

Localizing Unseen Activities in Video via Image Query

Zhu Zhang¹, Zhou Zhao^{1*}, Zhijie Lin¹, Jingkuan Song² and Deng Cai³

¹College of Computer Science, Zhejiang University, China

²University of Electronic Science and Technology of China, China

³Stata Key Lab of CAD&CG, Zhejiang University, China

{zhangzhu, zhaozhou, linzhijie}@zju.edu.cn, {jingkuan.song, dengcai}@gmail.com

Abstract

Action localization in untrimmed videos is an important topic in the field of video understanding. However, existing action localization methods are restricted to a pre-defined set of actions and cannot localize unseen activities. Thus, we consider a new task to localize unseen activities in videos via image queries, named Image-Based Activity Localization. This task faces three inherent challenges: (1) how to eliminate the influence of semantically inessential contents in image queries; (2) how to deal with the fuzzy localization of inaccurate image queries; (3) how to determine the precise boundaries of target segments. We then propose a novel self-attention interaction localizer to retrieve unseen activities in an end-to-end fashion. Specifically, we first devise a region self-attention method with relative position encoding to learn fine-grained image region representations. Then, we employ a local transformer encoder to build multi-step fusion and reasoning of image and video contents. We next adopt an order-sensitive localizer to directly retrieve the target segment. Furthermore, we construct a new dataset ActivityIBAL by reorganizing the ActivityNet dataset. The extensive experiments show the effectiveness of our method.

1 Introduction

With the rapid growth of video data, action understanding has attracted much attention in the computer vision community. Early works mainly focus on action recognition [Simonyan and Zisserman, 2014; Tran *et al.*, 2015] of trimmed video clips that only contain one activity. Recently, some works [Yeung *et al.*, 2016; Shou *et al.*, 2016; Chao *et al.*, 2018] have made great progress in action localization, which retrieves multiple activities from a long, untrimmed video and meanwhile predicts the temporal boundaries of these activities. However, existing action localization methods are restricted to a pre-defined set of actions. Although the action set can be relatively large [Caba Heilbron *et al.*, 2015], it is still difficult to cover various and complex activities in the




Figure 1: Image-Based Activity Localization.

real world. Thus, in practical applications, these approaches are hard to recognize those unseen activities, which do not appear in the training data. But if there is an image query as a reference, we can localize the unseen activities by matching visual contents with respect to objects, appearances, motions and interactions. As shown in Figure 1(a), the activity “people play soccer on the beach” has the similar semantic contents with the image query, including the object “soccer” and the interaction between “people” and “soccer”.

Motivated by these observations, we consider a new task to localize unseen activities in video via image queries, named Image-Based Activity Localization. Specifically, the inputs of this task are one untrimmed video and one image query. The image query depicts an activity that users are interested in. And the video contains one segment semantically corresponding to this activity. Then, we aim to identify the start and end boundaries of the targeted segment, even if this activity has never appeared in training data before.

Localizing unseen activities via image queries is a very challenging task, which faces three inherent difficulties. One is how to eliminate the influence of semantically inessential contents in image queries, such as playing soccer “on the grass” or “on the beach” in Figure 1(a). It requires the precise image representation learning and video-based attention method to highlight the crucial query information. Recently, the self-attention mechanism [Vaswani *et al.*, 2017; Wang *et al.*, 2018] has been applied to capture semantic dependencies in sequential data. Inspired it, we devise a novel region self-attention method with relative position encoding to learn fine-grained image region representations, where the basic region features are extracted using Faster R-CNN [Ren *et al.*, 2015]. Subsequently, we can perform the frame-level

*Zhou Zhao is the corresponding author.

attention mechanism to focus on important regions and ignore unrelated parts.

The second obstacle of this task is how to deal with the fuzzy localizing of inaccurate image queries. Sometimes users may not be able to precisely describe the activities they intend to localize, and only provide an image with close semantics. As shown in Figure 1(b), users aim to retrieve the activity “playing the violin” but give an image query about “playing the guitar”. It requires the multi-level matching of both accurate and rough semantics. Following the transformer architecture [Vaswani *et al.*, 2017], we employ a local transformer encoder with multiple stacked layers to build multi-step fusion and reasoning of image and video contents, which sufficiently explores the query semantics, regardless of the accurate “playing the guitar” or the rough “playing the instrument”. Different from the original transformer, we devise a multi-head local self-attention sublayer to focus on local relation modeling, and apply a more informative fusion sublayer for query understanding.

The last difficulty of this task is how to determine the precise boundaries of target segments, which are often indistinct in videos. We consider leveraging the contextual information to boost the localization accuracy. Concretely, we employ an order-sensitive localizer to directly retrieve temporal boundaries, where forward and backward context is incorporated for the prediction of start and end boundaries, respectively.

