

Fine-grained Semantic Alignment Network for Weakly Supervised Temporal Language Grounding

Yuechen Wang Wengang Zhou Houqiang Li*

University of Science and Technology of China

{wyc9725, zhwg, lihq}@ustc.edu.cn

Abstract

Temporal language grounding (TLG) aims to localize a video segment in an untrimmed video based on a natural language description. To alleviate the expensive cost of manual annotations for temporal boundary labels, we are dedicated to the weakly supervised setting, where only video-level descriptions are provided for training. Most of the existing weakly supervised methods generate a candidate segment set and learn cross-modal alignment through a MIL-based framework. However, the temporal structure of the video as well as the complicated semantics in the sentence are lost during the learning. In this work, we propose a novel *candidate-free* framework: Fine-grained Semantic Alignment Network (FSAN), for weakly supervised TLG. Instead of view the sentence and candidate moments as a whole, FSAN learns token-by-clip cross-modal semantic alignment by an iterative cross-modal interaction module, generates a fine-grained cross-modal semantic alignment map, and performs grounding directly on top of the map. Extensive experiments are conducted on two widely-used benchmarks: ActivityNet-Captions, and DiDeMo, where our FSAN achieves state-of-the-art performance.

1 Introduction

Given an untrimmed video and a natural language sentence, Temporal Language Grounding (TLG) aims to localize the temporal boundaries of the video segment described by a referred sentence. TLG is a challenging problem with great importance in various multimedia applications, *e.g.*, video retrieval (Shao et al., 2018), visual question answering (Tapaswi et al., 2016; Antol et al., 2015; Yu et al., 2020), and visual reasoning (Yang et al., 2018). Since its first proposal (Gao et al., 2017; Hendricks et al., 2017), tremendous success has

*Corresponding author.

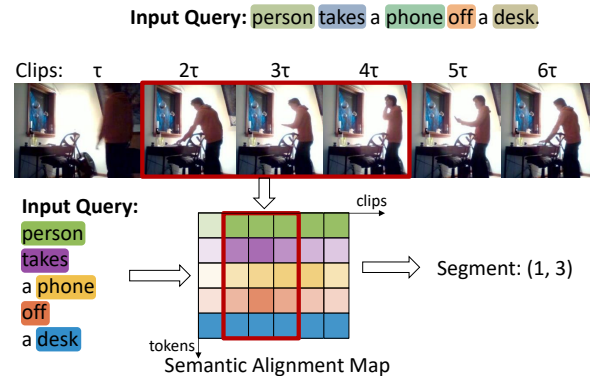


Figure 1: Illustration of fine-grained semantic alignment map for temporal language grounding.

been made on this problem (Wu and Han, 2018; Chen et al., 2018; Ge et al., 2019; Yuan et al., 2018; Zhang et al., 2019a; He et al., 2019; Wang et al., 2019; Zhang et al., 2020b; Ning et al., 2021). Despite the achievements with supervised learning, the temporal boundaries for every sentence query need to be manually annotated for training, which is expensive, time-consuming, and potentially noisy. On the other hand, it is much easier to collect a large amount of video-level descriptions without detailed temporal annotations, since video-level descriptions naturally appear with videos simultaneously on the Internet (*e.g.*, YouTube). To this end, some prior works are dedicated to weakly supervised setting, where only video-level descriptions are provided, without temporal labels.

Most of the previous weakly supervised methods follow a Multiple Instance Learning (MIL) paradigm, which samples matched and non-matched video-sentence pairs, and learn a matching classifier to implicitly learn the cross-modal alignment. However, during the matching classification, the input sentence is often treated as a single feature query, neglecting the complicated linguistic semantics. VLANet (Ma et al., 2020) treats tokens in the input sentence separately, and performs cross-modal attention on token-moment

pairs, where the moment candidates are carefully selected by a surrogate proposal selection module to reduce computation cost. SCN (Lin et al., 2020) proposes to generate and select moment candidates and performs semantic completion for the sentence to rank selected candidates. Nevertheless, the generation and selection process of moment candidates also involves high computational costs. In addition, the moment candidates are considered separately, while the temporal structure of the video is also important for grounding. Figure 1 shows an example to localize the query “*person takes a phone off a desk*” in the given video. If the model views the sentence as a whole and performs matching classification, it is hard to learn undistinguished words like “*off*” during the training. However, the neglected words may play important roles to determine the temporal boundaries of the described moment.

In this paper, we propose a novel framework named a Fine-grained Semantic Alignment Network (FSAN), for weakly supervised temporal language grounding. The core idea of FSAN is to learn token-by-clip cross-modal semantic alignment presenting as a token-clip map, and ground the sentence on video directly based on it. Specifically, given an untrimmed video and a description sentence, we first extract their features by visual encoder and textual encoder independently. Then, an Iterative Cross-modal Interaction Module is devised to learn the correspondence between visual and linguistic representations. To make temporal predictions for grounding, we further devise a semantic alignment-based grounding module. Based on the learned cross-modal interacted features, a token-by-clip semantic alignment map is generated, where the (i, j) -th element on the map indicates relevance between the i -th token in the sentence and j -clip in the video. Finally, an alignment-based grounding module predicts the grounding result corresponding to the input sentence.

