

# AEI: Actors-Environment Interaction with Adaptive Attention for Temporal Action Proposals Generation

Khoa Vo<sup>1</sup>  
khoavoho@uark.edu

Hyekang Joo\*<sup>2</sup>  
hkjoo@cs.umd.edu

Kashu Yamazaki\*<sup>1</sup>  
kyamazak@uark.edu

Sang Truong<sup>1</sup>  
sangt@uark.edu

Kris Kitani<sup>3</sup>  
kkitani@cs.cmu.edu

Minh-Triet Tran<sup>4,5</sup>  
tmtriet@hcmus.edu.vn

Ngan Le<sup>1</sup>  
thile@uark.edu

<sup>1</sup> AICV Lab, University of Arkansas  
Fayetteville, AR, USA

<sup>2</sup> University of Maryland  
College Park, MD, USA

<sup>3</sup> Carnegie Mellon University  
Pittsburgh, PA, USA

<sup>4</sup> University of Science, VNU-HCM  
Ho Chi Minh City, Vietnam

<sup>5</sup> Vietnam National University  
Ho Chi Minh City, Vietnam

---

## Abstract

Humans typically perceive the establishment of an action in a video through the interaction between an actor and the surrounding environment. An action only starts when the main actor in the video begins to interact with the environment, while it ends when the main actor stops the interaction. Despite the great progress in temporal action proposal generation, most existing works ignore the aforementioned fact and leave their model learning to propose actions as a black-box. In this paper, we make an attempt to simulate that ability of a human by proposing Actor Environment Interaction (**AEI**) network to improve the video representation for temporal action proposals generation. AEI contains two modules, i.e., perception-based visual representation (PVR) and boundary-matching module (BMM). PVR represents each video snippet by taking human-human relations and humans-environment relations into consideration using the proposed **adaptive attention mechanism**. Then, the video representation is taken by BMM to generate action proposals. AEI is comprehensively evaluated in ActivityNet-1.3 and THUMOS-14 datasets, on temporal action proposal and detection tasks, with two boundary-matching architectures (i.e., CNN-based and GCN-based) and two classifiers (i.e., Unet and P-GCN). Our AEI robustly outperforms the state-of-the-art methods with remarkable performance and generalization for both temporal action proposal generation and temporal action detection. Source code is available at <sup>1</sup>.

---

© 2021. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

<sup>1</sup><https://github.com/vhvkhoa/TAPG-AgentEnvInteration.git>

\*These authors contributed equally.

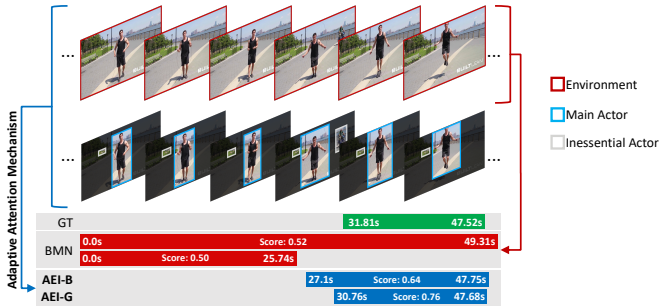


Figure 1: TAPG comparison. The existing approaches (e.g., BMN [26]) apply a network backbone to the entire spatial domain (red boxes); Our AEI takes main actor(s) into consideration via AAM. Our AEI is implemented with CNN-based BMM (AEI-B) and GCN-based BMM (AEI-G).

## 1 Introduction

Temporal action proposals generation (TAPG) is one of the most important problems in video analysis and video understanding [1, 3, 4, 36]. Particularly, TAPG is the fundamental step for other downstream tasks, including temporal action detection [9, 20], action recognition [21], and video dense captioning [22]. Given an untrimmed video, TAPG aims to propose temporal segments with specific starting and ending timestamps for each action of interest appearing in the video.

Recently, state-of-the-art (SOTA) methods [23, 25, 26, 38] follow a paradigm where a set of possible starting and ending timestamps of all actions are detected separately, then a proposal evaluation module is employed to evaluate every possible pair of starting and ending timestamps by predicting its confidence score. A non-maximum suppression (NMS) [4, 30] function is finally used to eliminate redundant candidate proposals based on their confidence scores and overlapping metrics.

As we observe, a human has a capacity to perceive an action being established in a video [4, 15] in two steps. First, the main actors at each temporal period are identified; then, the interactions between main actors and the environment are observed to specify when the action starts and ends. Despite good achievements on benchmarking datasets [9, 20], the SOTA approaches [23, 26, 38] disregard the above perception process of humans by only applying a backbone network (pre-trained on action recognition task) to extract the video representation, leading to a potential loss of some proposals. For instance, in Fig. 1, the works in the literature take the whole spatial region of video frames (e.g., red boxes) to propose action intervals; this, however, may lead to inaccurate results because the background occupies much bigger region than the actor who performs the action (e.g., blue boxes). In Fig. 1, the "rope skipping" action can trick an action proposal model into missing the time at which this action starts or ends due to a subtle difference in shape between "jumping" and "standing".

