

Multi-Scale 2D Temporal Adjacency Networks for Moment Localization with Natural Language

Songyang Zhang, Houwen Peng, Jianlong Fu, Yijuan Lu, and Jiebo Luo

Abstract—We address the problem of retrieving a specific moment from an untrimmed video by natural language. It is a challenging problem because a target moment may take place in the context of other temporal moments in the untrimmed video. Existing methods cannot tackle this challenge well since they do not fully consider the temporal contexts between temporal moments. In this paper, we model the temporal context between video moments by a set of predefined two-dimensional maps under different temporal scales. For each map, one dimension indicates the starting time of a moment and the other indicates the duration. These 2D temporal maps can cover diverse video moments with different lengths, while representing their adjacent contexts at different temporal scales. Based on the 2D temporal maps, we propose a Multi-Scale Temporal Adjacency Network (MS-2D-TAN), a single-shot framework for moment localization. It is capable of encoding the adjacent temporal contexts at each scale, while learning discriminative features for matching video moments with referring expressions. We evaluate the proposed MS-2D-TAN on three challenging benchmarks, *i.e.*, Charades-STA, ActivityNet Captions, and TACoS, where our MS-2D-TAN outperforms the state of the art.



1 INTRODUCTION

TEMPORAL localization is a fundamental problem in video content understanding. There are several related tasks, such as temporal action localization [1]–[5], anomaly detection [6], video summarization [7], [8], and moment localization with natural language [9]–[13]. Among them, moment localization with natural language is the most challenging one due to the complexity of moment description. This task is introduced recently by Gao *et al.* and Hendricks *et al.* [9], [10]. It aims to retrieve a temporary segment from an untrimmed video, as queried by a given natural language sentence. For example, given a query “a guy is playing the saxophone” and a paired video, the task is to localize the best matching moment described by the query, as shown in Fig. 1 (Query A). Video moment localization with natural language has a wide range of applications, such as video question answering [14], video content retrieval [15], dense video captioning [16], as well as video storytelling [17].

Most current language-queried moment localization methods follow a two-step pipeline [9], [10], [18]–[20]. It first utilizes sliding windows to generate moment segments from the input videos. Next, each moment is matched with the query sentence to estimate whether it is the target moment. This pipeline models different moments independently and neglects their temporal contexts. Thus it is difficult to precisely localize a moment that requires context information from other moments. For example, as shown in Fig. 1 (Query C), it targets to localize the query “the guy plays the saxophone again” in the video. If the model only watches the temporal moments in the latter part of the video, it cannot localize the described “again” moment accurately. Moreover, as shown in Fig. 1 (Query B), there are many temporal candidates overlapping with the target moment (the visualized time

slots). These candidates are relevant with respect to visual content, but they depict different semantics. Consequently, it may be difficult for existing methods to distinguish these visually similar moments since they process each temporal candidate individually.

To address these critical problems in moment localization with natural language, we propose a novel method, 2D Temporal Adjacency Networks (2D-TAN). The core idea is to localize a target moment from a two-dimensional temporal map, instead of a conventional one-dimensional sequential modeling as presented in Fig. 1. More specifically, for a temporal map with time unit τ , the (i, j) -th location on it represents a candidate moment starting from the timestamp $i\tau$ and lasting for $(j + 1)\tau$. This kind of 2D temporal map covers various video moments with different lengths, while representing their adjacent relations. In this fashion, 2D-TAN can perceive more moment context information when predicting whether a moment is related to other temporal segments. On the other hand, the adjacent moments in the map have content overlap but may depict different semantics. Considering them as a whole, 2D-TAN is able to learn discriminative features to distinguish them.

The time unit τ on the 2D temporal map decides the localization granularity. To guarantee a fine-grained localization of action instances, especially for precise temporal boundaries, we further build up multi-scale 2D temporal maps, in which multiple choices of time unit are considered. We construct K 2D temporal maps, each of which enumerates possible moments based on its time unit and longest duration. Such multi-scale modeling allows the localization networks to perceive moment candidates with more diverse temporal ranges and richer context. Moreover, the multi-scale maps can be regarded as a sparse sampling with different intervals on a dense single-scale temporal map, thus reducing the computation costs of moment feature extraction and temporal context modeling. Compared with the dense single-scale modeling, sparse sampling reduces the total number of moment candidates from $O(N^2)$ to

- S. Zhang and J. Luo are with the Department of Computer Science, University of Rochester, Rochester, NY, 14627.
E-mail: szhang83, jluo@cs.rochester.edu
- H. Peng, J. Fu and Y. Lu are with Microsoft.
E-mail: {houwen.peng,jianf, yijlu}@microsoft.com