Instead of aggregating sentence semantics into one representation and generating video moment candidates, FSAN learns a fine-grained cross-modal alignment map that helps to retain both the temporal structure among video clips and the complicated semantics in the sentence. Furthermore, the grounding module in FSAN makes predictions mainly based on the cross-modal alignment map, which alleviates the computation cost of candidate moment representation generation. We demonstrate the effectiveness of the proposed method

on two widely-used benchmarks: ActivityNet-Captions (Krishna et al., 2017) and DiDeMo (Hendricks et al., 2017), where state-of-the-art performance is achieved by FSAN.

2 Related Work

2.1 Temporal Language Grounding

Temporal language grounding is proposed (Gao et al., 2017; Hendricks et al., 2017) as a new challenging task, which requires deep interactions between two visual and linguistic modalities. Previous methods have explored this task in a fully supervised setting (Gao et al., 2017; Hendricks et al., 2017; Chen et al., 2018; Ge et al., 2019; Xu et al., 2019; Chen and Jiang, 2019; Yuan et al., 2018; Zhang et al., 2019b,a; Lu et al., 2019). Most of them follow a two-stage paradigm: generating candidate moments with sliding windows and subsequently matching the language query. Reinforcement learning has also been leveraged for temporal language grounding (He et al., 2019; Wang et al., 2019; Cao et al., 2020).

Despite the boom of fully supervised methods, it is very time-consuming and labor-intensive to annotate temporal boundaries for a large number of videos. And due to the annotation inconsistency among annotators, temporal labels are often ambiguous for models to learn. To alleviate the cost of fine-grained annotation, weakly supervised setting is explored lately (Mithun et al., 2019; Gao et al., 2019; Lin et al., 2020; Ma et al., 2020; Zhang et al., 2020c). TGA (Mithun et al., 2019) exploits maps video candidate features and query features into a latent space to learn cross-modal similarity. In (Ma et al., 2020), a video-language attention network is proposed to learn cross-modal alignment between language tokens and video segment candidates. Differently, our FSAN gets rid of the trouble of generating candidates and learns fine-grained token-by-clip semantic alignment.

2.2 Transformer in Language and Vision

Since it is first proposed by Vaswani *et al.* (Vaswani et al., 2017) for machine translation, transformer has become a prevailing architecture in NLP. The basic block of transformer is the multi-head attention module, which aggregates information from the whole input in both transformer encoder and decoder module. Transformer demonstrates superior performance in language model pretraining methods (Devlin et al., 2019; Radford et al.,




Figure 2: The architecture of FSN. It consists of four main components: (1) a text encoder, (2) a visual encoder, (3) an iterative cross-modal interaction module, and (4) a proposal module.

2018; Yang et al., 2019), and achieves competitive performance on diverse NLP problems. Recently, transformer has been introduced to various computer vision tasks, such as image classification (Chen et al., 2020b), image generation (Parmar et al., 2018), object detection (Carion et al., 2020), semantic segmentation (Wang et al., 2021a), tracking (Wang et al., 2021b), *etc.* Comparing to CNN, the attention mechanism learns more global dependencies, therefore, transformer also shows great performance in low-level tasks (Chen et al., 2020a). Transformer has also been proved effective in multi-modal area, including multi-modal representations (Zhang et al., 2020a; Tan and Bansal, 2019; Su et al., 2020; Sun et al., 2019) and applications (Shi et al., 2020; Ju et al., 2020; Liang et al., 2020). Inspired by the great success, we devise an iterative cross-modal interaction module mainly based on the multi-head attention mechanism.

3 Our Approach

Given an untrimmed video and a text-sentence query, a temporal grounding model aims to localize the most relevant moment in the video, represented by its beginning and ending timestamps. In this paper, we consider the weakly supervised setting, *i.e.*, for each video V , a textual query S is provided for training. The query sentence describes a specific moment in the video, yet the temporal boundaries are not provided for training. In the inference stage, the weakly trained model is required to predict the beginning and ending timestamps of the video moment that corresponds to the input sentence S .

We present a novel framework named Fine-grained Semantic Alignment Network (FSAN) for the temporal language grounding problem. As

shown in Figure 2, given a video and text query, we first encode them separately. The resulting representations then interact with each other through an iterative cross-modal interaction module. The outputs are used to learn a Semantic Alignment Map (SAP) between the two modalities. Finally, the SAP is fed into an alignment-based grounding module to predict scores for all possible moments.

In the following subsections, we will first introduce the visual and language encoder, then describe the Iterative Cross-Modal Interaction Module. Finally, we will elaborate on the semantic alignment map and the grounding module based on it.

3.1 Input Representation

Language Encoder. We use a standard transformer encoder (Vaswani et al., 2017) to extract the semantic information for the input query sentence S . Each token in the input query is first embedded using GloVe (Pennington et al., 2014). The resulting vectors are mapped to dimension of d_s by a linear layer and fed into a transformer encoder to obtain context-aware token features $\mathbf{S} = \{\mathbf{w}_i\}_{i=1}^{N_s}$, where N_s is the number of tokens and $\mathbf{w}_k \in \mathbb{R}^{d_s}$ denotes the feature of k -th token in the sentence.

Video Encoder. For the input videos, we extract visual features using a pretrained feature extractor and then apply a temporal pooling on frame features to divide it into N_v clips. Hence the video can be represented by $\mathbf{V} = \{\mathbf{v}_j\}_{j=1}^{N_v}$, where $\mathbf{v}_j \in \mathbb{R}^{d_v}$ denotes the feature of j -th video clip, and $d_v = d_s$ is the dimension of visual feature. Experimental results illustrate that the computation cost is considerably reduced by the temporal pooling.