Furthermore, we construct a new dataset ActivityIBAL for image-based activity localization by reorganizing the ActivityNet [Caba Heilbron *et al.*, 2015], where we split 200 action classes into three parts: 160 classes for training, 20 classes for validation and the remaining 20 classes for testing. This setting guarantees that the videos of a class used in testing do not appear during the training process.

In summary, the main contributions of this paper are as follows:

- We introduce a new task, named Image-based Activity Localization, which aims to localize unseen activities in untrimmed videos via an image query.
- We propose a self-attention interaction localizer (SAIL) to deal with this challenging task. We first devise region self-attention method to learn fine-grained image region representations, and then employ a local transformer encoder to build multi-step fusion and reasoning of image and video contents. Finally, we adopt an order-sensitive localizer to directly retrieve the target segment.
- We reorganize the ActivityNet dataset to construct a new dataset ActivityIBAL for image-based activity localization, and validate the effectiveness of our SAIL method through extensive experiments.

2 Related Work

In this section, we briefly review some related work on action localization.

Temporal action localization aims to detect the action instances in untrimmed videos. In the fully-supervised fashion, some works [Yeung *et al.*, 2016; Singh *et al.*, 2016] apply recurrent neural networks to capture the temporal dynamics of

video contents. [Yeung *et al.*, 2016] propose an RNN-based agent to observe video frames and decide both where to look next and when to emit a prediction. [Singh *et al.*, 2016] devise two additional streams on motion and appearance for fine-grained action detection. With the development of 3D convolution [Tran *et al.*, 2015], [Shou *et al.*, 2016] employ three segment-based 3D ConvNets to explicitly consider temporal overlap in videos. [Xu *et al.*, 2017] directly encodes the video streams using a 3D fully convolutional network without pre-defined segments. And [Shou *et al.*, 2017] employ temporal upsampling and spatial downsampling operations simultaneously. Furthermore, [Zhao *et al.*, 2017] model the temporal structure of each action instance via a temporal pyramid. [Lin *et al.*, 2017] skip the proposal generation and directly detect action instances based on temporal convolutional layers. And inspired by the Faster R-CNN object detection framework [Ren *et al.*, 2015], [Chao *et al.*, 2018] develop an improved method for temporal action localization.

In the non-fully supervised setting, [Wang *et al.*, 2017] directly learn action recognition from untrimmed videos without the temporal action annotations. And [Nguyen *et al.*, 2018] identify a sparse subset of key segments associated with target actions and fuse them through adaptive temporal pooling. Moreover, [Soomro and Shah, 2017] learn unsupervised action localization as a Knapsack problem.

To go beyond the pre-defined set of actions, [Gao *et al.*, 2017; Hendricks *et al.*, 2017; Chen *et al.*, 2019; Zhang *et al.*, 2019] focus on temporal localization of actions by natural language queries. These methods can localize complex actions according to sentence queries, but are still difficult to recognize unseen activities. Thus, we consider localizing unseen activities in untrimmed videos via an image query.

3 Preliminary

In this section, we introduce the scaled dot-product attention and multi-head attention [Vaswani *et al.*, 2017].

We assume the input of the scale dot-product attention is a sequence of queries $\mathbf{Q} \in \mathbb{R}^{d_1 \times n_1}$, keys $\mathbf{K} \in \mathbb{R}^{d_1 \times n_2}$ and values $\mathbf{V} \in \mathbb{R}^{d_2 \times n_2}$, where n_1 , n_2 and n_2 represent the number of queries, keys and values, and d_1 , d_1 and d_2 are respective dimensions, the scaled dot-product attention is then calculated by

$$\text{DotAtten}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}(\mathbf{Q}^\top \mathbf{K} / \sqrt{d_1}) \mathbf{V}^\top,$$

where the Softmax operation is performed for every row.

The multi-head attention consists of H paralleled units. Each unit first projects the queries, keys and values to d_1 , d_1 and d_2 dimenstions, then independently perform the scale dot-product attention, and finally concatenate their outputs to another linear projection, given by

$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{W}^o \text{Concat}(\text{head}_1, \dots, \text{head}_H)$$

$$\text{where } \text{head}_i = \text{DotAtten}(\mathbf{W}_i^q \mathbf{Q}, \mathbf{W}_i^k \mathbf{K}, \mathbf{W}_i^v \mathbf{V}),$$

where $\mathbf{W}_i^q \in \mathbb{R}^{d_1 \times d_q}$, $\mathbf{W}_i^k \in \mathbb{R}^{d_1 \times d_k}$, $\mathbf{W}_i^v \in \mathbb{R}^{d_2 \times d_v}$, $\mathbf{W}^o \in \mathbb{R}^{d_v \times H d_2}$ are linear projection matrices. The d_q , d_k and d_v are the initial input dimensions and d_v is the output dimension. If the queries, keys and values are set to the same sequences, it is called multi-head self-attention.

