.com, {wuwenhao17, xiaboyang20}@mailsucas.edu.cn,
{h.yao, hjhuang}@hit.edu.cn

Abstract

Temporal action proposal generation (TAPG) is a challenging task that aims to locate action instances in untrimmed videos with temporal boundaries. To evaluate the confidence of proposals, the existing works typically predict action score of proposals that are supervised by the temporal Intersection-over-Union (tIoU) between proposal and the ground-truth. In this paper, we innovatively propose a general auxiliary *Background Constraint* idea to further suppress low-quality proposals, by utilizing the background prediction score to restrict the confidence of proposals. In this way, the Background Constraint concept can be easily plug-and-played into existing TAPG methods (e.g., BMN, GTAD). From this perspective, we propose the Background Constraint Network (BC-Net) to further take advantage of the rich information of action and background. Specifically, we introduce an Action-Background Interaction module for reliable confidence evaluation, which models the inconsistency between action and background by attention mechanisms at the frame and clip levels. Extensive experiments are conducted on two popular benchmarks, i.e., ActivityNet-1.3 and THUMOS14. The results demonstrate that our method outperforms state-of-the-art methods. Equipped with the existing action classifier, our method also achieves remarkable performance on the temporal action localization task.

Introduction

With the rapid development of mobile devices and the Internet, a massive amount of video content is being uploaded to the Internet every second. The volume of video information has far exceeded the processing capacity of the conventional manual system, thus video content analysis has attracted the extensive interest of academic and industrial communities.

One of the most active research topics in video understanding is temporal action detection, which focuses on both classifying the action instances present in an untrimmed video and localizing them with temporal boundaries. The temporal action detection task, like object detection, is divided into two parts: temporal action proposal generation

*Co-first authorship. This work was done when Haosen Yang was a research intern at Baidu Inc, Wenhao Wu was the project leader.

†Corresponding author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

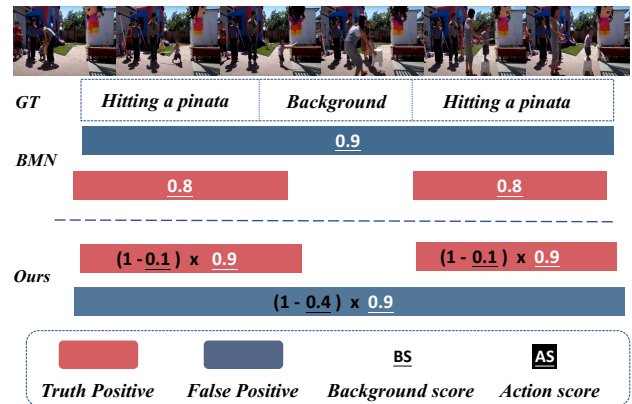


Figure 1: Illustration of the background constraint concept. We introduce a background score for the confidence of the proposal, which helps to restrict false-positive proposals.

(TAPG) and action recognition. Deep learning has recently been shown to significantly improve action recognition performance (Simonyan and Zisserman 2014; Lin, Gan, and Han 2019; Wu et al. 2021a). However, the performance of the two-stage temporal action detectors in mainstream benchmarks (Jiang et al. 2014; Caba Heilbron et al. 2015a) still has much room for improvement, which is mostly influenced by the quality of proposals from temporal action proposal generation.

Hence, great efforts have been devoted to TAPG task (Lin et al. 2018, 2019). These research generally use the temporal Intersection-over-Union (tIoU) between the proposal and instance, called the action score, to evaluate the confidence of the proposal in order to develop high-quality temporal action proposals with dependable confidence scores. However, the background information is also significant but was previously overlooked. For instance, as illustrated in Figure 1, we can easily restrict false-positive proposals by detecting the background “Move a chair”. Furthermore, we can evaluate the inconsistency between the background “Move a chair” and the action “Hitting a pinata”, leading to better action score and background score of the proposal.

Motivated by above observations, we propose a general auxiliary *Background Constraint* idea to reduce localization

errors. Specifically, we introduce a background score for the proposal’s confidence, and its supervision signal is defined by the temporal Intersection-over-Ancor (tIoA) between the proposal and the background. This concept can be flexibly integrated into existing TAPG methods (*e.g.*, BMN (Lin et al. 2019), GTAD (Xu et al. 2020)) to improve the performance in a plug-and-play fashion.

To further mine the rich information of action and background, in this paper, we propose the *Background Constraint Network* (BCNet) to generate high-quality temporal action proposals. An essential component of BCNet is the *Action-Background Interaction* (ABI) module, which performs both frame-level and clip-level action-background interaction to obtain reliable confidence scores of proposals. To do so, we first generate action features and background features for each frame using self-attention and difference-attention. Sliding windows are then used to generate multi-scale anchors from action and background features. The clip-level interaction then discovers the complex relationships between action-anchors and background-anchors, and outputs the action and background scores for these anchors. We also propose a *Boundary Prediction* (BP) module for precisely locating action boundaries. To capture the complex long-term temporal relationships while avoiding the influence of global noise, we aggregate the original feature sequence using self-attention and cross-attention mechanisms. The output representation is then used as the global representation for the boundary prediction task. Finally, we feed the boundary probabilities, action scores and background scores into the post processing module to get the final proposal set.