3.2 Iterative Cross-Modal Interaction Module

Inspired by the great success of transformer encoder on vision-language pretraining (Li et al., 2020; Tan and Bansal, 2019), we devise an Iterative Cross-modal Interaction Module (ICIM) to learn the semantic relevance between visual and textual representations. The module is composed of a stack of 6 layers, and each layer consists of cross-modal attention, inner-modal attention, and feed-forward layers. The core component for cross-modal interaction is the multi-head attention, which is also vital in the transformer structure. Formally, given two sequences of d -dimensional features $X = [\mathbf{x}_1, \dots, \mathbf{x}_{N_x}]$ and $Y = [\mathbf{y}_1, \dots, \mathbf{y}_{N_y}]$, the calculation of multi-head attention is as follows:

$$\begin{aligned} \mathbf{Q}_i &= \mathbf{W}_i^Q \mathbf{X}, \mathbf{K}_i = \mathbf{W}_i^K \mathbf{Y}, \mathbf{V}_i = \mathbf{W}_i^V \mathbf{Y}, \\ \mathbf{A}_i &= \text{softmax}\left(\frac{\mathbf{Q}_i \mathbf{K}_i^T}{\sqrt{d}}\right) \mathbf{V}_i, \end{aligned} \quad (1)$$

$$MA(\mathbf{X}, \mathbf{Y}) = \mathbf{W}^M (\mathbf{A}_1 || \mathbf{A}_2 || \dots || \mathbf{A}_m), \quad (2)$$

where \mathbf{A}_i is output the i -th of m attention heads. The final output $MA(X, Y)$ of multi-head attention is of same dimension as the input X .

For the textual representation $\mathbf{S} \in \mathbb{R}^{N_x \times d^S}$ and visual representation $\mathbf{V} \in \mathbb{R}^{N_x \times d^V}$ input to the iterative cross-modal interaction module, we first adopt cross-modal attention, *i.e.*,

$$\begin{aligned} \mathbf{S}' &= LN(MA(\mathbf{S}, \mathbf{V}) + \mathbf{S}), \\ \mathbf{S}'' &= LN(FFN(\mathbf{S}') + \mathbf{S}'), \end{aligned} \quad (3)$$

$$\begin{aligned} \mathbf{V}' &= LN(MA(\mathbf{V}, \mathbf{S}) + \mathbf{V}), \\ \mathbf{V}'' &= LN(FFN(\mathbf{V}') + \mathbf{V}'), \end{aligned} \quad (4)$$

where LN denotes layer normalization and FFN denotes feed-forward layer. To retain the temporal structure of the video and the grammar of the sentence, we add learnable positional encodings to the input of each modality. Through the above attention operation, the features of different modalities are able to freely interact with the other modality to learn a fine-grained semantic alignment.

To model the inner-modal context after cross-modal interaction, we further apply an inner-modal attention, which is similar with the calculation in Equation (3) and (4), except that the multi-head attention is applied only on single-modal representation, *i.e.*, self-attention on single-modal features. After 6 iterations of cross-modal interaction and inner-modal modeling, the enhanced features \mathbf{S}''

and \mathbf{V}'' are fed into subsequent proposal module to predict a cross-modal semantic alignment map, and perform grounding based on it.

3.3 Semantic Alignment Map

After iterative cross-modal and inner-modal attention in ICIM, the correspondence can be fully explored between each pair of textual tokens and video clips. Therefore, a token-by-clip Semantic Alignment Map (SAP) \mathbf{P} with size $N_s \times N_v$ can be learned for temporal grounding. Formally,

$$\begin{aligned} \mathbf{P} &= SAP(\mathbf{V}, \mathbf{S}) \\ &= \text{Norm}(\mathbf{W}^s \mathbf{S}''^T \cdot \mathbf{W}^v \mathbf{V}''^T), \end{aligned} \quad (5)$$

where $\mathbf{W}^s \in \mathbb{R}^{d_t \times d_s}$ and $\mathbf{W}^v \in \mathbb{R}^{d_t \times d_v}$ are learnable parameters, \cdot denotes dot product, and $\text{Norm}(\cdot)$ denotes max-min normalization operation along the visual axis. Note that the normalization operation is important here to avoid a trivial distribution, *i.e.*, the semantic alignment map \mathbf{P} is invariant to different sentence and video inputs.

The value of the (i, j) -th element on the SAP \mathbf{P}_{ij} represents the relevance between the i -th textual token and the j -th video clip. However, the SAP can learn the cross-modal relationship only with supervision that indicates semantic alignment. To this end, we adopt a video-level matching loss, which is calculated as:

$$\mathcal{L}^t = \max(0, \delta - S(\mathbf{V}, \mathbf{S}) + S(\mathbf{V}, \mathbf{S}^-)), \quad (6)$$

where $S(\mathbf{V}, \mathbf{S})$ is the matching score function between video \mathbf{V} and sentence \mathbf{S} , which is defined as:

$$S(\mathbf{V}, \mathbf{S}) = \frac{\sum_{i \in [0, N_s]} \max_{j \in [0, N_v]} \mathbf{P}_{ij}}{N_s}, \quad (7)$$

where \mathbf{S}^- is a non-matching description sentence randomly sampled from the dataset.