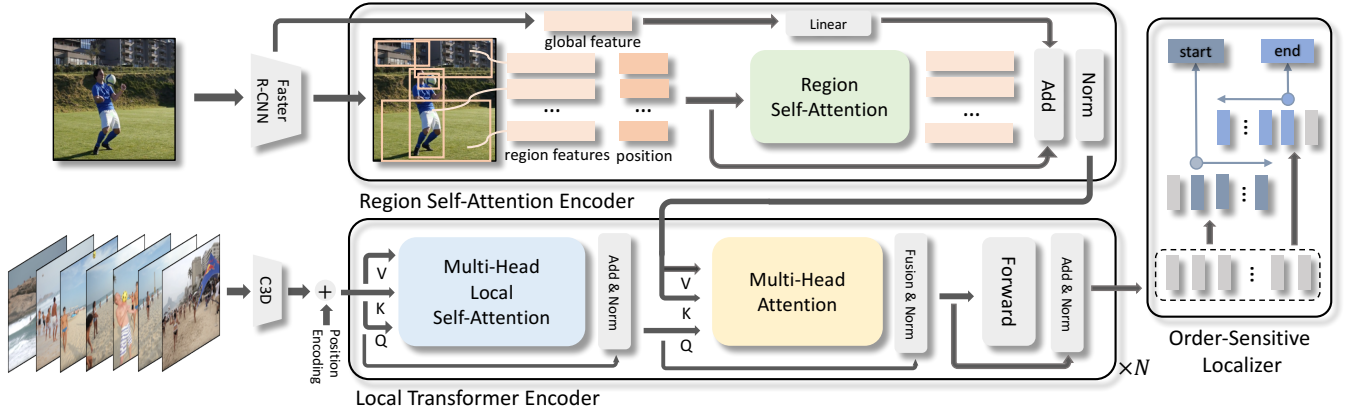


Figure 2: The Framework of Self-Attention Interaction Localizer for Image-Based Activity Localization.

4 Self-Attention Interaction Localizer

4.1 Problem Formulation

We denote a video by $\mathbf{v} = \{\mathbf{v}_i\}_{i=1}^n$, where \mathbf{v}_i is the feature of the i -th frame and n is the frame number. Each video is associated with an image query \mathbf{q} , which corresponds to a video segment depicting the similar semantic contents. And the start and end boundaries of this target segment are denoted by $\tau = (s, e)$. Thus, given a set of videos, image queries and segment annotations, the goal of the self-attention interaction localizer is to predict the boundaries $\hat{\tau} = (\hat{s}, \hat{e})$ of unseen activities during the inference process. The overall framework is shown in Figure 2.

4.2 Region Self-Attention Encoder

In this section, we introduce the region self-attention encoder with relative position encoding to learn fine-grained image region representations.

The image query often contains a series of objects in different regions, and the interactions between these objects imply crucial semantic information. Thus, we develop the region self-attention to explore the region relations for fine-grained understanding. Similar to the region proposal [Anderson *et al.*, 2018] for image caption, we first extract basic region features from the image query using a pre-trained Faster R-CNN [Ren *et al.*, 2015], denoted by $\mathbf{q} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m)$ of m regions. Each region \mathbf{r}_i contains an object and associates with a 4-dimensional position vector $\mathbf{p}_i = (x_i, y_i, w_i, h_i)$, where x_i and y_i are center coordinates, and w_i and h_i represent width and height of this region. Moreover, we pool the overall feature maps of the last shared convolutional layer of Faster R-CNN to obtain the global image feature \mathbf{f}^g .

We then develop the region self-attention method to establish the potential relations between objects. The original self-attention approach [Vaswani *et al.*, 2017] is based on scaled dot-product attention. Setting \mathbf{Q} , \mathbf{K} and \mathbf{V} to the same region matrix $\mathbf{R} = [\mathbf{r}_1; \mathbf{r}_2; \dots; \mathbf{r}_m] \in \mathbb{R}^{d_r \times m}$, the $\text{DotAtten}(\mathbf{Q}, \mathbf{K}, \mathbf{V})$ becomes a self-attention operation, where d_r is the dimension of region features. Furthermore, we consider the relative position of each region pair to build more informative interactions, shown in Figure 3(a). Specifi-

cally, given the position vector \mathbf{p}_i and \mathbf{p}_j of the i -th and j -th regions, the relative position vector $\mathbf{p}_{ij}^r = (x_{ij}, y_{ij}, w_{ij}, h_{ij})$ is computed by

$$\begin{aligned} x_{ij} &= (x_i - x_j)/w_j, & y_{ij} &= (y_i - y_j)/h_j, \\ w_{ij} &= \log(w_i/w_j), & h_{ij} &= \log(h_i/h_j). \end{aligned}$$