In this paper, we propose a novel **Actor Environment Interaction network (AEI)** in an attempt to simulate and explore the capability of human-perception process. Our AEI consists of a visual representation module (PVR) and a boundary-matching module (BMM). The PVR is comprised of three components: (i) environment spectator; (ii) actors spectator; and (iii) actors-environment interaction spectator. As illustrated in Fig. 1, the environment spectator processes the entire spatial dimensions of a snippet (red boxes) to capture the global environmental information. The actors spectator focuses on actors (blue and light gray boxes

around humans) to capture local appearance and motion information of actors. Additionally, we introduce a novel **adaptive attention mechanism (AAM)** in the actors spectator to select the main actors (blue boxes) who mainly commit the action. Given a video snippet, the features corresponding to the global environment and the local main actors are first extracted by the first two spectators. The relationship between the environment and the main actors is then modeled by the third component (i.e., the actors-environment interaction spectator).

**Our contributions are summarized as follows:**

- We propose a video representation network, AEI, which follows the human-perception process to understand human action.
- We introduce a novel *adaptive attention mechanism (AAM)* that simultaneously selects main actors and eliminates inessential actor(s) and then extracts semantic relations between main actors.
- We investigate the effectiveness of the proposed AEI by implementing the BMM under two network architectures: CNN-based and GCN-based.
- Our proposed AEI network achieves the SOTA performance on common benchmarking datasets of ActivityNet-1.3 and THUMOS-14 in both TAPG and TAD tracks with a large margin compared to the previous works.

## 2 Related Works

### Temporal Action Proposal Generation (TAPG)

TAPG aims to propose intervals that tightly contain actions in a long untrimmed video. Previous works can be divided into two main groups: anchor-based and boundary-based. Anchor-based methods [8, 11, 17, 35, 36], which are inspired by anchor-based object detection methods in 2D images [24, 33, 34], predefine a set of fixed segments and try to fit them into ground-truth action segments in the video. Although a regression network may be applied in some of those methods to refine the proposals, a finite number of anchors cannot fit all ground-truth actions with diverse lengths. Boundary-based methods [23, 25, 26, 38, 40, 41, 45] address this problem by localizing the starting and ending timestamps of all actions appearing in the video and matching them by a boundary-matching module. Our boundary-matching module belongs to the second group.

### Attention Networks

Attention Networks (AN) have a long history in the artificial neural networks literature [13]. We can divide AN into two main groups: Soft-Attention Networks (Soft-AN) and Hard-Attention Networks (Hard-AN). [10] was one of the first Soft-AN that was applied to machine translation. Because of its differentiable architecture, which helps the whole model learn in an end-to-end fashion, Soft-AN has become an essential component in a large number of applications (e.g., speech [7], NLP [10], computer vision [6]). Hard-AN was first introduced in [44] and [8] for digit and object classifications, respectively. Hard-AN aims to mask out irrelevant elements of the inputs to reduce the distractions. This is an advanced benefit over Soft-AN; however, Hard-AN in [44] is indifferentiable. Recently, [32] proposes a Hard-AN that can be trained by normal gradient back-propagation, with a fundamental observation that the L2-norm values of more important features are usually higher than those of less important features in a feature map. In this work, we propose AAM, a module to leverage both the differentiable Hard-AN [32] and the self-attention network [39] to select the main actors of the video and to learn the relations between main actors, respectively.

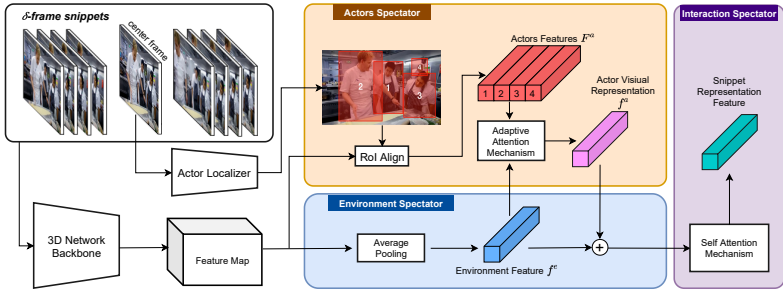


Figure 2: The architecture of our proposed PVR. Given a  $\delta$ -snippet, the corresponding snippet visual representation is obtained by three modules: (i) environment spectator to extract global environment feature; (ii) actors spectator to extract local actor representation; and (iii) actors-environment interaction spectator to model the relationship between the environment feature and the actor feature.