Experimental results show the superiority of our system on two popular datasets, *i.e.*, ActivityNet (Caba Heilbron et al. 2015b) and THUMOS14 (?). Our BCNet achieves significant performance and outperforms existing state-of-the-art methods on both datasets. Our contributions are summarized as follows:

- We introduce a Background Constraint concept, which can be integrated easily with existing TAPG methods (*e.g.*, BMN, GTAD) and improve performance significantly.
- We propose a Background Constraint Network, which consists of multiple attention units, *i.e.*, self-attention, cross-attention and difference-attention, and generates high-quality proposals by exploiting inconsistency between action and background.
- Extensive experiments demonstrate that our method outperforms the existing state-of-the-art methods on THUMOS14 and achieves comparable performance on ActivityNet-1.3, in both temporal action proposal generation task and temporal action detection task.

Related Work

Video Action Recognition

Action recognition is a fundamental task in the video understanding area. Currently, there are two types of end-to-end action recognition methods: 3D CNN-based methods and

2D CNN-based methods. 3D CNNs are natural extensions of their 2D counterparts and are intuitive spatiotemporal networks that directly tackle 3D volumetric video data (Tran et al. 2015; Carreira and Zisserman 2017) but have a high computational cost. Other alternative efficient architectures, such as TSM (Lin, Gan, and Han 2019), TEI (Liu et al. 2020), MVFNet (Wu et al. 2021a), DSANet (Wu et al. 2021b), *etc.*, have been developed to capture temporal information with reasonable training resources. These methods aim to design efficient temporal modules to perform efficient temporal modeling. There is also ongoing research into dynamic inference (Wu et al. 2020), adaptive frame sampling techniques (Wu et al. 2019; Korbar, Tran, and Torresani 2019), which we believe can complement the end-to-end video recognition approaches.

Temporal Action Proposal Generation

Temporal action proposal generation aims to detect action instances with temporal boundaries and confidence in untrimmed videos. Existing methods can be mainly divided into Top-down and Bottom-up methods. The Top-down methods (Gao et al. 2017; Gao, Chen, and Nevatia 2018; Gao et al. 2020; Chen et al. 2019) generate proposals using sliding windows or pre-defined anchors. The Bottom-up methods mainly focus on evaluating “actionness”, which indicates the probability of a potential action, for each frame or clip in a video. These works (Shou, Wang, and Chang 2016; Zhao et al. 2017) use snippet-wise probability to generate candidate proposals. BSN (Lin et al. 2018) first proposes to predict start, end and actionness of each frame, then proposals are generated by constructing start and end points with high probabilities, with low confidence ones further abandoned by an evaluation module. (Lin et al. 2019; Su et al. 2020; Xu et al. 2020; Lin et al. 2020) generate all possible combinations of temporal locations to evaluate confidence of proposals. (Liu et al. 2019) generates coarse segment proposals by perceiving the whole video sequence and predicts the frame actionness by densely evaluating each video frame. These methods evaluate action scores of proposals with rich clip-level context. However, these methods fail to take full advantage of background by focusing only on the action score. In our work, we predict the extra background score for the confidence of proposals to reduce low-quality proposals.

Transformer and Self-Attention Mechanism

Transformers (Vaswani et al. 2017) has achieved great success in natural language processing. Transformer architectures are based on a self-attention mechanism that summarizes content from the source sequence and is capable of modeling complex and arbitrary dependencies within a limited number of layers. Recently, many works (Dosovitskiy et al. 2020; Carion et al. 2020; Liu et al. 2021b; Tan et al. 2021; Wang et al. 2021) have revealed the great potential of Transformers in the computer vision task. Inspired by the successful application of Transformers in various fields, we intuitively take advantage of Transformers in modeling long-range contextual information. In this paper, we utilize the Transformer-alike structure to devise three attention units.