We then project the 4-dimensional vector to d_r dimension, denoted by $\tilde{\mathbf{p}}_{ij}^r = \mathbf{W}^r \mathbf{p}_{ij}^r$, where $\mathbf{W}^r \in \mathbb{R}^{d_r \times 4}$ represents the projection weight. Thus, for each region \mathbf{r}_i , we can obtain a relative position matrix $\tilde{\mathbf{P}}_i = [\tilde{\mathbf{p}}_{i1}^r; \tilde{\mathbf{p}}_{i2}^r; \dots; \tilde{\mathbf{p}}_{im}^r] \in \mathbb{R}^{d_r \times m}$, and the region self-attention result for the i -th region is given by

$$\tilde{\mathbf{h}}_i^r = \text{Softmax}(\mathbf{r}_i^\top (\mathbf{R} + \tilde{\mathbf{P}}_i) / \sqrt{d_r}) (\mathbf{R} + \tilde{\mathbf{P}}_i)^\top,$$

where the region representation $\tilde{\mathbf{h}}_i^r$ incorporates relevant information from other regions with relative position encoding. Next, we employ a residual connection [He *et al.*, 2016] to keep basic region information, and add the global feature to each region by a linear projection, followed by a layer normalization [Ba *et al.*, 2016]:

$$\mathbf{h}_i^r = \text{LayerNorm}(\tilde{\mathbf{h}}_i^r + \mathbf{r}_i + \mathbf{W}_g \mathbf{f}^g).$$

Hence, by the region self-attention encoder, we finally obtain the fine-grained region representations $\mathbf{H}^r \in \mathbb{R}^{d_r \times m}$, consisting of $[\mathbf{h}_1^r; \mathbf{h}_2^r; \dots; \mathbf{h}_m^r]$.

4.3 Local Transformer Encoder

In this section, we propose a local transformer encoder to build multi-step fusion and reasoning of image and video contents, where multi-level interactions can sufficiently explore the query semantics for fuzzy localization.

Following the original transformer framework [Vaswani *et al.*, 2017], we first add a temporal information encoding to the frame sequences $\{\mathbf{v}_i\}_{i=1}^n$, then develop the K -layer local transformer encoder, where the output of $k-1$ -th layer is regarded as the input of the k -th layer. Each layer consists of a multi-head local self-attention sublayer, a multi-head attention sublayer and a feed-forward sublayer. And the residual connection is applied from input to output of each sublayer, followed by layer normalization. Next, we detailedly

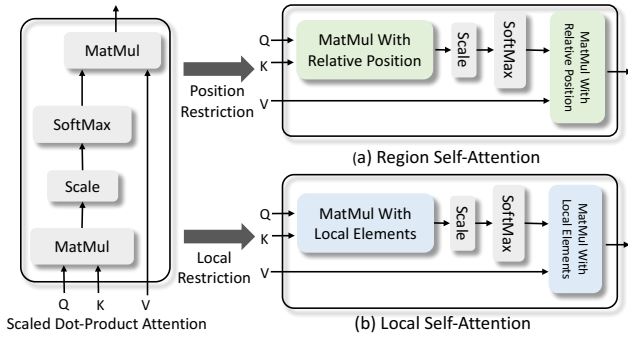


Figure 3: The variants of scaled dot-product attention. (a) The region self-attention contains relative position encoding. (b) The local self-attention only focuses on its surrounding elements.

illustrate the 1-th layer of our encoder, which takes the initial frame sequence $\{\mathbf{v}_i\}_{i=1}^n$ as input.

We first introduce the multi-head local self-attention sublayer. Given the feature matrix $\mathbf{F} = [\mathbf{v}_1; \mathbf{v}_2; \dots; \mathbf{v}_n] \in \mathbb{R}^{d_f \times n}$ with n elements of d_f dimension, the original multi-head self-attention $\text{MultiHead}(\mathbf{F}, \mathbf{F}, \mathbf{F})$ captures global frame-to-frame dependencies. However, retrieving the target segment for image queries mainly depends on local relation modeling. Thus, we augment the local restrictions for the underlying scaled dot-product attention, shown as Figure 3(b). For the i -th frame, the local self-attention only focuses on its surrounding elements, given by