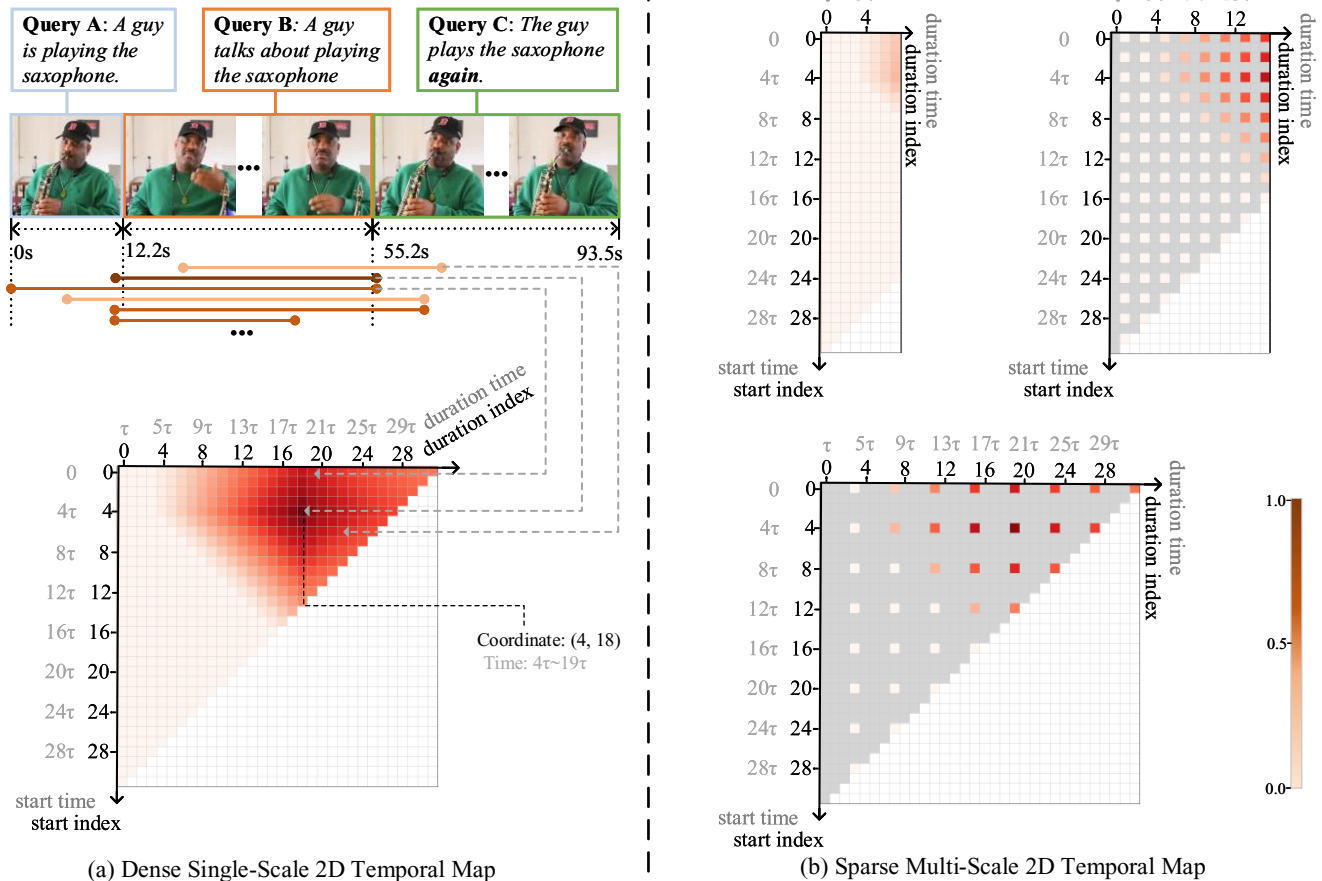


Fig. 1: Examples of the 2D temporal map. (a) Dense Single-scale 2D Temporal Map: the **black** vertical and horizontal axes represent the start and the duration indices while the **gray** axes represent the corresponding start timestamp and duration of a moment. The values in the 2D map, highlighted by **red** color, indicate the matching scores between the moment candidates and the target moment. Here, τ is a predefined short duration. The white boxes in the map indicate invalid moments. (b) Sparse Multi-Scale 2D Temporal Map: the sparse multi-scale 2D temporal map is composed of a series of 2D maps under different time units (τ , 2τ and 4τ in this figure). The color of boxes and axes follow the previous definition. In addition, the **gray** boxes on the 2D map represent the valid moments that are not selected. In this configuration, we are able to reduce the computational cost by modeling on smaller maps.

$O(N)$. Modeling context on multiple smaller maps, instead of a single large one, can further reduce the complexity from quadratic to linear (see Sec. 4.5.1 for detailed proof).

Some preliminary ideas of this paper have appeared in our earlier work [21]. In this paper, we extend the previously proposed single-scale 2D moment localization method to a multi-scale version. The new model considers the adjacent relations of moments at different temporal scales while achieving better performance with faster speed and less memory consumption. The main contributions of this paper are summarized as following:

- We introduce a novel *two-dimensional temporal map* for modeling the temporal adjacent relations of video moments. Compared with previous methods, the 2D temporal map enables the model to perceive more video context information and learn discriminative features to distinguish the moments with complex semantics.
- We propose a *Multi-Scale 2D Temporal Adjacency Network*, i.e., MS-2D-TAN, for moment localization with natural language. The multi-scale modeling allows our

model to have a larger receptive field and obtain richer context. Meanwhile, it reduces the complexity of generating moments from quadratic to linear, which makes the dense video prediction more efficient. Without any sophisticated video-language cross-modality fusion, MS-2D-TAN achieves competitive performance in comparison with the state-of-the-art methods on three benchmark datasets. Our code and models are publicly available at <https://github.com/microsoft/2D-TAN>.

2 RELATED WORK

There are two major sub-fields in the temporal localization in untrimmed videos: temporal action localization and moment localization with natural language. Temporal action localization aims to predict the start and the end time and the category of an activity instance in untrimmed videos. The representative methods include two-stage temporal detection methods [1] and one-stage single shot methods [22]. This task is limited to pre-defined simple actions and cannot handle complex events in the real world. Therefore, moment

localization with natural language [9], [10] is introduced recently to tackle this problem. In the remainder of this section, we first review recent clip-based and moment-based methods, and then discuss their cross-modality alignment. In addition, we introduce some related methods in other tasks and their difference from our preliminary work.

Localizing moments in videos by query sentences is very challenging. It not only needs to understand video content, but also requires to align the semantics between video and language. A video clip is a set of short video segments with fixed duration and interval, while a moment is a sequence of video clips. Depending on whether the moment features are extracted, we divide current methods into two categories, clip-based methods (not extracted) and moment-based methods (extracted).