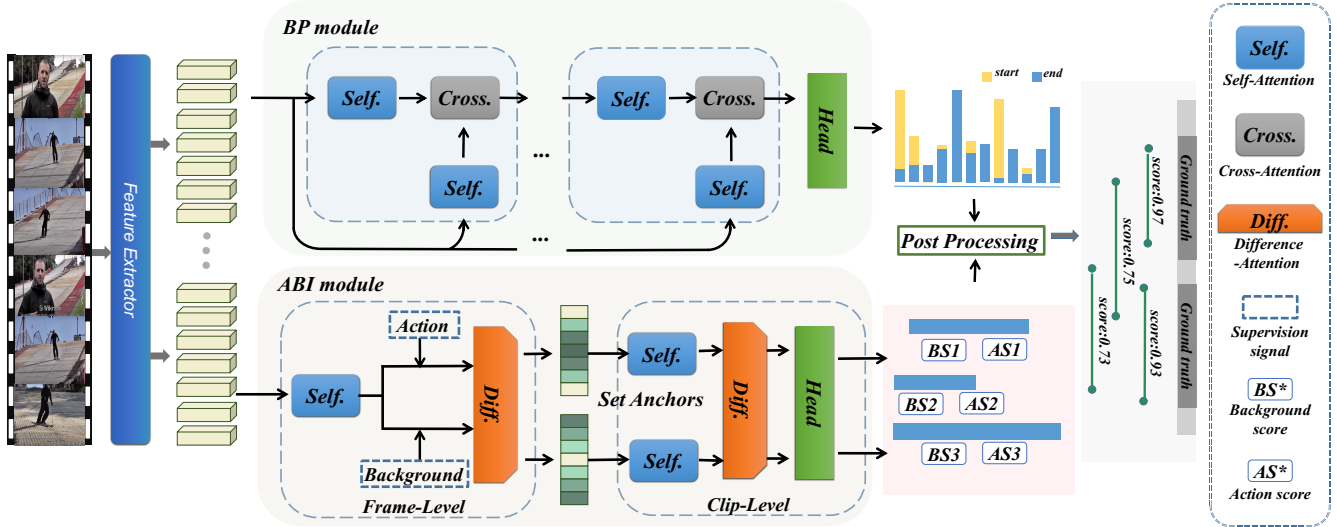


Figure 2: Illustration of the proposed Background Constraint Network. First we apply the feature extractor to encode video frames. Boundary Prediction (BP) module takes the feature sequence as input, and outputs boundary probability sequence. Action-Background Interaction (ABI) module takes the feature sequence as input, and outputs features of action and background at the frame level. Then, we set anchors on features of action and background, and feed them into clip-level interaction to generate action and background scores of anchors. Finally, we construct proposals based on boundary probabilities sequence and refine them using the corresponding anchor.

Background Modeling on Temporal Action Localization

Background modeling in Temporal action localization has received some attention. Several previous works (Shou, Wang, and Chang 2016; Yuan et al. 2016) generate proposals by sliding window and classify them into $C + 1$ classes for C action classes plus background class. Also, several studies attempt to explicit background modeling for weakly-supervised temporal action localization. Some works (Nguyen, Ramanan, and Fowlkes 2019; Lee, Uh, and Byun 2020) try to classify background frames as a separate class. (Lee et al. 2020) formulates background frames as out-of-distribution samples. Essentially, all the above works aim to perform classification for these proposals. Unlike them, in our work, we propose a Background Constraint concept to predict an additional background score for proposal confidence evaluation. To supervise the background score, we use temporal Intersection-over-Ancor (tIoA) between the proposal and the background. Our work concentrates on utilizing the background prediction score to restrict the confidence of proposals.

Background Constraint Network

As shown in Figure 2, we propose a Background Constraint Network (BCNet) to generate high-quality proposals, which mainly consists of two main modules: Action-Background Interaction Module and Boundary Prediction Module. Firstly, the Action-Background Interaction (ABI) module is adopted to perform both frame-level and clip-level action-background interaction to obtain reliable confidence scores of proposals. The Boundary Prediction (BP) module is then utilized to locate the boundaries of the proposals

by exploiting complex long-term temporal relationships for boundary regression.

Problem Definition

An untrimmed video U can be denoted as a frame sequence $U = \{u_t\}_{t=1}^{l_v}$ with l_v frames, where u_t denotes the t -th RGB frame of video U . The temporal annotation set of U is made up of a set of temporal action instances as $\Psi_g = \{\phi_n^g\}_{n=1}^{N_g}$ and $\phi_n^g = (t_{s_n}, t_{e_n})$, where N_g is the number of ground-truth action instances, t_{s_n} and t_{e_n} are the starting and ending time of the action instance ϕ_n^g , respectively. During training phase, the Ψ_g is provided. While in the testing phase, the predicted proposal set Ψ_p should cover the Ψ_g with high recall and high temporal overlapping.

Background Constraint

To evaluate the confidence of the proposal, existing methods primarily use the temporal Intersection-over-Union (tIoU) between the proposal and instance, called *action score*. The temporal Intersection-over-Union (tIoU) is used to define the label of action score, which can be computed by:

$$A_{label} = \max\left\{\left|\frac{G_i \cap P}{G_i \cup P}\right|_{i=1}^n\right\}, \quad (1)$$

where G_i is the i -th ground truth and P is the proposal, n is number of ground truth. In this paper, we propose a novel Background Constraint concept to suppress low-quality proposals (false positive proposal). Specifically, we predict an extra *background score* for evaluating the confidence of the proposal besides the action score. The label of the background score is generated using temporal Intersection-over-




Figure 3: Illustration of three different attention units.

Anchor (tIoA), which can be computed by:

$$B_{label} = 1 - \sum_{i=1}^n \frac{G_i \cap P}{P}, \quad (2)$$

where G_i is the i -th ground-truth and P is the proposal, n is number of ground-truth. It is worth reminding that A_{label} plus B_{label} is not equal to 1 in most circumstances.

Action-Background Interaction Module

In this section, we describe the Action-Background Interaction (ABI) module in our BCNet. The ABI module conducts action-background interaction on two temporal granularities: frame-level and clip-level. Inspired by the Transformer (Vaswani et al. 2017), ABI module models the interactions with two attention units, *i.e.*, self-attention unit and difference-attention unit. The details are described next.

Frame-level Interaction. We obtain the frame-wise feature sequence $F_o \in \mathbb{R}^{T \times C}$ via the feature encoder, where T is the length of the feature sequence and C is the feature dimension. First, we use the *Self-attention Unit* to learn the relationships between frames and enhance the feature representations. The structure of the *Self-attention Unit* is shown in Figure 3(a). We utilize the Transformer-alike structure to devise the self-attention unit, which consists of two sub-layers: self-attention layer and feed-forward network (FFN). Specifically, the input sequence F_o is projected onto three learnable linear transformations to get queries Q_{F_o} , keys K_{F_o} , and values V_{F_o} . Then the self-attention map can be calculated as follows:

$$A(Q_{F_o}, K_{F_o}) = \text{Softmax} \left(\frac{Q_{F_o} K_{F_o}^T}{\sqrt{C}} \right). \quad (3)$$

We update each component of the sequence by aggregating global information from the complete frame sequence by

$$F_{att} = \varphi^g (A(Q_{F_o}, K_{F_o}) V_{F_o}), \quad (4)$$

where φ^g is a linear projection function. Also, a residual connection around each of the two sub-layers and layer normalization is adopted to generate enhanced feature F' , which can be written as:

$$F' = \text{LayerNorm} (F_{att} + F_o), \quad (5)$$

$$F = \text{LayerNorm} (F' + \text{FFN}(F')).$$

Next, the enhanced features F are fed into the *Difference-attention Unit* to aggregate features based on the inconsistency between the action feature F_a and the background feature F_b under action and background supervision. The structure of the *Difference-attention Unit* is shown in Figure 3(c).