$$\tilde{\mathbf{v}}_i = \text{Softmax}(\mathbf{v}_i^\top \mathbf{F}_{:,i-w:i+w} / \sqrt{d_f}) \mathbf{F}_{:,i-w:i+w}^\top,$$

where the w is the window size of local self-attention and $\mathbf{F}_{:,i-w:i+w}$ represents consecutive $2w + 1$ frames around the i -th element. This local strategy is similar to the convolutional neural networks, but CNNs share the convolution kernel at each time step, while our local self-attention relies on the current element. That is, the CNNs are location-based modeling and our local self-attention belongs to content-based modeling. Furthermore, we apply the multi-head setting for local self-attention just like the multi-head dot-product attention, called LocalMultiHead , and obtain the output of the 1-th sublayer by $\mathbf{F}^{11} = \text{LayerNorm}(\text{LocalMultiHead}(\mathbf{F}, \mathbf{F}, \mathbf{F}) + \mathbf{F})$.

We then propose a multi-head attention sublayer over the image region representations \mathbf{H}^r . Specifically, given $\mathbf{H}^r = [\mathbf{h}_1^r; \mathbf{h}_2^r; \dots; \mathbf{h}_m^r] \in \mathbb{R}^{d_r \times m}$ and $\mathbf{F}^{11} \in \mathbb{R}^{d_f \times n}$, we aggregate the frame-aware image representations by $\tilde{\mathbf{F}}^{11} = \text{MultiHead}(\mathbf{F}^{11}, \mathbf{H}^r, \mathbf{H}^r)$, where $\tilde{\mathbf{F}}^{11}$ incorporates the crucial query information for each frame and eliminates the influence of semantically inessential query contents. After that, we apply a more informative fusion [Mou *et al.*, 2015] instead of the additive residual connection, given by

$$\text{Fusion}(\mathbf{F}_{:,i}^{11}, \tilde{\mathbf{F}}_{:,i}^{11}) = \tanh(\mathbf{W}_f \bar{\mathbf{F}}_{:,i}^{11} + \mathbf{b}_f),$$

$$\text{where } \bar{\mathbf{F}}_{:,i}^{11} = [\mathbf{F}_{:,i}^{11}; \tilde{\mathbf{F}}_{:,i}^{11}; \mathbf{F}_{:,i}^{11} \odot \tilde{\mathbf{F}}_{:,i}^{11}; \mathbf{F}_{:,i}^{11} - \tilde{\mathbf{F}}_{:,i}^{11}],$$

where \odot represents the element-wise product and $[\cdot; \cdot]$ is the feature concatenation. Thus, we get the output of the 2-th sublayer by $\mathbf{F}^{12} = \text{LayerNorm}(\text{Fusion}(\mathbf{F}^{11}, \tilde{\mathbf{F}}^{11}))$.

In addition to the two attention sublayers, each layer contains a feed-forward sublayer, composed of two linear transformations with ReLU activation [Nair and Hinton, 2010]:

$$\tilde{\mathbf{F}}^{12} = \mathbf{W}_{13}^2 \text{ReLU}(\mathbf{W}_{13}^1 \mathbf{F}^{12} + \mathbf{b}_{13}^1 \mathbf{1}^\top) + \mathbf{b}_{13}^2 \mathbf{1}^\top.$$

At the end, we obtain the output of the 1-th layer by $\mathbf{F}^{13} = \text{LayerNorm}(\mathbf{F}^{12} + \tilde{\mathbf{F}}^{12})$.

By stacking N layers, our local transformer encoder builds multi-step fusion and reasoning of image and video contents, and obtain the final semantic representations \mathbf{F}^{N3} , also denoted by $\mathbf{H}^v = [\mathbf{h}_1^v; \mathbf{h}_2^v; \dots; \mathbf{h}_n^v]$.

4.4 Order-Sensitive Localizer

In this section, we apply an order-sensitive localizer to directly retrieve the start and end boundaries of the target segment, which leverages the bi-directional contextual information to boost the localization accuracy.

While predicting the start and end boundaries, the information from different directions has different influences [Chen *et al.*, 2019]. For example, we suppose the i -th frame is the start frame of the target activity, the forward information after this frame can reflect whether a complete activity just starts at the i -th frame, and vice versa. Based on this observation, we propose a bi-directional aggregation method. Specifically, we first aggregate forward semantic representations for the i -th frame, given by

$$\mathbf{h}_i^{fw} = \sum_{t=i}^n \alpha_t^{fw} \mathbf{h}_t^v, \quad \alpha_{it}^{fw} = \frac{\exp(a_{it}^{fw})}{\sum_{t=i}^n \exp(a_{it}^{fw})},$$