Clip-based methods. The main idea of clip-based methods is to align video clips with language directly and predict matching scores without extracting moment features. In general, there are three common ways to map a clip to a moment score: anchor-based methods, anchor-free methods, and RL-based methods. Anchor-based methods define a set of anchors with a fixed length for each clip [23]–[25], while anchor-free methods directly predict the start and the end time through classification [26], [27] or regression [19], [28]–[35]. RL-based methods model the task as a sequential decision-making problem and solve it by reinforcement learning [36]–[39]. In terms of their context modeling, most approaches [19], [23], [24], [27]–[29], [32] gradually aggregate the context information through a recurrent structure. Some approaches [25], [30], [36] model surrounding clips as the local context using 1D convolution layers, while other approaches model the entire clip as the global context through self-attention modules [26], [31], [33], [34]. Since clips are the shortest moments, the clip-level context is a subset of moment-level context. Thus, all these context modeling methods ignore a large part of moment-level context. In contrast, our focus is to model contexts on the moment level and to learn discriminative moment features to boost the localization performance.

Moment-based methods. Moment-based methods extract moment-level features and learn a matching function between features and the query. Most approaches predict moment scores in one stage, while some approaches follow a cascaded strategy and filter moments by multiple stages [40], [41]. Different from clip-based methods, existing moment-based methods integrate temporal context in two common ways. One way is to use the whole video as the global context. For example, Hendricks *et al.* [10] and He *et al.* [36] concatenate each moment feature with the global video feature [10] as the moment representation. Wang *et al.* concatenate semantic features with the global video feature [37]. Another way is to use the surrounding clips as the local context for a moment. Gao *et al.*, Liu *et al.*, Song *et al.* and Ge *et al.* concatenate the moment feature with clip features before and after current clip as its representation [9], [18]–[20]. Since these methods only consider one or two specific moments, the rich context information from other possible moments is ignored. One recent exploration [42] of aggregating the context from other moments is Graph Convolutional Network (GCN). However, one inherent feature of graph is *node permutation invariance* [43], switching

any two nodes does not change the result. Therefore, the graph modeling ignores the temporal ordering of different moments. In contrast, our method considers all the neighboring moments as the context and models the moments on a 2D convolution network. The 2D convolution network can naturally preserve the relative position of different moments. This design enables the model to perceive more context information and learn more discriminative features.

Video and language cross-modality alignment. The simplest way for cross-modal alignment is to directly multiply or concatenate the clip/moment features with sentence features [9], [10], [30]. Recent approaches improve video and language alignment from several different directions, such as cross-modal attention, sentence syntactic modeling, and compositional reasoning. The key idea of cross-modal attention is to attend relevant video clips/moments or query words from another modality. Some methods attend relevant video features through words [19], [35], while most other methods attend both the relevant video features and words via the co-attention module [20], [23], [24], [26], [27], [30], [32], [34], [35], [42], [44], [45]. For sentence syntactic modeling, Zhang *et al.* [24] enhance the sentence modeling with the queries' syntactic graph. For compositional reasoning, Liu *et al.* [44] explicitly use the compositionality in natural languages for temporal reasoning in videos. Instead of using the complex attention modules, our proposed MS-2D-TAN model only adopts a simple multiplication operation for visual and language feature fusion.

Discussion with related methods in other tasks The self-attention operation in document modeling enumerate all possible word pairs. This enumeration is similar to our 2D temporal map, where the start and the end time pairs are enumerated. Due to the enumeration, both tasks face the same problem: memory cost and speed scale up quadratically with the sequence length. Several methods are recently proposed to improve the efficiency in the long document tasks. Beltagy *et al.* [46] combine global attention, dilated convolution and sliding window together to reduce the complexity to linear. Sukhbaatar *et al.* [47] design an adaptive attention span, which learns a hard mask to ignore the long-term relations. In this paper, we focus on reducing the computational cost in the video domain. Our study is also related to some contemporaneous and recent work [48], [49] that apply 2D temporal map to temporal action localization task. In contrast to their work, our work focus on moment localization with natural language. Besides, we also propose a multi-scale 2D temporal map to improve the efficiency and the performance of the dense 2D temporal map.

Difference from our preliminary work. Some preliminary ideas in this paper appeared in the conference version [21], where 2D-TAN is proposed to learn discriminative moment features through a 2D temporal map. A sparse single-scale map is introduced to reduce the computational cost of feature extraction from quadratic to linear. However, the complexity of the 2D temporal adjacent network remains quadratic. Compared with [21], the proposed MS-2D-TAN further reduce the complexity of the 2D temporal adjacent network from quadratic to linear. The experimental results (Sec. 4) show clearly that the new model is more efficient (Fig. 4 and Fig. 5) and achieves better performance in three benchmark datasets (Table. 1, 2 and 3).

3 OUR APPROACH

In this section, we first introduce the problem formulation of moment localization with natural language. Next, we present the pipeline of our proposed network, including language representation, video representation and multi-scale 2D temporal adjacency network, as shown in Fig. 2. Finally, we present the training and the inference of our network.

3.1 Problem Formulation

Given an untrimmed video V and a sentence S as a query, the task is to retrieve the best matching temporary moment M specified by the query. More specifically, we denote the query sentence as $S = \{s_i\}_{i=0}^{l^S-1}$, where s_i represents a single word, and l^S is the total number of the words. The input video is a frame sequence, i.e. $V = \{x_i\}_{i=0}^{l^V-1}$, where x_i represents a frame in a video and l^V is the total number of frames. The retrieved moment starting from frame x_i to x_j delivers the same semantic meaning as the query S .