The action feature F_a is projected onto two learnable linear transformations to get queries Q_{F_a} and values V_{F_a} . Also, we transform the background features F_b to queries K_{F_b} and values V_{F_b} . Then, we compute the difference map $A(Q_{F_a}, K_{F_b})$ as:

$$A(Q_{F_a}, K_{F_b}) = \text{Softmax} \left(\frac{Q_{F_a} K_{F_b}^T}{\sqrt{C}} \right). \quad (6)$$

In this way, $A_{i,j}$ represents the difference between the action frame i and the background frame j . The smaller the value of $A_{i,j}$, the bigger the difference between the background features and action features. Then we use the difference map to reweight V_{F_a} , V_{F_b} , and obtain enhanced feature F'_a , F'_b respectively.

Finally, we append a prediction head which encode the F'_a, F'_b with multi-layer perceptron (MLP) network and followed by a *Sigmoid* layer to generate the action and background probability sequence.

Clip-level Interaction. We first use sliding window group to generate the action anchors and background anchors with different scales. Following BMN (Lin et al. 2019), we construct weight term $w_{i,j} \in \mathbb{R}^{N \times T}$ via uniformly sampling N points between the temporal region for each anchor. First, we conduct dot product in temporal dimension between $w_{i,j}$ and F'_a with the shape $C \times N$ to generate the action anchor. Then, we get action anchor sequence $F_a^c \in \mathbb{R}^{L \times S}$ where L is number of clip and $S = C \times N$. Similarly, we generate background anchors sequence F_b^c in the same way.

Next, the anchor sequences F_a^c and F_b^c are fed into the clip-level action-background interaction to generate action score and background score. Specifically, we first utilize two independent *Self-attention Units* to capture the relationships among action/background anchors, respectively. The two self-attention units output updated anchor sequence $F_a^{c'}$ and $F_b^{c'}$. Similar to the frame-level interaction, $F_a^{c'}$ and $F_b^{c'}$ are then fed into the *Difference-attention Unit* to obtain difference map and reweighted anchor sequence $\widetilde{F_a^{c'}}$ and $\widetilde{F_b^{c'}}$. Note that the difference map $A_{i,j}$ represents the feature difference between i -th action anchor and j -th background anchor.

Finally, we add a clip-level predictor which encodes the $\widetilde{F_a^{c'}}$ and $\widetilde{F_b^{c'}}$ with multi-layer perceptron (MLP) and a *Sigmoid* layer to predict action scores and background scores.

Boundary Prediction Module

Long-term temporal modeling is a critical factor in proposal boundary prediction. It is natural to use self-attention mechanism to model long dependencies. However, global modeling is easy to introduce global noise then leads to the over-smoothing. To this end, we propose a *Boundary Prediction* (BP) module which introduces original features to alleviate this phenomenon. This module is built using the Transformer-alike structure which consist of multiple layers. Each layer contains a *Self-attention Unit*, a *Cross-attention Unit* and a feed-forward network. Specifically, we first obtain the feature F_i (i represents the input features of layer i ,

if $i = 1$, $F_i = F_o$) and F_o . Then, we feed them to the *Self-attention Unit* and generate augmented global features F_i^g and F_o^g . As shown in Fig. 3(b), we use *Cross-attention Unit* to generate the attention map $A(F_i^g, F_o^g)$ which represents the similarity between the aggregated feature F_i^g and the original aggregated feature F_o^g , called the originality score. To get F_{i+1} , we aggregate features which have high originality scores and discard features which have low originality scores.

The final output representation is then used as the global representation for the boundary prediction task. Specifically, we utilize a boundary predictor which encode the output representation with multi-layer perceptron (MLP) network and followed by a *Sigmoid* layer to generate boundary probability sequence.

Training

The overall objective of our framework is defined as:

$$\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2, \quad (7)$$

where \mathcal{L}_1 and \mathcal{L}_2 are the objective functions of the ABI module and the BP module respectively.

Objective of BP module. The BP module generates the starting and ending probability sequence P_s, P_e . Thus, the loss function consists of starting loss and ending loss:

$$\mathcal{L}_1 = \mathcal{L}_{bl}(P_s, G_s) + \mathcal{L}_{bl}(P_e, G_e), \quad (8)$$

where G_s and G_e are the ground truth labels of boundary sequence, and \mathcal{L}_{bl} is the binary logistic regression loss.