### 3 Our Proposed AEI

Given an input video  $\mathcal{V} = \{v_i\}_{i=1}^N$ , where  $N$  is the number of frames, we follow the common paradigm from previous works to divide it into a sequence of *snippets*, each of which consists of  $\delta$  consecutive frames from the video, resulting in a total of  $T = \lceil \frac{N}{\delta} \rceil$  snippets. Let  $\phi(\cdot)$  be an encode function to extract visual representation of a  $\delta$ -frame snippet  $s_i$ , the entire video is presented as:

$$f_i = \phi(s_i), \text{ and } \mathcal{F} = \{f_i\}_{i=1}^T \quad (1)$$

Prior works [25, 26, 45] employ a pre-trained backbone network (e.g., C3D network [19] or Two-Stream network [37]) to model  $\phi(\cdot)$ . However, simply applying those networks for video representation may have some drawbacks as mentioned in Section 1. In Section 3.1, our proposed perception-based visual representation (PVR) is discussed as an alternative to the former strategy. Then, boundary-matching module for temporal action proposals generation is discussed in Section 3.2.

#### 3.1 Perception-based Visual Representation (PVR)

The PVR module aims to extract video visual representation based on how a human perceives an action, i.e., identifying the main actors at each temporal period and interactions between main actors and the environment to specify when the action starts and ends. PVR consists of three main components: (i) environment spectator; (ii) actors spectator; and (iii) actors-environment interaction spectator. The overall architecture of PVR is shown in Fig. 2.

**(i) Environment Spectator** aims to extract global semantic information of the input  $\delta$ -frame snippet. To extract both the spatial and temporal details of the snippet, we adopt a 3D network pre-trained on action recognition benchmarking datasets as a backbone feature extractor. The snippet is processed through all convolutional blocks of the 3D network to obtain a feature map  $\mathcal{M}$ ; then, an average pooling operator is employed to produce a spatio-temporal feature vector  $f^e$ .

**(ii) Actors Spectator** aims to semantically extract main actor(s) representation. An action cannot happen in the absence of a human (actor). However, when an action occurs, it does not necessarily signal that every actor in the scene has committed the action. First, the actors spectator detects all existing actors in the snippet by an *actor localization* module.




Figure 3: Illustration of our proposed AAM. Given an environment feature  $f^e$  and a set of actor features  $F^a$ , this module aims to select main actor features, followed by fusing arbitrary main actor features to obtain an actor visual representation  $f^a$ .

**Algorithm 1:** Adaptive Attention Mechanism (AAM) to extract representation of main actors in a snippet.

---

**Data:** Feature vector  $f^e$  and features set  $F^a$  represent environment and all actors that appear in input snippet, respectively.

**Result:** Feature vector  $f^a$  represents main actors in input snippet.

```

 $\hat{f}^e \leftarrow \text{MLP}_{\theta_e}(f^e)$  // embed  $f^e$  to common space with every  $f_i^a$  in  $F^a$ 
set  $S^a$  and  $\tilde{F}^a$  to empty list //  $S^a$  will store scores of every actor
set  $\tilde{F}^a$  to empty list //  $\tilde{F}^a$  will store selected main actors
for each  $f_i^a$  in  $F^a$  do
   $\hat{f}_i^a \leftarrow \text{MLP}_{\theta_a}(f_i^a)$  // embed  $f_i^a$  to common space with  $f^e$ 
   $s_i^a \leftarrow \|\hat{f}_i^a \oplus \hat{f}^e\|_2$  to  $S^a$  // compute main actor score corresponding to  $f_i^a$ 
  append  $s_i^a$  to  $S^a$ 
 $S^a \leftarrow \text{softmax}(S^a)$  // rescale scores to sum up to 1.0
 $\tau \leftarrow \frac{1}{|S^a|}$  // compute adaptive threshold  $\tau$ 
for each  $f_i^a$  in  $F^a$  do
  if  $S_i^a > \tau$  then
    append  $f_i^a$  to  $\tilde{F}^a$  // select  $f_i^a$  if its score higher than  $\tau$ 
 $f^a \leftarrow \text{self\_attention}(\tilde{F}^a)$  // fuse main actor features by self-attention [59]

```

---

Then, an *adaptive attention mechanism* (AAM) is proposed to adaptively select an arbitrary number of main actor(s) and extract their mutual relationships to represent them as a single feature vector.

**Actor Localization:** To localize all actors in a  $\delta$ -frame snippet, we apply a human detector onto the middle frame of it with the assumption that, with a small  $\delta$ , the actors would not move fast enough to be mis-located. We denote  $\mathcal{B} = \{b_i\}_{i=1}^{N_B}$  as a set of detected human bounding boxes, where  $N_B \geq 0$ . Afterwards, each of the detected bounding boxes,  $b_i$ , is aligned onto feature map  $\mathcal{M}$  (obtained by the 3D network backbone from environment spectator) using RoIAlign [16] and then average-pooled into a single feature vector  $f_i^a$ . Finally, we obtain a set of actor features  $F^a = \{f_i^a\}_{i=1}^{N_B}$ .

**Adaptive Attention Mechanism (AAM):** Given  $N_B$  detected actors, there are only a few of detected actors (called main actors) who actually contribute to the action if it presents. Because the number of main actors is unknown and continuously changes throughout the input video, we propose an adaptive attention mechanism (AAM) that inherits the merits from adaptive hard attention to select an arbitrary number of main actors and a soft self-attention mechanism [59] to extract relationships among them. AAM is described by pseudocode in Algorithm 1 and illustrated in Fig. 3; more details are provided in the supplementary.