3.2 Language Representation via Sequential Embedding

We first extract the feature of an input query sentence. For each word s_i in the input sentence S , we generate its embedding vector $\mathbf{w}_i \in \mathbb{R}^{d^S}$ by the GloVe word2vec model [50], where d^S is the vector length. We then sequentially feed the word embeddings $\{\mathbf{w}_i\}_{i=0}^{l^S-1}$ into a three-layer bidirectional LSTM network [51], and use its average output as the feature representation of the input sentence, i.e. $\mathbf{f}^S \in \mathbb{R}^{d^S}$. The extracted feature encodes the language structure of the query sentence, thus describes the moment of interest.

3.3 Video Representation via 2D Temporal Feature Map

In this section, we extract moment features from the input video stream and then construct a 2D temporal feature map.

Given an input video, we first segment it into small non-overlapped video clips, where each clip consists of T contiguous frames. For each video clip, we extract its feature using pre-trained CNN model (see experiments section for details). To generate more compressed video clip representation in the channel dimension, we then feed the video clip feature into a fully connected layer with d^V output channels. Therefore, the final representation of compressed video clips is represented as $\{\mathbf{f}_i^V\}_{i=0}^{N-1}$, where d^V is the number of output channels and N is the total number of video clips.

The N video clips serve as the basic elements for moment candidate construction. Thus, we build up feature maps of moment candidates by the video clip features $\{\mathbf{f}_i^V\}_{i=0}^{N-1}$. Previous works extract moment features from clip features in two ways: pooling [10] or stacked convolution [42]. In this work, we follow the stacked convolution design. In the stacked convolution, the output of each layer are moment features, where these moments have equal length but start at different time, as shown in Fig. 2. In order to generate moment features, a simple way is to stack N convolution layers, where each layer has a kernel of size 2 and stride 1 (except for the first layer, which has kernel size 1 and stride 1). However, when N is large, it is infeasible

to fit the data into memory. Instead of enumerating all of them, we conduct a sparse sampling strategy as following: We first define two constant number A and K . For the first convolution layer, we set both the kernel size and stride to 1. For the $(i+1)\frac{A}{2}$ -th layers ($1 \leq i \leq K-1$), the kernel size and stride are set to 3 and 2, respectively. All the rest layers have kernels of size 2 and strides 1. We also add batch norm after each convolution layer and use \tanh as the activation function. Therefore, there are $\frac{(K+1)A}{2}$ convolution layers in total. In this way, we densely sample moments of short duration, and gradually increase the sampling interval when the moment duration becomes long. In more details, when the number of sampled clips is small, i.e. $N \leq A$, we enumerate all possible moments as candidates. When N becomes large, i.e. $N > A$, a moment starting from clip v_a to v_b is selected as the candidate when satisfying the following condition $G(a, b)$:

$$G(a, b) \Leftarrow (a \bmod s=0) \ \& \ (b \bmod s=0), \quad (1)$$

where a and b are the indexes of clips, s is defined as $s = 2^{k-1}$, $k = \lceil \log_2(\frac{b-a+1}{A}) + 1 \rceil$ and $\lceil \cdot \rceil$ is the ceiling function. If $G(a, b) = 1$, the moment is selected as the candidate, otherwise, it is not selected. This sampling strategy can largely reduce the number of moment candidates, as well as the computational cost.

Different from previous methods which directly operate on an individual video moment, we restructure the whole sampled moments to a 2D temporal feature map, denoted as $\mathbf{F}^M \in \mathbb{R}^{N \times N \times d^V}$. The 2D temporal feature map \mathbf{F}^M consists of three dimensions: the first two dimensions N represent the start and duration clip indexes respectively, while the third one d^V indicates the feature dimension. The feature of a moment starting from clip v_a and during b clips is located at $\mathbf{F}^M[a, b, :]$ on the feature map. Denoting the i -th output at j -th convolution layer as $\mathbf{f}^M[i, j] \in \mathbb{R}^{d^V}$, it corresponds to the (a, b) -th location on the feature map \mathbf{F}^M , where

$$\begin{aligned} a &= 2^k i, \\ b &= \begin{cases} j & \text{if } j < A, \\ A + 2^{\lceil \frac{2(j-A+1)}{A} \rceil} \cdot (j - A + 1) - 1 & \text{otherwise.} \end{cases} \end{aligned} \quad (2)$$

Noted that, the moment's start and duration clip indexes a and b should satisfy $a + b < N$. Therefore, on the 2D temporal feature map, all the moment candidates locating at the region of $a + b \geq N$ are invalid, i.e. the lower triangular part of the map, as shown in Figure 2. The values in this region are padded with zeros in implementation.

3.4 Multi-Scale 2D Temporal Adjacency Network

After obtaining the language and video features, we then predict the best matching moment queried by the sentence from all candidates.

We first fuse the 2D temporal feature maps \mathbf{F}^M with the sentence feature \mathbf{f}^S . Specifically, we project these two cross-domain features to the same subspace via fully connected layers, and then fuse them through Hadamard product and ℓ_2 normalization as

$$\mathbf{F}^F = \|(\mathbf{w}^S \cdot \mathbf{f}^S \cdot \mathbf{1}^T) \odot (\mathbf{W}^M \cdot \mathbf{F}^M)\|_F, \quad (3)$$