Some same textual expressions may appear in both the positive description and the negative one, confusing the model in the matching classification procedure. To this end, we mask the repeated tokens on the semantic alignment map \mathbf{P}^- of the input video \mathbf{V} and the sampled sentence \mathbf{S}^- .

3.4 Alignment-Based Grounding Module

The elements on the semantic alignment map indicate relevance between video clips and textual tokens, which leads to the idea of a fine-grained alignment-based grounding module. The core idea

is that if a specific clip v_i is part of the described moment $V_{s,e}$, where $V_{s,e}$ denotes the video segment from the s -clip to e -th clip. The semantic of v_i tends to be highly relevant with all tokens in the description. While if v_i is out of the correct moment, at least one token in query is irrelevant to it. Therefore, we score all possible temporal segments formed with video clips by the relevance scores of clips both in and out the segment. Considering the clips in the segment as positive clips, and those not in it as negative clips. Then the relevance score of $V_{s,e}$ and the query is then defined as:

$$\begin{aligned} SC_{s,e} &= \text{Score}(V_{s,e}) \\ &= \frac{\sum_{i \in [0, N_s]} \sum_{j \in [s, e]} P_{ij}}{N_s(e-s+1)} \\ &\quad - \frac{\sum_{j \in \mathcal{L}_{[s, e]}} \min_{i \in [0, N_s]} P_{i,j}}{N_v - (e-s+1)}, \end{aligned} \quad (8)$$

where $\mathcal{L}_{[s, e]}$ denotes the aggregation of negative clips, *i.e.*, the complementary set of $V_{s,e}$. The higher the average scores among positive clips for each token, the more possible that $V_{[s, e]}$ is relevant to the query. While the lower the average response of negative clips, the less possible that $V_{[s, e]}$ is redundant. Through the two-fold filtering, the moment that is more relevant with all tokens in the query will be given a higher score, and therefore more likely to be proposed as the grounding result.

Although the contrastive loss in Equation (6) enables the model to learn the cross-modal semantic alignment, the temporal discrimination can not be learned under coarse video-level supervision, which is vital for grounding. To provide fine-level temporal supervision for the fine-grained cross-modal alignment, we further devise a novel two-fold loss on the semantic alignment map \mathbf{P} , including an inner-sample loss and an outer-sample loss. Specifically, the inner-sample loss aims to enhance the grounded moment on the fine-level alignment map \mathbf{P} . We promote the weakest response among clips in all possible segments, with a weight representing the confidence of the prediction:

$$\mathcal{L}_{s,e}^i = -\frac{\sum_{i \in [0, N_s]} \log(\min_{j \in [s, e]} P_{i,j})}{N_s}, \quad (9)$$

$$\mathcal{L}^i = \sum_{s \in [0, N_v]} \sum_{e \in [0, N_v]} SC_{s,e} \mathcal{L}_{s,e}^{inner}. \quad (10)$$

On the other hand, the outer-sample loss aims to suppress complementary part of the video by low-

Benchmark	Num. of Videos	Num. of Descriptions	Vocab Size
ActivityNet-Captions	14 950	51 567	15 406
DiDeMo	10 464	41 206	7 523
Charades-STA	6 670	16 128	1 289

Table 1: Statistics of TLG benchmarks.

ering its weakest response of each clip:

$$\mathcal{L}_{s,e}^o = -\frac{\sum_{j \in \mathcal{L}_{[s, e]}} \log(\min_{i \in [0, N_s]} P_{i,j})}{N_v - (e-s+1)}, \quad (11)$$

$$\mathcal{L}^o = \sum_{s \in [0, N_v]} \sum_{e \in [0, N_v]} SC_{s,e} \mathcal{L}_{s,e}^{outer}. \quad (12)$$

At the beginning of training, the model is uncertain about the grounding results, therefore, the weight $SC_{s,e}$ varies a little among moment options. As the training continues, the model can give positive moment higher scores easily, hence the loss weight will be larger for positive moments and smaller for negative ones. Therefore, the model will not deviate much from the correct solution.

3.5 Training and Inference

The overall training objective is an aggregation of aforementioned losses, given by:

$$\mathcal{L} = \lambda_1 \mathcal{L}^t + \lambda_2 \mathcal{L}^i + \lambda_3 \mathcal{L}^o, \quad (13)$$

where λ_* are hyper-parameters, and satisfy the condition $\lambda_1 + \lambda_2 + \lambda_3 = 1$.

During the inference, the moment with the highest score $SC_{s,e}$ is selected as the grounding result.

4 Experiments

4.1 Datasets and Metrics

ActivityNet-Captions. The ActivityNet-Captions dataset (Krishna et al., 2017) is developed based on ActivityNet dataset (Heilbron et al., 2015), which contains 20 000 untrimmed videos and corresponding language descriptions. The released ActivityNet-Captions dataset contains 37 421 moment-description pairs for training, and 17 505, 17 031 in val_1 and val_2 sets, respectively. Following SCN (Lin et al., 2020), we use val_1 as validation set and val_2 as test set.