$$a_{it}^{fw} = \mathbf{w}_{fw}^T \tanh(\mathbf{W}_{fw}^1 \mathbf{h}_i^v + \mathbf{W}_{fw}^2 \mathbf{h}_t^v),$$

where \mathbf{h}_i^{fw} is the forward context representation of the i -th frame and α_{it}^{fw} is the attention weight from the i -th frame to the t -th frame, where $t \geq i$. Instead of preceding dot-product attention, we apply the additive attention here, which is slower but more effective. Likewise, the backward context representation \mathbf{h}_i^{bw} is generated by aggregating those frame representations before the index i . Given the bi-directional context representation $\mathbf{H}^{fw} = [\mathbf{h}_1^{fw}; \mathbf{h}_2^{fw}; \dots; \mathbf{h}_n^{fw}]$ and $\mathbf{H}^{bw} = [\mathbf{h}_1^{bw}; \mathbf{h}_2^{bw}; \dots; \mathbf{h}_n^{bw}]$, we can directly predict the start and end boundaries as follows

$$\mathbf{p}^s = \text{Softmax}(\mathbf{w}_s^T \tanh(\mathbf{W}_s \mathbf{H}^{fw} + \mathbf{b}_s \mathbf{1}^\top)),$$

$$\mathbf{p}^e = \text{Softmax}(\mathbf{w}_e^T \tanh(\mathbf{W}_e \mathbf{H}^{bw} + \mathbf{b}_e \mathbf{1}^\top)),$$

$$\hat{s} = \arg \max_i (\mathbf{p}_i^s), \quad \hat{e} = \arg \max_i (\mathbf{p}_i^e).$$

where \mathbf{p}^s and \mathbf{p}^e are the probability distributions of start and end boundaries over frame sequences, and the $\hat{\tau} = (\hat{s}, \hat{e})$ is the retrieval result. Finally, we train our self-attention interaction localizer by minimizing the negative log probabilities of the ground truth boundaries by

$$\mathcal{L} = -\frac{1}{N_b} \sum_1^{N_b} (\log \mathbf{p}_s^s + \log \mathbf{p}_e^e),$$

where \mathbf{p}_s^s is the probability of ground truth start, \mathbf{p}_e^e is the end probability and N_b is the sample numbers of a batch.

	Sample	Class	Video Time	Target Time
Train	30,677	160	63.17	22.37
Valid	2,995	20	73.38	25.21
Test	3,013	20	61.20	21.61
All	36,685	200	63.84	22.54

Table 1: Summaries of the ActivityIBAL Dataset, including sample number, class number, average video duration and average segment duration.

5 Dataset

Since none of existing datasets can be directly used for image-based activity localization, we construct a new dataset ActivityIBAL by reorganizing the ActivityNet [Caba Heilbron *et al.*, 2015]. The ActivityNet is a large-scale video dataset for action localization, which provides 14950 public videos with temporal annotations of 200 classes. Videos in ActivityNet contain one or multiple action segments, and these segments in one video often have the same action label.

Following similar reconstruction [Feng *et al.*, 2018], we preprocess these videos as follows. We first merge every two overlapped segments belonging to the same labels in videos, and then split each original video with multiple action segments into several independent videos. Next, we discard some videos with too short or long action segments, and divide these newly generated videos into three subsets according to their action classes: 160 classes for training, 20 classes for validation and the remaining 20 classes for testing. This setting guarantees these videos used in validation and testing do not appear during the training process.

We then assume that the segments belonging to the same class describe the semantically related activities. The ActivityIBAL dataset contains two parts of data: we first construct the **Simple** part, where an untrimmed video corresponds to an image query selected from another video with the same class; then select the image query from the similar but not same class to compose the **Difficult** part. Here we manually select the most informative frame as the image query. The similar class can be found in the tree structure of these classes in ActivityNet. Note that, to keep the diversity of image queries, we construct 5 samples for every video, corresponding to 3 simple queries and 2 difficult queries. The details are summarized in Table 1.

6 Experiments

6.1 Implementation Details

For image queries, we first extract the 4,096-d region feature and 4-d position vector for each region using a pre-trained Faster R-CNN [Ren *et al.*, 2015] with VGG-16 backbone networks, and then apply the non-maximum suppression with threshold 0.7 to reduce redundant regions. By the way, we obtain the 512-d global feature. As for video sequences, we first resize each frame to 112×112 , and then extract 4,096-d spatio-temporal features by a pre-trained 3D-ConvNet [Tran *et al.*, 2015] for each unit. A unit contains consecutive 16 frames and overlaps 8 frames with adjacent units. Next, we reduce the dimension from 4,096 to 500 using PCA. To avoid

Method	mIoU	IoU@0.3	IoU@0.5	IoU@0.7
Random	14.68%	18.13%	10.28%	6.87%
FLP	23.73%	29.81%	19.49%	12.37%
SST+	40.63%	59.92%	32.17%	14.50%
TAL+	41.93%	64.55%	36.50%	17.86%
SAIL	49.66%	73.98%	48.25%	24.37%

Table 2: Evaluation Results on the whole ActivityIBAL Dataset.