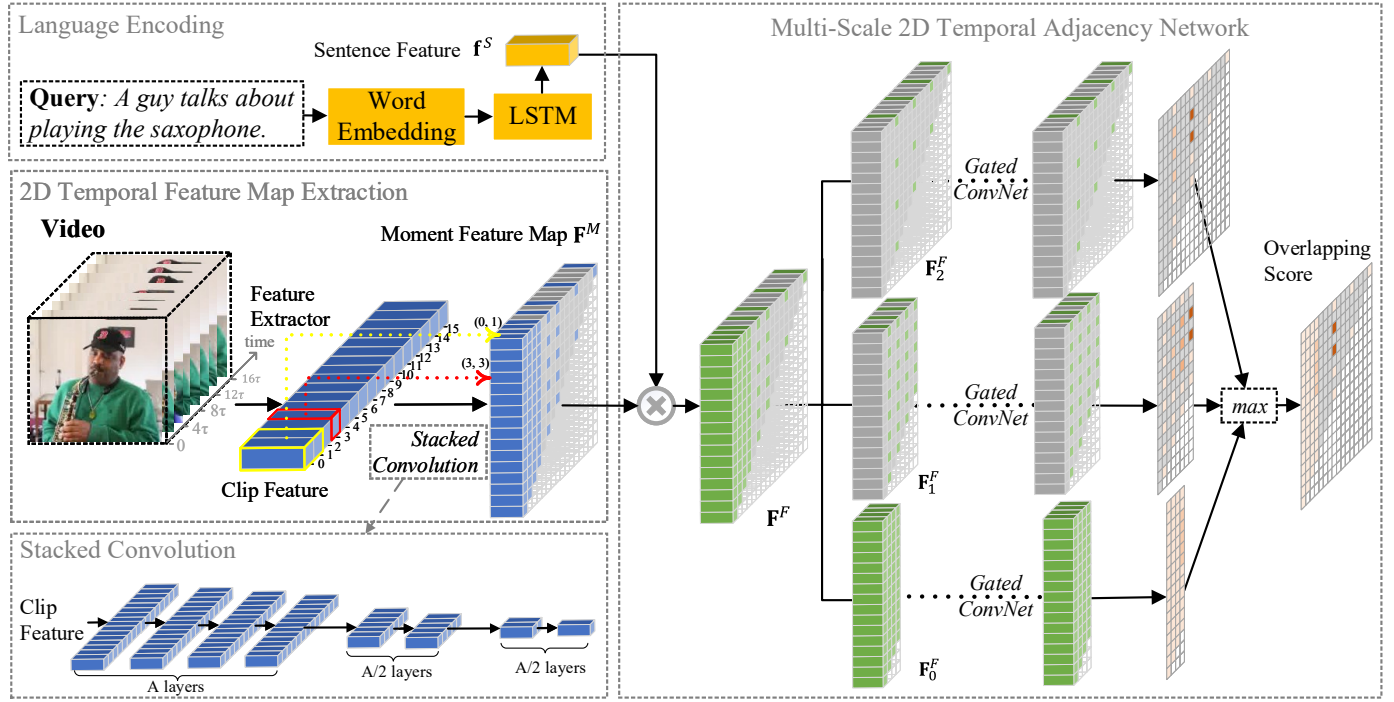


Fig. 2: Our proposed framework. The framework is composed of three modules: the language encoding module, the 2D temporal feature map extraction module, and the multi-scale 2D temporal adjacency network.

where \mathbf{w}^S and \mathbf{W}^M represent the learnt parameters of the fully connected layers, $\mathbb{1}^T$ is the transpose of an all-ones vector, \odot is Hadamard product, and $\|\cdot\|_F$ denotes Frobenius normalization. The fused feature map is denoted as $\mathbf{F}^F \in \mathbb{R}^{N \times N \times d^F}$, where d^F is the feature dimension after fusion.

Different from our preliminary work which directly builds up a temporal adjacency network over the single-scale map \mathbf{F}^F , in this paper, we build up multi-scale temporal adjacency networks over the sparse multi-scale maps, which executes faster and takes lower memory footprints. In more details, we first restructure the single-scale map \mathbf{F}^F to K multi-scale maps. The k -th sparse map is sampled from the fused feature map with interval 2^k . Specifically, all the $(2^k i - 1, 2^k j - 1)$ -th locations on \mathbf{F}^F are sampled, where $1 \leq i \leq N$ and $1 \leq j \leq A$ and A defines the number of anchors at each scale. In this way, the shortest and longest moments in the k -th map are of size 2^k and $2^k A$, respectively. We denote the multi-scale map as $\{\mathbf{F}_k^F | \mathbf{F}_k^F \in \mathbb{R}^{N \times 2^k A \times d^F}, 0 \leq k \leq K-1\}$.

Next, we build up temporal adjacency networks over the multi-scale 2D feature maps. In more details, each map corresponds to L gated convolutional layers [52] with kernel size of κ , dilation size of 2^k and stride size of 2^k . The output of the L layers keeps the same shape as the input fused feature map through zero padding at each layer. This design enables the model to gradually perceive more context of adjacent moment candidates, while learn the difference between moment candidates. Moreover, the receptive field of the network is large, thus it can observe a large portion of video content, resulting in learning the temporal contexts.

Finally, we predict the matching scores of moment candidates with the given sentence on the multi-scale 2D tem-

poral maps. We feed the output feature maps into fully connected layers and the *sigmoid* function separately, then generate multi-scale 2D score maps. The valid scores on each score map are then collected, denoted as $P_k = \{p_k^i\}_{i=1}^{C_k}$, where C_k is the total number of valid moment candidates of the k -th map. Each value p_k^i on the k -th map represents the matching score of a moment candidate with the sentence.

3.5 Training and Inference

During training, we adopt a scaled *IoU* value as the supervision signal, rather than a hard binary score. Specifically, for each moment candidate, we compute its *IoU* o_k^i with the ground truth moment. The *IoU* score o_k^i is then scaled with thresholds 0.5 as

$$y_k^i = \begin{cases} 0 & o_k^i \leq 0.5, \\ 2o_k^i - 1 & o_k^i > 0.5, \end{cases} \quad (4)$$

and y_k^i serves as the supervision label. Our network is trained by a binary cross entropy loss as

$$Loss = \sum_{k=0}^{K-1} \sum_{i=1}^{C_k} y_k^i \log p_k^i + (1 - y_k^i) \log(1 - p_k^i), \quad (5)$$

where K is the total number of maps, p_k^i and C_k are the output score of a moment and the total number of valid candidates on the k -th map.