DiDeMo. The Distinct Describable Moments (DiDeMo) dataset is first proposed in (Hendricks et al., 2017). It contains over 10k videos

Method	R@1	R@5	mIoU
fully supervised methods			
MCN (2017)	28.10	78.21	41.08
TGN (2018)	28.23	79.26	42.97
MAN (2019a)	27.02	81.70	41.16
I ² N (2021)	29.00	73.09	44.32
weakly supervised methods			
Random	3.75	22.5	22.64
TGA (2019)	12.19	39.74	24.92
WLLN (2019)	19.40	53.10	<u>25.40</u>
VLANet (2020)	19.32	65.68	25.33
FSAN	19.40	<u>57.85</u>	31.92

Table 2: Performance comparison on DiDeMo dataset. The best and second best numbers are highlighted in **bold** and underlined, respectively.

selected from Flickr, and all of them are trimmed to a maximum of 30 seconds and equally divided into six 5-second segments. Therefore, there are only 21 possible moment candidates for each video. The DiDeMo dataset is randomly split into training, validation and test set containing 33 005, 4 180 and 4 021 video-sentence pairs, respectively. For each video, at least 4 annotators are assigned to label text description boundaries.

We exclude the Charades-STA (Gao et al., 2017) dataset because of its limited scale. Charades-STA (Gao et al., 2017) contains about 6k videos, and the contents are mainly indoor activities. However, as shown in Table 1, comparing to other datasets, Charades-STA is limited in terms of total video amount, number of video-sentence pairs, and vocabulary size. The vocabulary size is critical to enriching the linguistic semantics, hence the semantic diversity is limited in the dataset.

Evaluation Metrics. We follow the settings of previous methods (Mithun et al., 2019; Lin et al., 2020). For the ActivityNet-Captions dataset, we report results for intersection-over-union (IoU) $\in \{0.5, 0.3, 0.1\}$ and Recall $\in \{1, 5\}$. On the DiDeMo dataset, considering the limited number of candidates (21) and variance among different annotators, we measure the performance with metrics: Rank@1, Rank@5, and mean intersection over union (mIoU). Here Rank@ k means the percentage of samples where ground truth moment labeled by different annotators are on average ranked higher than k . Following (Hendricks et al., 2017), we discard the worst-ranked ground truth label to reduce the influence of outliers.

4.2 Implementation Details

For fair comparison, we utilize released visual features as previous methods (Mithun et al., 2019; Lin et al., 2020). For videos in ActivityNet-Captions, we adopt C3D (Tran et al., 2015) features. For DiDeMo, we adopt VGG (Simonyan and Zisserman, 2014) features. Note that we report the performance of baseline models using the same features as ours. The dimension of these features are reduced from 4096 to 500 using PCA. The loss weights in to train FSAN are set to 1/3 equally. The hidden dimensions are set to 512 for all datasets. We adopt adam algorithm with an initial learning rate of 0.0001. The batch size is set to 128 for all datasets, and the dropout rate is set to 0.1. The code of FSAN is implemented in pytorch, and is trained on one RTX 3090 GPU.

4.3 Comparisons with State-of-the-art Methods

We compare the proposed FSAN with multiple baselines, including 1) recently published fully supervised state-of-the-arts methods: MCN (Hendricks et al., 2017), ABLR (Yuan et al., 2019), DEBUG (Lu et al., 2019), CMIN (Zhang et al., 2019b), 2D-TAN (Zhang et al., 2020b), TGN (Chen et al., 2018), MAN (Zhang et al., 2019a), I²N (Ning et al., 2021); and 2) some representative weakly supervised methods: TGA (Mithun et al., 2019), SCN (Lin et al., 2020), WLLN (Gao et al., 2019), VLANet (Ma et al., 2020).

Experiments on DiDeMo. Table 2 illustrates the performance comparisons on the DiDeMo dataset. It can be observed from the numbers that FSAN outperforms TGA and WLLN on all three metrics. And comparing to VLANet, FSAN performs overall better, except for R@5. This may be due to the surrogate proposal selection module introduced in VLANet (Ma et al., 2020), which in fact performs a two-stage candidate selection and gets rid of temporally overlapped candidates.

Experiments on ActivityNet-Captions. Table 3 illustrates the performance comparisons on the ActivityNet-Captions dataset. It can be observed that FSAN surpasses previous weakly supervised methods on all metrics. Especially, the FSAN gains about 7% relative improvement on R@1 with IoU=0.1 over SCN. This is because the average length of videos in the ActivityNet-Captions dataset is relatively long, and the candidate generation-selection-criterion framework of

Method	R@1			R@5		
	IoU=0.1	IoU=0.3	IoU=0.5	IoU=0.1	IoU=0.3	IoU=0.5
fully supervised methods						
ABLR (2019)	73.30	55.67	36.79	-	-	-
DEBUG (2019)	-	55.91	39.72	-	-	-
CMIN (2019b)	-	63.61	43.40	-	80.54	67.95
2D-TAN (2020b)	-	58.75	44.05	-	85.65	76.65
weakly supervised methods						
Random	38.23	18.64	7.63	75.74	52.78	24.49
WSLLN (2019)	<u>75.4</u>	42.8	22.7	-	-	-
SCN (2020)	71.48	<u>47.23</u>	<u>29.22</u>	<u>90.88</u>	<u>71.45</u>	<u>55.69</u>
FSAN	78.45	55.11	29.43	92.59	76.79	63.32

Table 3: Performance comparison on ActivityNet-Captions dataset. The best and second best numbers are highlighted in **bold** and underlined, respectively. And “-” means the result on the metric is not reported in the original paper.