Objective of ABI module. The ABI module generates frame-level and clip-level scores: $P_a^f, P_b^f, \widehat{P}_a^c, \widehat{P}_a^c$ and P_b^c . P_a^f and P_b^f are frame-level action and background classification scores. P_b^c is clip-level background classification scores. Following BMN, \widehat{P}_a^c is clip-level action classification scores and \widehat{P}_a^c is regression action scores. The loss function \mathcal{L}_2 consists of frame-level loss and clip-level loss:

$$\mathcal{L}_2 = \mathcal{L}_{frame} + \mathcal{L}_{clip}. \quad (9)$$

The frame-level loss is

$$\mathcal{L}_{frame} = \mathcal{L}_c(P_a^f, G_a^f) + \mathcal{L}_c(P_b^f, G_b^f), \quad (10)$$

where G_a^f and G_b^f are the ground truth labels of action and background probability at frame-level. The clip-level loss is formulated as follows:

$$\mathcal{L}_{clip} = \mathcal{L}_c(\widehat{P}_a^c, G_a^c) + \mathcal{L}_r(\widehat{P}_a^c, G_a^c) + \mathcal{L}_c(P_b^c, G_b^c), \quad (11)$$

where G_a^c and G_b^c are the ground truth labels of action and background scores at clip-level. \mathcal{L}_c denotes the binary logistic regression loss function and \mathcal{L}_r is a smooth L_1 loss.

Inference

As mentioned above, the BP module generates boundary probability and the ABI module generates the action and background scores. Then we take the boundary probability, action scores and background scores into the Post-processing module. Firstly, we construct a proposals set ψ_p^c based on boundary probabilities. Second, the proposal

is refined by a corresponding pre-set anchor. The proposal $\varphi = [t'_s, t'_e] \in \psi_p^c$ is taken as an example, we compute the temporal Intersection over Union (tIoU) between proposal φ and anchors, then select a matching anchor $p_m = [t_s^m, t_e^m]$ to refine proposals. we refine the proposal as:

$$[t_s, t_e] = \begin{cases} [\frac{t'_s+t_s^m}{2}, \frac{t'_e+t_e^m}{2}], & \text{if } \widehat{p}_m^a > \alpha_1 \text{ and } \widehat{p}_m^a > \alpha_2, \\ [t'_s, t'_e], & \text{others} \end{cases}, \quad (12)$$

where \widehat{p}_m^a is the anchor action classification score, \widehat{p}_m^a is the action regression score, α_1 and α_2 are the adjustment thresholds. Finally, we get a proposal set $\psi_p = \{\phi_n = (t_s, t_e, p_{t'_s}^s, p_{t'_e}^e, \widehat{p}_m^a, \widehat{p}_m^a, p_m^b)\}_{n=1}^N$, where $p_{t'_s}^s, p_{t'_e}^e$ are the starting and ending probabilities and p_m^b is anchor background score.

Following the previous practices, we also perform score fusion and redundant proposal suppression to further obtain final results. Specifically, in order to make full use of various predicted scores for each proposal φ_n , we fuse its boundary probabilities and action-background scores of matching anchor by multiplication. The confidence score p^f can be defined as :

$$p^f = p_{t'_s}^s \cdot p_{t'_e}^e \cdot \widehat{p}_m^a \cdot \widehat{p}_m^a \cdot (1 - p_m^b). \quad (13)$$

Hence, the final proposal set as

$$\psi = \{\varphi_n = (t_s, t_e, p^f)\}_{n=1}^N. \quad (14)$$

Moreover, we also use the Soft-NMS algorithm for post-processing to remove the proposals which highly overlap with each other.

Experiments

Datasets and Evaluation Metrics

ActivityNet-v1.3 is a large-scale video dataset for action recognition and temporal action detection tasks. It contains 10K training, 5k validation, and 5k testing videos with 200 action categories, and the ratio of training, validation and testing sets is 2:1:1.

THUMOS14 contains 200 validation untrimmed videos and 213 test untrimmed videos, including 200 action categories. This dataset is challenging due to the large variations in the frequency and duration of action instances across videos.

Evaluation Metrics. Temporal action proposal generation aims to produce high-quality proposals with high tIoU, which have a high recall rate. To evaluate quality of proposals, Average Recall (**AR**) is the average recall rate under specified tIoU thresholds. Following the standard protocol, we use thresholds set [0.5:0.05:0.95] on ActivityNetv1.3 and [0.5:0.05:1.0] on THUMOS14. To evaluate the performance of temporal action detection task, mean Average Precision (**mAP**) under multiple tIoU is the widely-used evaluation metric. On ActivityNet-v1.3, the tIoU thresholds are set to {0.5, 0.75, 0.95}, and we also test the average mAP of tIoU thresholds between 0.5 and 0.95 with step of 0.05. On THUMOS14, these tIoU thresholds are set to {0.3, 0.4, 0.5, 0.6, 0.7}.

Method	@50	@100	@200	@500	@1000
TAG	18.6	29.0	39.6	-	-
CTAP	32.5	42.6	52.0	-	-
BSN	37.5	46.1	53.2	61.4	65.1
MGG	39.9	47.8	54.7	61.4	64.6
BMN	39.4	47.7	54.8	62.2	65.5
BSN++	42.4	49.8	57.6	65.2	66.8
TCANet	42.1	50.5	57.1	63.61	66.9
RTD-Net	41.1	49.0	56.1	62.9	-
Ours	45.5	53.6	60.0	67.0	69.8

Table 1: Performance comparison with state-of-the-art proposal generation methods on test set of THUMOS14 in terms of AR@AN.