(iii) **Actors-Environment Interaction Spectator** This module aims to model the rela-

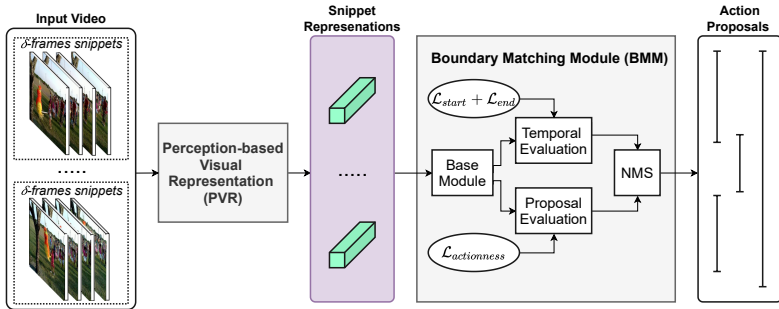


Figure 4: The overall architecture of our proposed AEI, consisting of perception-based visual representation module (PVR), and boundary-matching module (BMM).

tions between environment feature  $f^e$  and actors representation feature  $f^a$ , and then combine them into a single feature  $f$ . Herein, we employ the self-attention model [59] where  $f^e$  and  $f^a$  are the inputs. We denote  $f_i$  as a visual representation for snippet  $s_i$ .

### 3.2 Boundary-Matching Module (BMM)

BMM is responsible for generating action proposals, which are boundary-pairs of every possible action of interest appearing in the video. Our BMM contains three components: base module, temporal evaluation module, and proposal evaluation module as illustrated in Fig. 4. The base module aims to model the semantic relationship between snippets. The temporal evaluation module assesses each snippet  $0 \leq i \leq T$  in the video to estimate probabilities that any action starts or ends there, corresponding to  $P_i^S$  and  $P_i^E$ , respectively. Meanwhile, the proposal evaluation module evaluates every interval  $[i, j]$  in the video to estimate its actionness score, corresponding to  $P_{i,d}^A$ , where  $d = j - i$ .

At the inference stage, we search through  $P^S$  and  $P^E$  to select temporal locations  $i$  whose  $P_i^S$  or  $P_i^E$  are local maximums to form sets of potential starting and ending temporal locations, respectively. Then, starting and ending locations  $(s, e)$  from those lists that satisfy timing constraint (e.g.  $s \leq e \leq T$ ) are paired and become a candidate proposal with a score  $s = P_s^S \cdot P_e^E \cdot P_{s,e-s}^A$ . Based on the timestamps and scores of candidate proposals, we finally apply NMS [9, 60] to produce the final set of action proposals.

In our paper, we have conducted BMM under two different network architectures: CNN-based and GCN-based. Our CNN-based BMM, called **AEI-B**, is leveraged by [24] where the base module is comprised of 1D convolutional layers to learn and extract the temporal relations between snippets. On the other hand, our GCN-based BMM, called **AEI-G**, is leveraged by [45] to extract not only local relations, but also the relations of snippets that share close semantic features.

### 3.3 Training Methodology

Given a list of  $N_G$  ground-truth action segments  $G = \{g_i = (g_i^s, g_i^e)\}_{i=1}^{N_G}$  of input video  $\mathcal{V}$ , we generate the ground-truth starting labels  $L^S \in [0, 1]^T$ , ending labels  $L^E \in [0, 1]^T$  and actionness labels  $L^A \in [0, 1]^{T \times D}$  ( $D$  is a pre-defined maximum proposal length).  $L_i^S$  (or  $L_i^E$ ) carrying a value of 1 means that there is a ground-truth starting (or ending) boundary of any action at  $i$  and vice versa. Likewise,  $L_{i,d}^A$  carrying a value of 1 means that there is a ground-truth action starts at  $i$  with a length of  $d$ . To train our AEI network with the ground-truth labels,

Table 1: **TAPG** comparisons in terms of AR@AN and AUC between our AEI and other SOTA methods on **ActivityNet-1.3**.

Metrics	BSN	MGG	MR	BMN	DBG	BSN++	TSI++	AEN	ABN	AEI-B	AEI-G
	[43]	[44]	[45]	[46]	[47]	[58]	[48]	[49]	[50]		
AR@100 (val)	74.16	74.54	75.27	75.01	76.65	76.52	76.31	75.65	76.72	<b>77.25</b>	<u>77.24</u>
AUC (val)	66.17	66.43	66.51	67.10	68.23	68.26	68.35	68.15	69.16	<u>69.43</u>	<b>69.47</b>
AUC (test)	66.26	66.47	-	67.19	68.57	-	68.85	68.99	69.26	<u>69.94</u>	<b>70.09</b>

we define the loss function  $\mathcal{L}_{AEI}$  as follows:

$$\mathcal{L}_{AEI} = \mathcal{L}_{start}(P^S, L^S) + \mathcal{L}_{end}(P^E, L^E) + \mathcal{L}_{actionness}(P^A, L^A)$$