Settings	R@1			R@5			mIoU
	IoU=0.1	IoU=0.3	IoU=0.5	IoU=0.1	IoU=0.3	IoU=0.5	
w/o grounding module	82.58	47.99	21.09	86.34	52.60	22.85	34.57
w/o loss on SAP	65.59	46.08	26.31	88.92	66.11	36.72	30.15
w/o cross-modal attention	74.51	43.36	27.03	81.02	65.37	53.41	32.68
w/o inner-modal attention	76.75	50.40	28.71	91.39	71.89	50.81	34.18
full model	78.45	55.11	29.43	92.59	76.79	63.32	36.10

Table 4: Ablation studies on ActivityNet-Captions dataset.

SCN suffers from the limited number of candidates.

4.4 Ablation Study

To investigate the importance of each component in FSAN, we conduct ablation experiments. Results are shown in Table 4 and we give detailed discussions in the next subsections. Note that for comparison, mean intersection-over-union (mIoU) is not reported in the previous subsection, which calculates the average IoU of rank 1st predictions. However, we report and compare mIoU among FSAN variants in this section.

Impact of Grounding Module. To validate the effectiveness of the alignment-based grounding module, we devise a common yet competitive prediction layer upon the visual branch output of ICIM. Specifically, we apply an attention pooling on the text-aware visual features V'' , then apply a three-layer MLP to predict matching score for all possible temporal segments. Results are shown in the 1st row in Table 4. It can be observed that without the grounding module based on the semantic alignment map, the performance drops rapidly on strict

IoU metrics (IoU=0.5, 0.3), which demonstrates the temporal precision improvement by introducing the token-by-clip alignment map.

Impact of Loss on SAP. The 2nd row in Table 4 shows the result without inner-sample loss and outer-sample loss. Under this setting, the grounding performance drops on all metrics compared to full FSAN. The explanation is that without the two losses refining SAP, the FSAN is trained only by the video-level matching loss \mathcal{L}^{tri} . Hence the model can learn video-level coarse semantic alignment, while neglecting token-wise sentence semantics as well as the temporal structure of the video.

Impact of ICIM. We study the role of cross-modal attention and inner-modal attention in the iterative cross-modal interaction module. It can be observed in the 3th and 4th rows in Table 4 that both of them contribute to the grounding performance. Concretely, without inner-modal attention, mIoU drops by 1.92%. While without cross-modal attention, mIoU drops by 3.42%. These results demonstrate the importance of both two attention mechanisms in capturing temporal context among

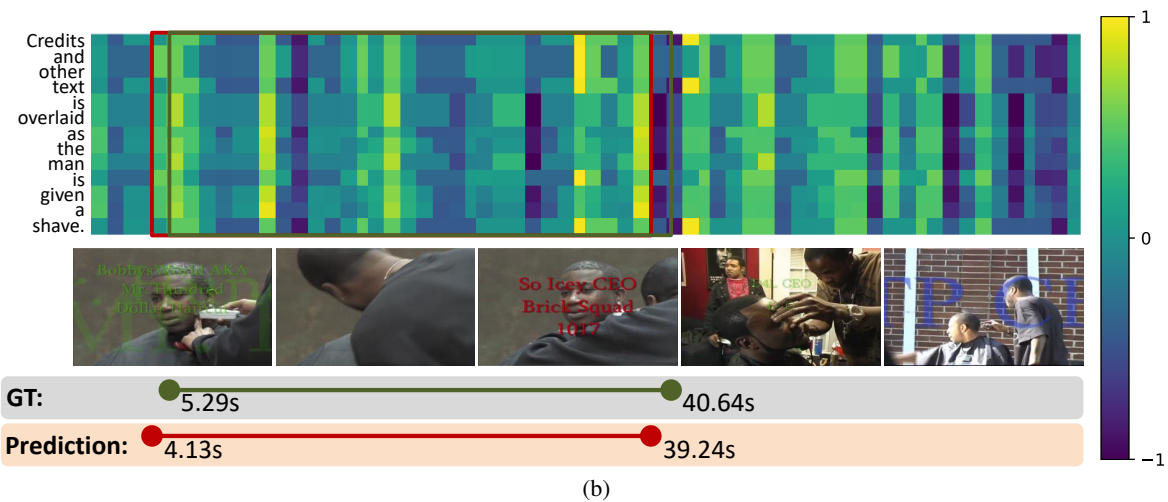
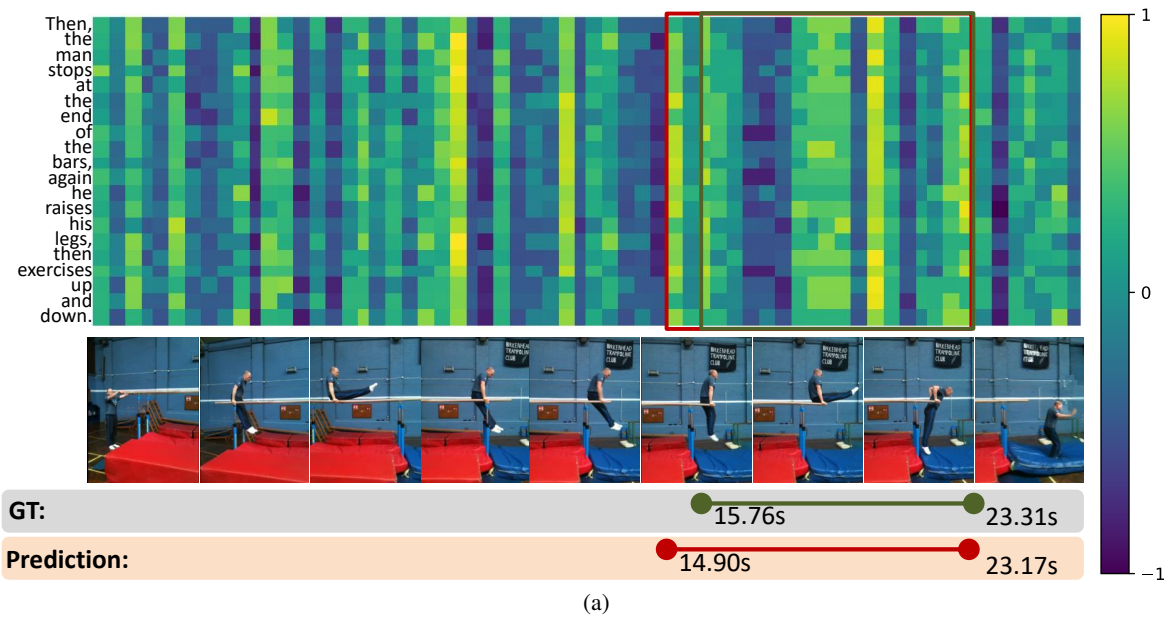