Method	Feature	0.3	0.4	0.5	0.6	0.7
SST	TSN	-	-	23.0	-	-
TURN	TSN	44.1	34.9	25.6	-	-
SSN	TSN	51.9	41.0	29.8	-	-
BSN	TSN	53.5	45.0	36.9	28.4	20.0
MGG	TSN	53.9	46.8	37.4	29.5	21.3
DBG	TSN	57.8	49.4	39.8	30.2	21.7
BMN	TSN	56.0	47.4	38.8	29.7	20.5
G-TAD	TSN	54.5	47.6	40.2	30.8	23.4
BSN++	TSN	59.9	49.5	41.3	31.9	22.8
TCANet	TSN	60.6	53.2	44.6	36.8	26.7
Ours	TSN	66.5	60.0	51.6	41.0	29.2
BU-TAL	I3D	53.2	48.5	42.8	33.8	20.8
P-GCN	I3D	63.6	57.8	49.1	-	-
AFSD	I3D	67.3	62.4	55.5	43.7	31.1
RTD-Net	I3D	53.9	48.9	42.0	33.9	23.4
G-TAD+P-GCN	I3D	66.4	60.4	51.6	37.6	22.9
RTD+P-GCN	I3D	68.3	62.3	51.9	38.8	23.7
MUSES	I3D	68.9	64.0	56.9	46.3	31.0
Ours+P-GCN	I3D	69.8	62.9	52.0	39.8	24.0
Ours+MUSES	I3D	71.5	67.0	60.0	48.9	33.0

Table 2: Performance comparison with state-of-the-art action detection methods on test set of THUMOS14, in terms of mAP (%) at different tIoU thresholds.

Implementation Details

Feature Encoding. Following previous works (Lin et al. 2019; Xu et al. 2020), we adopt the TSN (Wang et al. 2016) and I3D (Simonyan and Zisserman 2014) for feature encoding. For THUMOS14, the interval σ is set to 8 and 5 for I3D and TSN respectively. We crop each video feature sequence with overlapped windows of size $T = 256$ and stride 128. As for ActivityNet-1.3, the sampling frame stride is 16, and each video feature sequence is rescaled to $T = 100$ snippets using linear interpolation.

Training and Inference. The number of layers in Boundary Prediction module is 12. Due to the limit of computation resource, we apply 1D Conv for dimension reduction, then take the features as the input to the Boundary Prediction module and Action-Background Interaction module. For each anchor, we use sampling points $N = 32$. For post-processing module, we set adjustment thresholds $\alpha_1 = 0.9$ and $\alpha_2 = 0.8$. We train our model from scratch using the Adam optimizer and the learning rate is set to 10^{-4} and decayed by a factor of 0.1 after every 10 epoch.

Method	AR@1 (val)	AR@100 (val)	AUC (val)
CTAP	-	73.2	65.7
BSN	32.2	74.2	66.2
MGG	-	75.5	66.4
BMN	-	75.0	67.0
BSN++	34.3	76.5	68.3
TCANet	34.6	76.1	68.1
RTD-Net	32.8	73.1	65.7
Ours	35.2	76.6	68.7

Table 3: Performance comparison with state-of-the-art proposal generation methods on validation set of ActivityNet-1.3 in terms of AUC and AR@AN.

Method	0.5	0.75	0.95	Average
Singh et al.	34.5	-	-	-
SCC	40.0	17.9	4.7	21.7
CDC	45.3	26.0	0.20	23.8
R-C3D	26.8	-	-	-
BSN	46.5	30.0	8.0	30.0
BMN	50.1	34.8	8.3	33.9
GTAD	50.4	34.6	9.0	35.1
BSN++	51.3	35.7	8.3	34.9
TCANet w/ BSN	51.9	34.9	7.5	34.4
RTD-Net	46.4	30.5	8.6	30.5
Ours	53.2	36.2	10.6	35.5

Table 4: Performance comparison with state-of-the-art action detection methods on validation set of ActivityNet-1.3, in terms of mAP (%) at different tIoU thresholds and the average mAP.

Comparison with State-Of-the-Arts

Here we compare our BCNet with the existing state-of-the-art methods on ActivityNet-v1.3 and THUMOS14. For fair comparisons, we adopt the same two-stream features used by previous methods in our experiments.

Results on THUMOS14. BCNet is compared with state-of-the-art methods in Table 1 and Table 2, where our method improves the performance significantly for both temporal action proposal generation and action detection. For the temporal action proposal generation task, results are shown in Table 1, which demonstrate that BCNet outperforms state-of-the-art methods in terms of AR@AN with AN varying from 50 to 1000.

For the temporal action detection task, the proposed BCNet also achieves superior results, as shown in Table 2. The performance of our method exceeds state-of-the-art proposal generation methods by a big margin at different tIoU thresholds. Specially, BCNet based on TSN feature reaches an mAP of 51.6% at IoU 0.5. Besides, the performance of BCNet can be further boosted when it is combined with proposal post-processing methods: P-GCN (Zeng et al. 2019) and MUSES (Liu et al. 2021a). Now BCNet reaches **60.0%** at IoU 0.5, outperforming all the other methods. This signifies the advantage of BCNet proposals regardless of post-processing.