We use weighted binary log-likelihood loss  $\mathcal{L}_{wb}$  for  $\mathcal{L}_{start}$  and  $\mathcal{L}_{end}$ , which is defined as follows:

$$\mathcal{L}_{wb}(P, L) = \sum_{i=1}^N \left[ \frac{L_i}{N^+} \log P_i + \frac{(1-L_i)}{N^-} \log(1-P_i) \right] \quad (2)$$

where  $N^+$  and  $N^-$  are the number of positives and negatives in ground-truth labels, respectively. Conversely,  $\mathcal{L}_{actionness}(P, L) = \mathcal{L}_{wb}(P, L) + \lambda \mathcal{L}_2(P, L)$ , where  $\mathcal{L}_2$  is the mean squared error loss and  $\lambda$  is set to 10.

To reduce time cost in training phase of our proposed AEI network, actors features set  $F^a$  and environment feature  $f^e$  for actors spectator and environment spectator, respectively, are extracted in advance. Then, the AAM and Interaction Spectator of PVR module is trained with BMM module in an end-to-end fashion.

## 4 Experiments

### Datasets

We conduct experiments on ActivityNet-1.3 [9] and THUMOS-14 [20] datasets. The former contains 20,000 videos with 200 annotated activities while the latter consists of 414 videos with 20 actions. For both datasets, we follow previous works [23, 25, 26] to preprocess videos with the snippet length set to  $\delta = 16$ .

### Metrics

In TAPG, AR@AN and AUC are popular metrics to benchmark the performance of competing methods. The former is the average recall (AR) computed with the average number of proposals (AN) kept by each video. The latter is the score of the area under AR versus the AN curve. In ActivityNet-1.3, AR@100 and AUC are mainly used. On the other hand, in THUMOS-14, only AR@AN is used to compare between methods; however, multiple AN is selected from a list of [50, 100, 200, 500, 1000].

In TAD, both ActivityNet-1.3 and THUMOS-14 use mean Average Precision (mAP) to benchmark methods in this problem. ActivityNet-1.3 uses tIoU thresholds of {0.5, 0.75, 0.95} and average mAP, while THUMOS-14 uses tIoU thresholds of {0.3, 0.4, 0.5, 0.6, 0.7} for evaluation.

### Implementation Details

For all experiments on both ActivityNet-1.3 [9] and THUMOS-14 [20], we employ C3D [19] network pre-trained on Kinetics-400 [21] as the backbone network to extract features from video snippets. The features extracted from C3D backbone have 2048 dimensions.

In the actors spectator, for actor localization, we employ a Faster-RCNN model [52]

Table 2: **TAD** comparisons on **ActivityNet-1.3** in terms of mAP@tIoU and mAP, where the proposals are combined with video-level classification results generated by [43].

Metrics	BSN	BMN	GTAD	P-GCN	MR	ABN	TSI++	GTAN	AEI-B	AEI-G
	[25]	[26]	[45]	[46]	[47]	[48]	[28]	[60]		
0.5	46.5	50.1	50.4	42.9	43.5	51.8	51.2	<b>52.6</b>	52.3	<u>52.4</u>
0.75	30.0	<u>34.8</u>	34.6	28.1	33.9	34.2	<b>35.0</b>	34.1	34.5	34.5
0.95	8.0	8.3	9.0	2.5	9.2	<b>10.3</b>	6.6	8.9	<u>9.7</u>	9.6
Average	30.0	33.9	34.1	27.0	30.1	34.2	34.2	<u>34.3</u>	<b>34.7</b>	<b>34.7</b>

pre-trained on COCO [47] dataset to detect humans as discussed in 3.1. To train our AEI network, we utilize Adam optimizer with the initial learning rate set to 0.0001 and 0.001 for ActivityNet-1.3 and THUMOS-14, respectively.

In TAPG, Soft-NMS [9] is applied in post-processing for all experiments on ActivityNet-1.3, while on THUMOS-14, both Soft-NMS [9] and NMS [49] are evaluated. In TAD, following [25, 26], we use NMS [49] for post-processing on both datasets.

In the following experiments, we highlight the best performance in **bold** and the second-best performance in *italic*.

## 4.1 Temporal Action Proposal Generation (TAPG)

Table 1 demonstrates the comparison on ActivityNet-1.3 validation and testing sets. Based on the results, it can be observed that the performances of our methods, AEI-B and AEI-G, are standing out against those of SOTA methods in terms of AR@100 and AUC by large margins. Table 3 presents the comparison of SOTA TAPG methods on THUMOS-14 dataset. Compared to SOTA approaches, our AEI obtains better performance on all AR@ANs metrics regardless of the architecture of BMM. From Table 1, 3, we empirically observe that AEI-G, which employs GCN to model the relationship between snippets, marginally surpasses AEI-B on TAPG.