Method	mIoU	IoU@0.3	IoU@0.5	IoU@0.7
Random	14.32%	18.22%	10.19%	6.77%
FLP	17.89%	22.30%	14.85%	9.42%
SST+	32.37%	48.06%	26.79%	11.26%
TAL+	35.20%	56.44%	32.05%	13.39%
SAIL	45.71%	68.23%	45.17%	21.84%

Table 3: Evaluation Results on the Difficult Part of ActivityIBAL.

exceeding GPU load, we downsample overlong features sequences to 200. In the proposed model, we set the projection dimension of various kinds of attention to 256. During the training process, we adopt an Adam optimizer and the initial learning rate is set to 0.0005.

6.2 Baselines and Evaluation Criteria

Since these are no other approaches designed for image-based activity localization, we compare our SAIL method with three baselines, one designed by ourselves and two extended from relevant tasks.

- **FLP**: We design a simple baseline to perform the frame-level localization. For each frame, given frame feature \mathbf{v}_i and global image feature \mathbf{f}^g , we compute the confidence score by multilayer perceptron $p_i = \text{MLP}(\mathbf{v}_i, \mathbf{f}^g)$. The score of each segment (i, j) is given by $p_{ij} = \prod_{k=1}^{i-1} (1 - p_k) \prod_{k=i}^j (p_k) \prod_{k=j+1}^n (1 - p_k)$.
- **SST+** [Buch *et al.*, 2017]: The baseline is the extension of an action proposal method (SST). We add a fusion layer on the top of original GRU networks to incorporate the global image feature, and then choose the proposal with the largest confident score.
- **TAL+** [Chao *et al.*, 2018]: Based on an action localization method, we first generate some action proposals, then add an attention layer to aggregate the image region features for each proposal and finally estimate the proposal score.

Evaluation Criteria: We evaluate the performance of all methods on the criteria **mIoU** and **IoU@R**. Specifically, we first calculate the Intersection over Union (IoU) for each pair of predicted segments and ground-truth ones, and then define the **mIoU** is a mean IoU of all testing samples. And **IoU@R** represents the percentage of the samples, which $\text{IoU} > R$.

6.3 Performance Evaluation

We then evaluate all models on the ActivityIBAL dataset. The overall results are shown in Table 2, where the random results are yielded by selecting a segment randomly. And the Table 3




Figure 4: Effect of layer number of the local transformer encoder.

shows the performance only on the difficult testing samples of ActivityIBAL, where the training process is still performed on the whole training set. The experiment results reveal some interesting points:

- The frame-level FLP method achieves the worst performance on all criteria, and is only a little better than the random strategy. It demonstrates that the image-based activity localization cannot only depend on the frame-level modeling.
- The attention based method SAIL and TAL+ outperform other baselines, which suggests that the frame-aware aggregation of image region features is critical to localize unseen activities.
- In all the criteria, our SAIL method achieves the best performance. This fact shows that our proposed region self-attention encoder, local transformer encoder and order-sensitive localizer can effectively improve the localization accuracy of unseen activities.
- Although all methods have lower accuracy on the Difficult samples than the whole testing set, our SAIL still outperforms other baselines and has less performance degradation, which demonstrates our strategy has the stronger ability of fuzzy localization.

In our SAIL method, the stacked layer number of the local transformer encoder is an essential hyper-parameter. We next explore the impact of this hyper-parameter. Concretely, we vary the layer number from 1 to 7 and display the results in Figure 4. From these figures, we note that too many or too few layers will both affect the performance, and our SAIL method achieves the best performance when the layer number is set to 5. Moreover, we display a typical example of image-based activity localization in Figure 5, where we localize an unseen activity “grass hockey” by an image “ice hockey”.

Method	mIoU	IoU@0.3	IoU@0.5	IoU@0.7
w/o. RS	46.63%	70.35%	43.36%	20.43%
w/o. ML	48.40%	73.29%	44.07%	21.24%
w/o. LS	47.35%	71.17%	45.71%	22.29%
w/o. BA	48.04%	72.00%	45.76%	23.60%
full	49.66%	73.98%	48.25%	24.37%

Table 4: Evaluation Results of Ablation Study.



Figure 5: A typical example of image-based activity localization.