Results on ActivityNet-v1.3. In Table 3 and Table 4,

Method	w/ BC	0.3	0.4	0.5	0.6	0.7
BMN*	-	59.5	54.3	45.1	35.3	24.8
BMN*	✓	62.5 (↑ 3.0)	56.3 (↑ 2.0)	47.6 (↑ 2.5)	37.2 (↑ 1.9)	26.3 (↑ 1.5)
GTAD*	-	58.4	52.1	43.5	33.3	23.2
GTAD*	✓	60.5 (↑ 2.1)	54.1 (↑ 2.0)	45.7 (↑ 2.2)	35.1 (↑ 1.8)	24.5 (↑ 1.3)
Ours	-	63.2	58.7	51.2	39.9	28.3
Ours	✓	66.5 (↑ 3.3)	60.0 (↑ 1.3)	51.6 (↑ 0.4)	41.0 (↑ 1.1)	29.2 (↑ 0.9)

Table 5: The effectiveness of the Background Constraint (BC). * indicates our implementation with the publicly available code.

we compare the proposed BCNet with other methods on ActivityNet-v1.3. For the temporal action proposal generation task, as shown in Table 3, the performance of BCNet again outperforms state-of-the-art proposal generation methods in terms of AR@AN with AN varying from 1 to 100 and AUC. Especially when AN equals 1, we achieve 35.2% regarding the AR metric, which indicates that top-1 proposal has high quality. For the temporal action detection task, as summarized in Table 4, our method achieves notable improvements on mAP over other proposal generation methods such as BMN (Lin et al. 2019) and G-TAD (Xu et al. 2020) at all tIoU thresholds. When tIoU is 0.95, the mAP we obtain is 10.6%, indicating that the confidence of the generated proposals are more reliable.

Ablation Study

In this section, we conduct ablation studies on THUMOS14 to verify the effectiveness of each component in BCNet.

Multi-level ABI module. We perform ablation studies to verify the effectiveness of multi-level interaction in ABI module. Frame-level interaction is designed to generate features of action and background. Here, the ablation experiment demonstrates the necessity of frame-level interaction as shown in Table 6. Compared with single-level ABI module that only has a clip-level interaction, multi-level ABI module is improved by 3.6% at tIoU 0.3.

Frame	Clip	0.3	0.4	0.5	0.6	0.7
-	✓	62.9	57.9	49.6	40.4	28.6
✓	✓	66.5	60.0	51.6	41.0	29.2

Table 6: The effect of ABCNet in frame-level and clip-level.

The effectiveness of Background Constraint. We perform ablation studies to verify the effectiveness of the background constraint idea. To validate the generalizability of our proposed background constraint idea, we add it to the BMN, GTAD. The experimental results are shown in Table 5, which reveals that background constraint can also significantly improve the performance of existed methods.

Architecture of ABI module. We perform ablation studies to verify the effectiveness of the architecture of ABI module. To generate reliable confidence of proposal, ABI module is designed by exploiting rich information of action and background. Our proposed ABI module consists of two key units: self-attention unit and difference-attention unit. Results are shown in Table 7. The self-attention unit can

improve the performance by a large margin (almost 4.5%) at tIoU 0.5. Difference-attention unit also brings significant improvement at tIoU 0.3, as the inconsistency of action and background is captured between action and background.

Self.	Diff.	0.3	0.4	0.5	0.6	0.7
-	-	59.3	55.8	46.7	36.6	25.2
✓	-	63.2	58.7	51.2	39.9	28.3
✓	✓	66.5	60.0	51.6	41.0	29.2

Table 7: The effectiveness of ABI module.

Analysis on Runtime

To verify the efficiency of our BCNet, we report the latency of our method on THUMOS14. For the fair comparisons with other models, we measure the latency under the same environment (a single NVIDIA 2080Ti GPU). We use a batch size of 1 to measure the latency on the full testing set and report the average time. As shown in Table 8, our BCNet achieve the best mAP with smallest latency (141ms *v.s.* 298ms, 330ms). The main reason is that our model generates fewer proposals than these methods, which helps our model run faster.

Method	0.3	0.4	0.5	0.6	0.7	Latency
BMN	56.0	47.4	38.8	29.7	20.5	330ms
GTAD	54.5	47.6	40.2	30.8	23.4	298ms
Ours	66.5	60.0	51.6	41.0	29.2	141ms

Table 8: Quantitatively analysis on latency. The smaller latency represent higher efficiency.

Conclusion

In this paper, we introduce a Background Constraint concept, which can be integrated easily with existing TAPG method. Based on this concept, we propose a Background Constraint Network, which consists of multiple attention units *i.e.*, self-attention unit, cross-attention unit, and difference-attention unit, and generates high-quality proposals by exploiting inconsistency between action and background. Extensive experiments show that our model achieves new state-of-the-art performance in temporal action proposal generation and action detection on THUMOS14 and ActivityNet1.3 datasets.

Acknowledgments

The work was funded by National Natural Science Foundation of China (U1711265) and Heilongjiang Province Science Foundation (2020ZX14A02).