Generalizability is also considered as an important criterion to evaluate a method in TAPG. Following the same experiment setup in [25, 26], our AEI is trained on *Unseen+Seen* and *Seen* training sets, separately, and then is evaluated on the *Seen* and *Unseen* validation sets, separately as illustrated in Fig. 5. The performances on *Seen* validation set is shown in the first two charts, whereas the performances on *Unseen* validation set is given in the last two charts. Fig. 5 shows that our AEI achieves good performances on *Seen* validation set with an acceptable drop on *Unseen* validation set on both training configurations, suggesting that our AEI is highly generalizable to unseen action types.

## 4.2 Temporal Action Detection (TAD)

Following the experiment settings in [25, 26, 45], we adopt top-1 video-level classification results generated by the method in [43] on ActivityNet-1.3 to label the proposals produced by our method. We use top-2 video-level classification results generated by UntrimmedNet [42] to label proposals generated by our method on THUMOS-14.

Table 2 illustrates TAD performance comparison between AEI and other SOTA methods on ActivityNet-1.3 validation set. The results emphasize that our methods outperform SOTA methods in spite of CNN-based BMM or GCN-based BMM. The experiment results on THUMOS-14 test set in Table 4 demonstrate that our AEI-B and AEI-G are superior to other






Figure 5: **Generalizability** evaluation and comparisons with BMN [26] and DBG [23] on ActivityNet 1.3 in terms of AR@100 and AUC. Methods are trained on *Unseen+Seen* (blue columns) and *Seen* training sets (orange columns), respectively; and are evaluated on *Seen* (first two charts) and *Unseen* (last two charts) validation sets.

Table 3: **TAPG** comparisons on **THUMOS-14** in terms of AR@AN, where SNMS represents Soft-NMS [9].

Method	@50	@100	@200	@500	@1000
T-TAG [10]	18.55	29.00	39.61	-	-
CTAP [10]	32.49	42.61	51.97	-	-
BSN [15]	37.46	46.06	53.21	60.64	64.52
MGG [15]	39.93	47.75	54.65	61.36	64.06
BMN [26]	39.36	47.72	54.70	62.07	65.49
DBG+NMS [23]	40.89	49.24	55.76	61.43	61.95
DBG+SNMS [15]	37.32	46.67	54.50	62.21	66.40
TSI++ [28]	42.30	50.51	57.24	63.43	-
MR [10]	44.23	50.67	55.74	-	-
ABN+NMS [10]	44.89	51.86	57.36	61.67	62.59
ABN+SNMS [10]	40.87	49.09	56.24	63.53	67.29
<b>AEI-B+NMS</b>	<b>45.74</b>	<b>52.39</b>	57.74	62.49	63.38
<b>AEI-G+NMS</b>	45.12	<b>52.81</b>	<b>57.94</b>	62.11	63.17
<b>AEI-B+SNMS</b>	44.97	50.13	57.34	<b>64.43</b>	<b>67.78</b>
<b>AEI-G+SNMS</b>	<b>45.31</b>	51.12	<b>58.19</b>	<b>64.58</b>	<b>67.96</b>

Table 4: **TAD** comparisons on **THUMOS-14** in term of mAP@IoU. \* indicates P-GCN classifier (2<sup>nd</sup> group); otherwise, UntrimmedNet classifier.

Method	0.7	0.6	0.5	0.4	0.3
T-TAP [10]	6.3	14.1	24.5	35.3	46.3
BSN [15]	20.0	28.4	36.9	45.0	53.5
BMN [26]	20.5	29.7	38.8	47.4	56.0
MGG [15]	21.3	29.5	37.4	46.8	53.9
DBG [23]	21.7	30.2	39.8	49.4	57.8
GTAD [15]	<b>23.4</b>	30.8	40.2	47.6	54.5
TSI++ [28]	22.4	33.2	42.6	<b>52.1</b>	<b>61.0</b>
GTAN [10]	-	-	38.8	47.2	57.8
<b>AEI-B</b>	<b>23.4</b>	<b>35.9</b>	<b>44.7</b>	<b>52.7</b>	<b>58.7</b>
<b>AEI-G</b>	<b>22.9</b>	<b>34.2</b>	<b>43.4</b>	51.6	57.4
BSN* [15]	-	-	49.1	57.8	63.6
GTAD* [15]	<b>22.9</b>	37.6	51.6	<b>60.4</b>	66.4
MR* [10]	<b>28.5</b>	<b>38.0</b>	45.4	50.7	53.9
<b>AEI-B*</b>	22.4	37.8	<b>52.1</b>	<b>60.6</b>	<b>67.3</b>
<b>AEI-G*</b>	22.3	<b>37.9</b>	<b>52.0</b>	<b>60.4</b>	<b>67.6</b>

SOTA methods on most of the metrics regardless of UntrimmedNet [10] or P-GCN [15] classifiers. From Table 2 and 4, we empirically notice that both AEI-B and AEI-G obtain comparable TAD performance.