During inference, the score maps $\{P_k, 0 \leq k \leq K-1\}$ are recovered to a single-scale map P' based on the moment location at the original single-scale map. Noted that some moments are predicted by more than one score map. And we choose the highest score as its final prediction.

4 EXPERIMENTS

In this section, we first introduce the datasets used in our experiments and then go through the experiment settings. Next, we compare our proposed MS-2D-TAN with previous work. Finally, we conduct ablation studies and visualize experimental results.

4.1 Datasets

We evaluate our methods on the following three datasets:

Charades-STA [9]. It contains 9,848 videos of daily indoor activities. It is originally designed for action recognition and localization. Gao *et al.* [9] extend the temporal annotation (*i.e.* the start and end time of moments) of this dataset with language descriptions. Charades-STA contains 12,408 moment-sentence pairs in training set and 3,720 pairs in testing set.

ActivityNet Captions [53]. It consists of 19,209 videos, whose content are diverse and open. It is originally designed for dense video captioning, and recently introduced to moment localization with natural language, since these two tasks are reversible [23], [24]. Following the experimental setting in [24], we use val_1 as the validation set and val_2 as the testing set, which have 37,417, 17,505, and 17,031 moment-sentence pairs for training, validation, and testing, respectively. Currently, this is the largest dataset in this task.

TACoS [54]. It consists of 127 videos selected from the MPII Cooking Composite Activities video corpus [55], which contain different activities happened in the kitchen room. Regneri *et al.* extend the sentence descriptions by crowd-sourcing. A standard split [9] consists of 10,146, 4,589, and 4,083 moment-sentence pairs for training, validation and testing, respectively.

We also demonstrate the distribution of video duration, target moment duration and their ratio of these three datasets in Fig. 3. We can observe that videos in TACoS dataset have more variant and longer duration than others (see Fig. 3 Left). Charades-STA has shorter target moments compared to ActivityNet Captions and TACoS (see Fig. 3 Middle). By comparing the ratio between the target moment and the entire video (see Fig. 3 Right), we find that most target moments in TACoS have smaller ratios.

4.2 Evaluation Metric

Following the setting in previous work [9], we evaluate our model by computing *Rank n@m*. It is defined as the percentage of language queries having at least one correct moment retrieval in the top-*n* retrieved moments. A retrieved moment is correct when its IoU with the ground truth moment is larger than *m*. There are specific settings of *n* and *m* for different datasets. Specifically, we report the results as $n \in \{1, 5\}$ with $m \in \{0.5, 0.7\}$ for Charades-STA dataset, $n \in \{1, 5\}$ with $m \in \{0.3, 0.5, 0.7\}$ for ActivityNet Captions dataset, and $n \in \{1, 5\}$ with $m \in \{0.1, 0.3, 0.5, 0.7\}$ for TACoS dataset.

4.3 Implementation Details

The performance of video features plays an important role in moment localization with natural language. To make a

thorough and fair comparison, we collect the most commonly used features in previous works, including VGG [56], C3D [57] and I3D [58] features. The details about video feature extractor are illustrated as follows:

- **VGG.** Following [42], we use VGG16 pre-trained on ImageNet [59]. Specifically, videos are decoded at 24 fps and the output of *fc7* layer after *ReLU* are extracted at 6 fps. The clip feature corresponds to every 4 consecutive frame. Therefore, each clip corresponds to 1 second.
- **C3D.** Following [45], we use C3D network pre-trained on sport1M [60]. Specifically, videos are decoded at 16 fps and the output of *fc6* layer after *ReLU* are extracted for every 16 consecutive frames. Each video clip corresponds to 1 second.
- **I3D.** Following [25], we use I3D network pre-trained on Kinetics [58]. Specifically, videos are decoded at 25 fps and the output of the last average pooling layer are extracted for every 16 consecutive frames. Therefore, each video clip corresponds to 0.64 second. For Charades-STA, we also finetune the I3D on the Charades dataset [61].

Some predicted moments highly overlap with each other. To reduce the redundant prediction, we adopt non-maximum suppression (NMS) based on the predicted scores P' . We fix the IoU threshold for NMS to 0.49 for all the experiments. After NMS, we use the top-*n* ranked moments for evaluation. Please note that NMS does not affect the top-1 result.

During training, we adopt sliding window to random select *N* consecutive clips. For a fair comparison, we set the number of hidden states $H = 512$, the window size $N = 64$, the number of scales $K = 3$, the number of anchors $A = 16$, the convolution kernel size $\kappa = 17$ and the number of layers $L = 2$ for all three datasets as default. All these anchors and scale settings is able to cover at least 95% target moments of the training set with an IoU threshold of 0.7. Finetuning these hyperparameters on specific dataset and feature type could get better performance as demonstrated in Table 3 (the results of MS-2D-TAN*), where $H = 512$, $N = 512$, $A = 8$, $\kappa = 9$, and $L = 2$. To train MS-2D-TAN from scratch, the learning rate is set to 0.0001 without weight decay. The batch size is set to 32 and Adam [62] is used as the optimizer.

4.4 Comparison with State-of-the-Art Methods

We evaluate the proposed MS-2D-TAN approach on three benchmark datasets, and compare it with recently proposed state-of-the-art methods [1] including:

- **clip-based methods:**
 - **anchor based methods:** TGN [23], CMIN [24] and CBP [45], SCDM [25],
 - **anchor free methods:** ACRN [19], ROLE [30], SLTA [35], DEBUG [34], VSLNet [26], GDP [33], LGI [31], ABLR [27], TMLGA [32], ExCL [28] and DRN [29],
 - **reinforcement learning based methods:** RWM-RL [36], SM-RL [37], TripNet [38] and TSP-RPL [39],
- **moment-based methods:**

1. In addition to the video features, SLTA [35], SM-RL [37] use extra object features extracted from Faster R-CNN [64].




Fig. 3: Comparing Charades-STA, ActivityNet Captions and TACoS. **Left:** comparing video duration. **Middle:** comparing target moment duration. **Right:** comparing the ratio between target moment duration and video duration.