References

- Caba Heilbron, F.; Escorcia, V.; Ghanem, B.; and Carlos Niebles, J. 2015a. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 961–970.
- Caba Heilbron, F.; Escorcia, V.; Ghanem, B.; and Carlos Niebles, J. 2015b. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 961–970.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 213–229. Springer.
- Carreira, J.; and Zisserman, A. 2017. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*.
- Chen, P.; Gan, C.; Shen, G.; Huang, W.; Zeng, R.; and Tan, M. 2019. Relation Attention for Temporal Action Localization. *IEEE Transactions on Multimedia*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Gao, J.; Chen, K.; and Nevatia, R. 2018. Ctap: Complementary temporal action proposal generation. In *Proceedings of the European conference on computer vision (ECCV)*, 68–83.
- Gao, J.; Shi, Z.; Li, J.; Wang, G.; Yuan, Y.; Ge, S.; and Zhou, X. 2020. Accurate Temporal Action Proposal Generation with Relation-Aware Pyramid Network. *arXiv preprint arXiv:2003.04145*.
- Gao, J.; Yang, Z.; Sun, C.; Chen, K.; and Nevatia, R. 2017. Turn tap: Temporal unit regression network for temporal action proposals. In *ICCV*.
- Jiang, Y.-G.; Liu, J.; Zamir, A. R.; Toderici, G.; Laptev, I.; Shah, M.; and Sukthankar, R. 2014. THUMOS challenge: action recognition with a large number of classes (2014). URL <http://csrcv.ucf.edu/THUMOS14>.
- Korbar, B.; Tran, D.; and Torresani, L. 2019. Scsampler: Sampling salient clips from video for efficient action recognition. In *Int. Conf. Comput. Vis.*, 6232–6242.
- Lee, P.; Uh, Y.; and Byun, H. 2020. Background suppression network for weakly-supervised temporal action localization. In *Proceedings of the AAAI conference on artificial intelligence*, 11320–11327.
- Lee, P.; Wang, J.; Lu, Y.; and Byun, H. 2020. Weakly-supervised temporal action localization by uncertainty modeling. *arXiv preprint arXiv:2006.07006*.
- Lin, C.; Li, J.; Wang, Y.; Tai, Y.; Luo, D.; Cui, Z.; Wang, C.; Li, J.; Huang, F.; and Ji, R. 2020. Fast Learning of Temporal Action Proposal via Dense Boundary Generator. In *AAAI*, 11499–11506.
- Lin, J.; Gan, C.; and Han, S. 2019. TSM: Temporal Shift Module for Efficient Video Understanding. In *Int. Conf. Comput. Vis.*
- Lin, T.; Liu, X.; Li, X.; Ding, E.; and Wen, S. 2019. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, 3889–3898.
- Lin, T.; Zhao, X.; Su, H.; Wang, C.; and Yang, M. 2018. Bsn: Boundary sensitive network for temporal action proposal generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 3–19.
- Liu, X.; Hu, Y.; Bai, S.; Ding, F.; Bai, X.; and Torr, P. H. 2021a. Multi-shot Temporal Event Localization: a Benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12596–12606.
- Liu, Y.; Ma, L.; Zhang, Y.; Liu, W.; and Chang, S.-F. 2019. Multi-granularity generator for temporal action proposal. In *CVPR*, 3604–3613.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*.
- Liu, Z.; Luo, D.; Wang, Y.; Wang, L.; Tai, Y.; Wang, C.; Li, J.; Huang, F.; and Lu, T. 2020. TEINet: Towards an Efficient Architecture for Video Recognition. In *AAAI*, 11669–11676.
- Nguyen, P. X.; Ramanan, D.; and Fowlkes, C. C. 2019. Weakly-supervised action localization with background modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5502–5511.
- Shou, Z.; Wang, D.; and Chang, S.-F. 2016. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*.
- Simonyan, K.; and Zisserman, A. 2014. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 568–576.
- Su, H.; Gan, W.; Wu, W.; Yan, J.; and Qiao, Y. 2020. BSN++: Complementary Boundary Regressor with Scale-Balanced Relation Modeling for Temporal Action Proposal Generation. *arXiv preprint arXiv:2009.07641*.
- Tan, J.; Tang, J.; Wang, L.; and Wu, G. 2021. Relaxed Transformer Decoders for Direct Action Proposal Generation. *arXiv preprint arXiv:2102.01894*.
- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Int. Conf. Comput. Vis.*
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wang, L.; Xiong, Y.; Wang, Z.; Qiao, Y.; Lin, D.; Tang, X.; and Van Gool, L. 2016. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, 20–36. Springer.

Wang, L.; Yang, H.; Wu, W.; Yao, H.; and Huang, H. 2021. Temporal Action Proposal Generation with Transformers. *arXiv preprint arXiv:2105.12043*.

Wu, W.; He, D.; Lin, T.; Li, F.; Gan, C.; and Ding, E. 2021a. MVFNet: Multi-View Fusion Network for Efficient Video Recognition. In *AAAI*.

Wu, W.; He, D.; Tan, X.; Chen, S.; and Wen, S. 2019. Multi-Agent Reinforcement Learning Based Frame Sampling for Effective Untrimmed Video Recognition. In *Int. Conf. Comput. Vis.*

Wu, W.; He, D.; Tan, X.; Chen, S.; Yang, Y.; and Wen, S. 2020. Dynamic Inference: A New Approach Toward Efficient Video Action Recognition. In *Proceedings of CVPR Workshops*, 676–677.

Wu, W.; Zhao, Y.; Xu, Y.; Tan, X.; He, D.; Zou, Z.; Ye, J.; Li, Y.; Yao, M.; Dong, Z.; et al. 2021b. DSANet: Dynamic Segment Aggregation Network for Video-Level Representation Learning. In *Proc. ACM MM*.

Xu, M.; Zhao, C.; Rojas, D. S.; Thabet, A.; and Ghanem, B. 2020. G-TAD: Sub-Graph Localization for Temporal Action Detection. In *CVPR*, 10156–10165.

Yuan, J.; Ni, B.; Yang, X.; and Kassim, A. A. 2016. Temporal action localization with pyramid of score distribution features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3093–3102.

Zeng, R.; Huang, W.; Tan, M.; Rong, Y.; Zhao, P.; Huang, J.; and Gan, C. 2019. Graph convolutional networks for temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7094–7103.

Zhao, Y.; Xiong, Y.; Wang, L.; Wu, Z.; Tang, X.; and Lin, D. 2017. Temporal action detection with structured segment networks. In *ICCV*, 2914–2923.