### 4.3 Ablation Study

We further conduct a detailed ablation study on THUMOS-14 dataset to evaluate the contributions of different components of the proposed AEI framework. We conduct two ablation studies as shown in Fig. 6 on TAPG and THUMOS-14 in terms of AR@ANs.

First, we evaluate the contribution of each spectator to the overall performance of our proposed PVR (described in Section 3.1), i.e., environment spectator, actors spectator, and interaction spectator. As illustrated in Fig. 6 (a), "Environment spectator only", which only focuses on global information, plays an important role in TAPG, whereas "Actors spectator only", which takes only local information of the main actor(s) into account, achieves adequate performance. "W/o interaction spectator", which withdraws the interaction spectator by simply fusing global and local information using an averaging operation, gives an un-

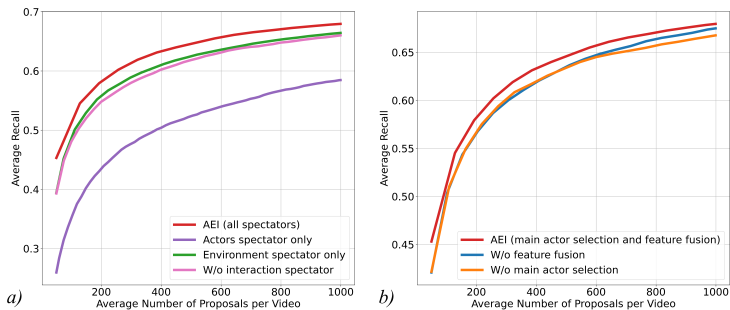


Figure 6: TAPG comparisons on different AEI configurations: (a) either only environment or only actors spectator or both; (b) either only main actor selection or only feature fusion or both.

desired performance that is even lower than "Environment only". The complete proposed model, e.g., "AEI (all spectators)", gives the best result thanks to the interaction spectator adaptively fusing global feature from environment spectator and local feature from actors spectator.

In addition, we also evaluate the effectiveness of main actor selection and feature fusion in our proposed AAM. Fig. 6 (b) shows the performance of the network without each of these components. In the "AEI w/o feature fusion" settings, we use an average pooling layer to fuse features obtained from main actor selection component. As illustrated, both configurations achieve similar performance with AN below 600, while main actor selection component plays a slightly more significant role than feature fusion component. This implies that having an appropriate main actor selection contributes significantly to the entire network.

## Conclusion

In this paper, we proposed a novel actors-environment interaction (AEI) network to simulate human perceiving ability in the temporal action proposals generation. Our AEI contains two modules: perception-based visual representation (PVR) and boundary-matching (BMM). PVR aims to extract visual representation of each snippet. To achieve the human perceiving ability, PVR is equipped with three spectators, each of which learns to perceive input snippet at a unique aspect, e.g. environment, main actors, and actors-environment interactions. An adaptive attention mechanism (AAM) is proposed in actors spectator to select an arbitrary number of main actor(s) in the snippet as well as learning the relationships between them.

Extensive experiments are conducted on ActivityNet-1.3 and THUMOS-14 datasets on TAPG and TAD tasks, which demonstrate that our proposed AEI outperforms SOTA methods regardless of BMM architecture (e.g., CNN-based or GCN-based). These results prove that replicating human perceiving ability in video understanding is a promising track to follow and further explore in the future.

Beside three proposed spectators in PVR, we can include additional spectators to observe human body parts and the interaction between them with objects to better handle egocentric videos, in which the main actor who perform the action does explicitly appear.

## Acknowledgment

This material is based upon work supported by the National Science Foundation under Award No. OIA-1946391; partially funded by Gia Lam Urban Development and Investment Company Limited, Vingroup and supported by Vingroup Innovation Foundation (VINIF) under project code VINIF.2019.DA19.

## References

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [2] V. Bazhanov and R. W. Scholz. Understanding human action. integrating meanings, mechanisms, causes, and contexts. 2016.
- [3] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S. Davis. Soft-nms – improving object detection with one line of code. In *ICCV*, Oct 2017.
- [4] Shyamal Buch, Victor Escorcía, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Niebles. Sst: Single-stream temporal action proposals. In *CVPR*, July 2017.
- [5] Y. Chao, S. Vijayanarasimhan, B. Seybold, D. A. Ross, J. Deng, and R. Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018.
- [6] Sneha Chaudhari, Gungor Polatkan, Rohan Ramanath, and Varun Mithal. An attentive survey of attention models. *arXiv preprint arXiv:1904.02874*, 2019.
- [7] Kyunghyun Cho, Aaron Courville, and Yoshua Bengio. Describing multimedia content using attention-based encoder-decoder networks. *IEEE Transactions on Multimedia*, 17(11):1875–1886, 2015.
- [8] Gamaleldin Elsayed, Simon Kornblith, and Quoc V Le. Saccader: Improving accuracy of hard attention models for vision. In *NeurIPS*, volume 32. Curran Associates, Inc., 2019.
- [9] Bernard Ghanem Fabian Caba Heilbron, Victor Escorcía and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015.
- [10] Andrea Galassi, Marco Lippi, and Paolo Torrioni. Attention in natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [11] Jiyang Gao, Zhenheng Yang, Kan Chen, Chen Sun, and Ram Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. In *ICCV*, Oct 2017.
- [12] Jiyang Gao, Zhenheng Yang, and Ram Nevatia. Cascaded Boundary Regression for Temporal Action Detection. *arXiv e-prints*, art. arXiv:1705.01180, May 2017.
- [13] Jiyang Gao, Kan Chen, and Ram Nevatia. Ctap: Complementary temporal action proposal generation. In *ECCV*, September 2018.