TABLE 1: Performance Comparison on Charades-STA.

Method	Rank1@		Rank5@	
	0.5	0.7	0.5	0.7
VGG features				
MCN [10]	17.46	8.01	48.22	26.73
MAN [42]	<i>41.24</i>	<i>20.54</i>	<i>83.21</i>	51.85
SM-RL [37]	24.36	11.17	61.25	32.08
SAP [41]	27.42	13.36	66.37	38.15
MS-2D-TAN	45.65	27.20	86.72	56.42
C3D features				
CTRL [9]	23.63	8.89	58.92	29.52
ACRN [19]	20.26	7.64	71.99	27.79
ROLE [30]	21.74	7.82	70.37	30.06
VAL [20]	23.12	9.16	61.26	27.98
ACL-K [18]	30.48	12.20	64.84	35.13
DEBUG [34]	37.39	17.69	—	—
GDP [33]	<i>39.47</i>	18.49	—	—
RWM-RL [36]	36.70	—	—	—
QSPN [40]	35.60	15.80	79.40	45.40
SLTA [35]	22.81	8.25	72.39	31.46
ABLR [27]	24.36	9.01	—	—
TripNet [38]	36.61	14.50	—	—
CBP [45]	36.80	<i>18.87</i>	70.94	50.19
TSP-PRL [39]	37.39	17.69	—	—
MS-2D-TAN	41.10	23.25	81.53	<i>48.55</i>
I3D features				
SCDM [25]	<i>54.44</i>	<i>33.43</i>	74.43	<i>58.08</i>
DRN [29]	53.09	31.75	<i>89.06</i>	<i>60.05</i>
MS-2D-TAN	56.64	36.21	89.14	61.13
I3D features finetuned on Charades				
ExCL [28]	44.10	22.40	—	—
TMLGA [32]	52.02	33.74	—	—
LGI [31]	<i>59.46</i>	<i>35.48</i>	—	—
VSLNet [26]	54.19	35.22	—	—
MS-2D-TAN	60.08	37.39	89.06	59.17

TABLE 2: Performance comparison on ActivityNet Captions.

Method	Rank1@			Rank5@		
	0.3	0.5	0.7	0.3	0.5	0.7
C3D features						
MCN [10]	39.35	21.36	6.43	68.12	53.23	29.70
CTRL [9]	47.43	29.01	10.34	75.32	59.17	37.54
TGN [23]	43.81	27.93	—	54.56	44.20	—
ACRN [19]	49.70	31.67	11.25	76.50	60.34	38.57
DEBUG [34]	55.91	39.72	—	—	—	—
GDP [33]	56.17	39.27	—	—	—	—
CMIN [24]	63.61	43.40	23.88	<i>80.54</i>	67.95	<i>50.73</i>
RWM-RL [36]	—	36.90	—	—	—	—
QSPN [40]	52.13	33.26	13.43	77.72	62.39	40.78
ABLR [27]	55.67	36.79	—	—	—	—
TripNet [38]	48.42	32.19	13.93	—	—	—
SCDM [25]	54.80	36.75	19.86	77.29	64.99	41.53
CBP [45]	54.30	35.76	17.80	77.63	65.89	46.20
TSP-PRL [39]	56.08	38.76	—	—	—	—
DRN [29]	—	<i>45.45</i>	<i>24.36</i>	—	<i>77.97</i>	50.30
LGI [31]	58.52	41.51	23.07	—	—	—
MS-2D-TAN	<i>61.04</i>	46.16	29.21	87.30	78.80	60.85
I3D features						
ExCL [28]	<i>62.30</i>	42.70	24.10	—	—	—
TMLGA [32]	51.28	33.04	19.26	—	—	—
VSLNet [26]	63.16	<i>43.22</i>	<i>26.16</i>	—	—	—
MS-2D-TAN	62.09	45.50	28.28	87.61	79.36	61.70

- **global/local context methods:** MCN [10], CTRL [9], ACL-K [18] and VAL [20],
- **graph based methods:** MAN [42],
- **cascade refinement methods:** QSPN [40] and SAP [41].

The results are summarized in Table 1–3. The values highlighted by **bold** and *italic* fonts indicate the top-2 methods, respectively. All results are reported in percentage (%).

The results show that MS-2D-TAN performs the best in various scenarios on all three benchmark datasets across different criteria. In most cases, MS-2D-TAN ranks the first or the second place [2]. It is worth noting that on TACoS

2. Comparing CBP on Rank5@0.7 in Table 1, our MS-2D-TAN is worse due to the IoU threshold for NMS. If the threshold is set 0.5, the Rank5@0.5 and Rank5@0.7 of our model are 77.58 and 53.04, which outperform CBP.

dataset (see Table 3), our MS-2D-TAN surpasses the previous best approach CBP [45], by approximate 18 points and 25 points in term of Rank1@0.3 and Rank5@0.3, respectively. Moreover, on the large-scale ActivityNet Captions dataset, MS-2D-TAN also outperforms the top ranked method DRN [29] and VSLNet [26] with respect to IoU@0.5 and 0.7. It validates that MS-2D-TAN is able to localize the moment boundary more precisely.