Figure 3: Visualization of the semantic alignment map and grounding results by FSN. The rows and columns on the map correspond to text tokens and video clips, respectively. The green boxes on the map refer to the ground truth temporal region, while the red boxes refer to output of FSN.

video clips, sentence structure among tokens, and token-by-clip cross-modal semantic alignment.

4.5 Analysis and Visualization

We also visualize some examples of the grounding result of FSN in ActivityNet-Captions dataset, as shown in Figure 3. For each video-sentence pair, we visualize the token-clip semantic alignment map, as well as the ground truth and predicted temporal boundaries.

In the first example, the description sentence is long and complicated, consisting of three sequential activities (*stop*, *raise* and *exercise up and down*). FSN achieves high IoU (0.88) on this difficult case, which indicates the strong ability of

FSN to learn fine-grained semantics from both visual and linguistic modalities. The second example shows the ability of FSN to not only detect objects and their actions in video, but also understand abstract descriptions for video (*credits, text*). To better understand abstract descriptions is one of the key points for TLG to develop from action localization. In addition, in the second example, the main action *shave* is blocked in some frames, which is challenging for grounding. Though the blocking reflects in the visualized alignment map, FSN manages to locate the complete moment.

To further analyze the performance of FSN, we plot a graph showing how the performance varies as video length grows. As shown in Figure 4, the




Figure 4: Visualization of the performance variation under different video length. blue points indicate individual videos, and the red line indicates the mIoU of videos in the corresponding length interval.

performance of FSAN is relatively stable when video length grows, with a trend of weakening. For example, the mIoU for shortest videos (2-12s, 130 cases) and longest videos (>230s, 140 cases) are 42.99 and 34.68, respectively.

5 Conclusion

In this paper, we present a novel framework for temporal language grounding, namely Fine-grained Semantic Alignment Network (FSAN). To capture fine-level video-language semantic alignment, we devise an iterative cross-modal interaction module, which enables single-modal representations to interact with each other. Furthermore, we propose to perform temporal grounding based on a semantic alignment map, which alleviates the generation of video candidates. We conduct experiments on two widely-used benchmarks: ActivityNet-Captions and DiDeMo, and achieve state-of-the-art performance on both datasets, which demonstrates the effectiveness of our proposed FSAN.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Contract 61632019.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Zitnick, and Devi Parikh. 2015. VQA: Visual question answering. *International Journal of Computer Vision*, 123:4–31.
- Da Cao, Yawen Zeng, Xiaochi Wei, Liqiang Nie, Richang Hong, and Zheng Qin. 2020. [Adversarial video moment retrieval by jointly modeling ranking and localization](#). In *Proceedings of ACM International Conference on Multimedia*, page 898–906.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229.
- Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. 2020a. Pre-trained image processing transformer. *arXiv preprint arXiv:2012.00364*.
- Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. 2018. [Temporally grounding natural sentence in video](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 162–171.
- Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. 2020b. Generative pretraining from pixels. In *Proceedings of the International Conference on Machine Learning*, pages 1691–1703.
- Shaoxiang Chen and Yu-Gang Jiang. 2019. Semantic proposal for activity localization in videos via sentence query. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8199–8206.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. 2017. TALL: Temporal activity localization via language query. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5277–5285.
- Mingfei Gao, Larry S. Davis, Richard Socher, and Caiming Xiong. 2019. WSLLN: Weakly supervised natural language localization networks. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 1481–1487.
- Runzhou Ge, Jiyang Gao, Kan Chen, and Ramakant Nevatia. 2019. MAC: Mining activity concepts for language-based temporal localization. In *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, pages 245–253.
- Dongliang He, Xiang Zhao, Jizhou Huang, Fu Li, Xiao Liu, and Shilei Wen. 2019. Read, watch, and move: Reinforcement learning for temporally grounding natural language descriptions in videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8393–8400.

- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. 2015. [ActivityNet: A large-scale video benchmark for human activity understanding](#). In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970.
- Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan C. Russell. 2017. Localizing moments in video with natural language. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5804–5813.
- Xincheng Ju, Dong Zhang, Junhui Li, and Guodong Zhou. 2020. Transformer-based label set generation for multi-modal multi-label emotion detection. In *Proceedings of the ACM International Conference on Multimedia*, page 512–520.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 706–715.
- Gen Li, Nan Duan, Yuejian Fang, Daxin Jiang, and Ming Zhou. 2020. Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, page 11336–11344.
- Jingjun Liang, Ruichen Li, and Qin Jin. 2020. Semi-supervised multi-modal emotion recognition with cross-modal distribution matching. In *Proceedings of the ACM International Conference on Multimedia*, page 2852–2861.
- Zhijie Lin, Zhou Zhao, Zhu Zhang, Qi Wang, and Huasheng Liu. 2020. Weakly-supervised video moment retrieval via semantic completion network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, page 11539–11546.
- Chujie Lu, Long Chen, Chile Tan, Xiaolin Li, and Jun Xiao. 2019. DEBUG: A dense bottom-up grounding approach for natural language video localization. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5144–5153.
- Minuk Ma, Sunjae Yoon, Junyeong Kim, Young-Joon Lee, Sunghun Kang, and Chang D. Yoo. 2020. Vlanet: Video-language alignment network for weakly-supervised video moment retrieval. In *Proceedings of the European Conference on Computer Vision*, pages 156–171.
- Niluthpol Chowdhury Mithun, Sujoy Paul, and Amit K. Roy-Chowdhury. 2019. Weakly supervised video moment retrieval from text queries. In *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pages 11584–11593.
- Ke Ning, Lingxi Xie, Jianzhuang Liu, Fei Wu, and Qi Tian. 2021. [Interaction-integrated network for natural language moment localization](#). *IEEE Transactions on Image Processing*, 30:2538–2548.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. 2018. Image transformer. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 4055–4064.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of Empirical Methods in Natural Language Processing*, volume 14, pages 1532–1543.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *Technical report, OpenAI*.
- Dian Shao, Yu Xiong, Yue Zhao, Qingqiu Huang, Yu Qiao, and Dahua Lin. 2018. Find and focus: Retrieve and localize video events with natural language queries. In *Proceedings of the European Conference on Computer Vision*, pages 202–218.
- Botian Shi, Lei Ji, Zhendong Niu, Nan Duan, Ming Zhou, and Xilin Chen. 2020. Learning semantic concepts and temporal alignment for narrated video procedural captioning. In *Proceedings of the ACM International Conference on Multimedia*, page 4355–4363.
- Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. VL-BERT: Pre-training of generic visual-linguistic representations. In *Proceedings of International Conference on Learning Representations*.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. VideoBERT: A joint model for video and language representation learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7463–7472.
- Hao Hao Tan and Mohit Bansal. 2019. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 5100–5111.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. MovieQA: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pages 4631–4640.

- Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4489–4497.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, page 6000–6010.
- Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. 2021a. MaX-DeepLab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. 2021b. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Weining Wang, Yan Huang, and Liang Wang. 2019. Language-driven temporal activity localization: A semantic matching reinforcement learning model. In *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pages 334–343.
- Aming Wu and Yahong Han. 2018. Multi-modal circulant fusion for video-to-language and backward. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1029–1035.
- Huijuan Xu, Kun He, Bryan A. Plummer, Leonid Sigal, Stan Sclaroff, and Kate Saenko. 2019. Multi-level language and vision integration for text-to-clip retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, page 9062–9069.
- Guangyu Robert Yang, Igor Ganiev, Xiao-Jing Wang, Jonathon Shlens, and David Sussillo. 2018. A dataset and architecture for visual reasoning with a working memory. In *Proceedings of European Conference on Computer Vision*, page 729–745.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In *Proceedings of the Neural Information Processing Systems*.
- Jing Yu, Weifeng Zhang, Yuhang Lu, Zengchang Qin, Yue Hu, Jianlong Tan, and Qi Wu. 2020. [Reasoning on the relation: Enhancing visual representation for visual question answering and cross-modal retrieval](#). *IEEE Transactions on Multimedia*, 22(12):3196–3209.
- Yitian Yuan, Tao Mei, and Wenwu Zhu. 2018. To find where you talk: Temporal sentence localization in video with attention based location regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9159–9166.
- Yitian Yuan, Tao Mei, and Wenwu Zhu. 2019. To find where you talk: Temporal sentence localization in video with attention based location regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9159–9166.
- Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S. Davis. 2019a. MAN: Moment alignment network for natural language moment retrieval via iterative graph adjustment. In *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pages 1247–1257.
- Shengyu Zhang, Tan Jiang, Tan Wang, Kun Kuang, Zhou Zhao, Jianke Zhu, Jin Yu, Hongxia Yang, and Fei Wu. 2020a. DeVLBERT: Learning deconfounded visio-linguistic representations. In *Proceedings of the ACM International Conference on Multimedia*, page 4373–4382.
- Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. 2020b. Learning 2d temporal adjacent networks for moment localization with natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, page 12870–12877.
- Zhu Zhang, Zhijie Lin, Zhou Zhao, and Zhenxin Xiao. 2019b. Cross-modal interaction networks for query-based moment retrieval in videos. In *Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3750–3762.
- Zhu Zhang, Zhijie Lin, Zhou Zhao, Jieming Zhu, and Xiuqiang He. 2020c. Regularized two-branch proposal networks for weakly-supervised moment retrieval in videos. In *Proceedings of the ACM International Conference on Multimedia*, page 4098–4106.