- [14] Jiyang Gao, Runzhou Ge, Kan Chen, and Ram Nevatia. Motion-appearance co-memory networks for video question answering. In *CVPR*, June 2018.
- [15] H. Ghosh. *Computational Models for Cognitive Vision*. Wiley, 2020. ISBN 9781119527862. URL <https://books.google.com/books?id=fnPtDwAAQBAJ>.
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In *ICCV*, Oct 2017.
- [17] Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In *CVPR*, June 2016.
- [18] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE TPAMI*, 20(11):1254–1259, 1998. doi: 10.1109/34.730558.
- [19] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE TPAMI*, 35(1):221–231, 2013. doi: 10.1109/TPAMI.2012.59.
- [20] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://crcv.ucf.edu/THUMOS14/>, 2014.
- [21] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [22] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, 2017.
- [23] Chuming Lin, Jian Li, Yabiao Wang, Ying Tai, Donghao Luo, Zhipeng Cui, Chengjie Wang, Jilin Li, Feiyue Huang, and Rongrong Ji. Fast learning of temporal action proposal via dense boundary generator. *AAAI*, pages 11499–11506, Apr. 2020.
- [24] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *ICCV*, pages 2999–3007, 2017.
- [25] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *ECCV*, September 2018.
- [26] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, October 2019.
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollar, and Larry Zitnick. Microsoft coco: Common objects in context. In *ECCV*, September 2014.
- [28] Shuming Liu, Xu Zhao, Haisheng Su, and Zhilan Hu. Tsi: Temporal scale invariant network for action proposal generation. In *ACCV*, November 2020.
- [29] Yuan Liu, Lin Ma, Yifeng Zhang, Wei Liu, and Shih-Fu Chang. Multi-granularity generator for temporal action proposal. In *CVPR*, June 2019.

- [30] Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. Gaussian temporal awareness networks for action localization. In *CVPR*, pages 344–353, 2019.
- [31] A. Neubeck and L. Van Gool. Efficient non-maximum suppression. In *ICPR*, volume 3, pages 850–855, 2006. doi: 10.1109/ICPR.2006.479.
- [32] Badri Patro and Vinay P Nambodiri. Differential attention for visual question answering. In *CVPR*, pages 7680–7688, 2018.
- [33] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv*, 2018.
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, page 91–99, 2015.
- [35] Alexander Richard and Juergen Gall. Temporal action detection using a statistical language model. In *CVPR*, June 2016.
- [36] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, June 2016.
- [37] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS’14*, page 568–576, Cambridge, MA, USA, 2014. MIT Press.
- [38] Haisheng Su, Weihao Gan, Wei Wu, Junjie Yan, and Yu Qiao. BSN++: complementary boundary regressor with scale-balanced relation modeling for temporal action proposal generation. In *ACCV*, 2020.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, volume 30. Curran Associates, Inc., 2017.
- [40] Khoa Vo, Kashu Yamazaki, Sang Truong, Minh-Triet Tran, Akihiro Sugimoto, and Ngan Le. Abn: Agent-aware boundary networks for temporal action proposal generation. *IEEE Access*, 9:126431–126445, 2021. doi: 10.1109/ACCESS.2021.3110973.
- [41] Viet-Khoa Vo-Ho, Ngan Le, Kashu Kamazaki, Akihiro Sugimoto, and Minh-Triet Tran. Agent-environment network for temporal action proposal generation. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2160–2164, 2021. doi: 10.1109/ICASSP39728.2021.9415101.
- [42] Limin Wang, Yuanjun Xiong, Dahua Lin, and Luc Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [43] Yuanjun Xiong, Limin Wang, Zhe Wang, Bowen Zhang, Hang Song, Wei Li, Dahua Lin, Yu Qiao, Luc Van Gool, and Xiaoou Tang. CUHK & ETHZ & SIAT submission to activitynet challenge 2016. *CoRR*, abs/1608.00797, 2016.
- [44] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015.

- 
- [45] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *CVPR*, June 2020.
- [46] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *ICCV*, pages 7094–7103, 2019.
- [47] Peisen Zhao, Lingxi Xie, Chen Ju, Ya Zhang, Yanfeng Wang, and Qi Tian. Bottom-up temporal action localization with mutual regularization. In *European Conference on Computer Vision*, pages 539–555. Springer, 2020.