6.4 Ablation Study

We next verify the contribution of each component by the ablation study. We first remove the region self-attention from the image encoder, denoted by **SAIL(w/o. RS)**. To validate the effect of multi-level interactions, **SAIL(w/o. ML)** discards the multi-head attention sublayer from the first $N - 1$ layers of the local encoder. We next replace the multi-head local self-attention with the normal multi-head self-attention, called as **SAIL(w/o. LS)**. Finally, we develop a model **SAIL(w/o. BA)** without the bi-directional aggregation method in the order-sensitive localizer.

As shown in Table 4, the SAIL(full) outperforms all ablation models, demonstrating the region self-attention, multi-level interactions, local self-attention and bi-directional aggregation method are all effective to boost the localization performance. Specifically, SAIL(w/o. RS) achieves the worst performance, indicating the fine-grained region representation learning is the basis of accurate localization. Moreover, SAIL(w/o. ML) has a more severe performance degradation on IoU@0.5 and IoU@0.7 than mIoU and IoU@0.3, demonstrating the multi-level interactions are more effective for accurate localization. On the contrary, we note that the local self-attention and bi-directional aggregation method has a greater impact on the rough localization.

7 Conclusion

In this paper, we introduce a new task to localize unseen activities in videos via image queries. To deal with the inherent challenges of this task, we propose a novel self-attention interaction localizer. Furthermore, we reorganize the ActivityNet to construct a new dataset ActivityIBAL. The extensive experiments show the effectiveness of our method.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant No.61602405, No.U1611461, No.61751209 and No.61836002, Joint Research Program of ZJU and Hikvision Research Institute. This work is also supported by Alibaba Innovative Research and Microsoft Research Asia.

References

- [Anderson *et al.*, 2018] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- [Ba *et al.*, 2016] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [Buch *et al.*, 2017] Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Niebles. Sst: Single-stream temporal action proposals. In *CVPR*, pages 6373–6382. IEEE, 2017.
- [Caba Heilbron *et al.*, 2015] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015.
- [Chao *et al.*, 2018] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018.
- [Chen *et al.*, 2019] Jingyuan Chen, Lin Ma, Xinpeng Chen, Zequn Jie, and Jiebo Luo. Localizing natural language in videos. In *AAAI*, 2019.
- [Feng *et al.*, 2018] Yang Feng, Lin Ma, Wei Liu, Tong Zhang, and Jiebo Luo. Video re-localization. In *ECCV*, pages 51–66, 2018.
- [Gao *et al.*, 2017] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. TALL: temporal activity localization via language query. In *ICCV*, pages 5277–5285. IEEE, 2017.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [Hendricks *et al.*, 2017] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, pages 5803–5812, 2017.
- [Lin *et al.*, 2017] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In *MM*, pages 988–996. ACM, 2017.
- [Mou *et al.*, 2015] Lili Mou, Rui Men, Ge Li, Yan Xu, Lu Zhang, Rui Yan, and Zhi Jin. Natural language inference by tree-based convolution and heuristic matching. *arXiv preprint arXiv:1512.08422*, 2015.
- [Nair and Hinton, 2010] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, pages 807–814, 2010.
- [Nguyen *et al.*, 2018] Phuc Nguyen, Ting Liu, Gautam Prasad, and Bohyung Han. Weakly supervised action localization by sparse temporal pooling network. In *CVPR*, pages 6752–6761, 2018.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015.
- [Shou *et al.*, 2016] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, pages 1049–1058, 2016.
- [Shou *et al.*, 2017] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-de-convolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, pages 1417–1426. IEEE, 2017.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, pages 568–576, 2014.
- [Singh *et al.*, 2016] Bharat Singh, Tim K Marks, Michael Jones, Oncel Tuzel, and Ming Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *CVPR*, pages 1961–1970, 2016.
- [Soomro and Shah, 2017] Khurram Soomro and Mubarak Shah. Unsupervised action discovery and localization in videos. In *CVPR*, pages 696–705, 2017.
- [Tran *et al.*, 2015] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, pages 5998–6008, 2017.
- [Wang *et al.*, 2017] Limin Wang, Yuanjun Xiong, Dahua Lin, and Luc Van Gool. Untrimmednets for weakly supervised action recognition and detection. In *CVPR*, 2017.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018.
- [Xu *et al.*, 2017] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: region convolutional 3d network for temporal activity detection. In *ICCV*, pages 5794–5803, 2017.
- [Yeung *et al.*, 2016] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, pages 2678–2687, 2016.
- [Zhang *et al.*, 2019] Zhu Zhang, Zhijie Lin, Zhou Zhao, and Zhenxin Xiao. Cross-modal interaction networks for query-based moment retrieval in videos. In *SIGIR*. ACM, 2019.
- [Zhao *et al.*, 2017] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, 2017.