In more details, by comparing MS-2D-TAN with other related methods, we obtain several observations. First, we compare MS-2D-TAN with previous moment-based methods: MCN [10], CTRL [9], ACL-K [18], VAL [20] and MAN [42]. From the results in Table 1–3, we observe that our MS-2D-TAN achieves superior results than concatenating local/global context. The reason is that independently matching the sentence with moment candidates ignores the temporal contexts, and cannot distinguish the small differences between the overlapped moments. Differently, our proposed MS-2D-TAN models the relations between moment candidates by a series of sparse 2D temporal maps, and enables the network to perceive more context information from the adjacent moment candidates. Hence, it gains

TABLE 3: Performance comparison on TACoS.

Method	Rank1@				Rank5@			
	0.1	0.3	0.5	0.7	0.1	0.3	0.5	0.7
VGG features								
MCN [10]	14.42	—	5.58	—	37.35	—	10.33	—
SM-RL [37]	26.51	20.25	15.95	—	50.01	38.47	27.84	—
SAP [41]	31.15	—	18.24	—	53.51	—	28.11	—
MS-2D-TAN	50.64	43.31	35.27	23.54	78.31	66.18	55.81	38.09
C3D features								
CTRL [9]	24.32	18.32	13.30	6.96	48.73	36.69	25.42	15.33
MCF [63]	25.84	18.64	12.53	—	52.96	37.13	24.73	—
TGN [23]	41.87	21.77	18.9	11.88	53.40	39.06	31.02	15.26
ACRN [19]	24.22	19.52	14.62	—	47.42	34.97	24.88	—
ROLE [30]	20.37	15.38	9.94	—	45.45	31.17	20.13	—
VAL [20]	25.74	19.76	14.74	—	51.87	38.55	26.52	—
ACL-K [18]	31.64	24.17	20.01	—	57.85	42.15	30.66	—
DEBUG [34]	41.15	23.45	—	—	—	—	—	—
GDP [33]	39.68	24.14	—	—	—	—	—	—
CMIN [24]	32.48	24.64	18.05	—	62.13	38.46	27.02	—
QSPN [40]	25.31	20.15	15.23	—	53.21	36.72	25.30	—
SLTA [35]	23.13	17.07	11.92	—	46.52	32.90	20.86	—
ABLR [27]	34.70	19.50	9.40	—	—	—	—	—
TripNet [38]	—	23.95	19.17	9.52	—	—	—	—
SCDM [25]	—	26.11	21.17	—	—	40.16	32.18	—
CBP [45]	—	27.31	24.79	19.10	—	43.64	37.40	25.59
DRN [29]	—	—	23.17	—	—	—	33.36	—
MS-2D-TAN	49.24	41.74	34.29	21.54	78.33	67.01	56.76	36.84
MS-2D-TAN*	52.39	45.61	35.77	23.44	79.26	69.11	57.31	36.09
I3D features								
ExCL [28]	—	45.50	28.00	13.80	—	—	—	—
TMLGA [32]	—	24.54	21.65	16.46	—	—	—	—
VSLNet [26]	—	29.61	24.27	20.03	—	—	—	—
MS-2D-TAN	48.66	41.96	33.59	22.14	75.96	64.93	53.44	36.12
MS-2D-TAN*	53.24	45.96	36.59	24.79	78.73	68.53	57.99	37.94

large improvements compared to the methods that only consider global or local context. The closest work, MAN [42], utilizes GCN to model the relations among the moments and outperforms its previous methods on Charades-STA with VGG features. The experiment results demonstrate the effectiveness of the moment-level relation modeling. However, the nodes in GCN are permutation invariant, thus they ignore the temporal ordering of different moments. In contrast, our MS-2D-TAN models the temporal relations through a 2D convolution network and it can naturally learn the temporal ordering through its convolution kernel. From Table 1, we can observe that MS-2D-TAN surpasses MAN in all evaluation metrics.

Moreover, we compare our approach with clip-based methods (including anchor-based methods, anchor-free methods and RL-based methods). Previous anchor-based methods include TGN [23], CMIN [24], CBP [45] and SCDM [25]. Due to the involvement of more context information during prediction, the anchor-based approaches perform better than the global/local context approaches, however, inferior to our proposed MS-2D-TAN method. Anchor-based approaches implicitly learn the moment context information through a recurrent memory module or a 1D convolution network, while our MS-2D-TAN explicitly exploits the long range context information via the sparse multi-scale 2D temporal map. It further verifies the effectiveness of our model in high quality moment localization. We also compare our method with the anchor-free methods [19], [26]–[35], reinforcement learning based methods [36]–[39] and cascade refinement methods [40],

[41]. Noted that anchor free methods and reinforcement learning methods do not rely on predefined anchors and directly predict the start and end time stamp. Therefore, their upper bounds under any IoU values should all be 100%, which is not directly comparable with our methods. Even with lower upper bound, our model still achieves superior performance compared to theirs. Compared with cascade refinement methods, our method predicts moments in one stage. While it is simple in design, it still outperforms these two methods. This demonstrates the significance of introducing moment context information. We also find that the performance improvement in TACoS dataset is larger than the other two datasets. This also verifies that our MS-2D-TAN is more efficient in modeling long videos.

4.5 Ablation Study

We conduct all our experiments with C3D features in our ablation study. In this section, we first verify the efficiency of our MS-2D-TAN from both theoretical proof and experiments. We then verify the effectiveness of the sparse multi-scale map and gated convolution of our model in Table 4. We also evaluate how does our MS-2D-TAN perform under different types of queries in Table 5, and compare the sparse multi-scale map and the dense single-scale map under the same receptive field in Table 6. Finally, we conduct ablation studies on the hyperparameters of our model in Table 7, i.e., the kernel and layer settings, the number of anchors, window size and the number of hidden states. We also compare the number of candidate moments with previous methods.

In Table 4 - 7, Pool and Conv in the “Feat” column indicate whether using max pooling or stacked convolutions for moment feature extraction. DS, SS and MS in the “Map” column represent the dense single-scale map, sparse single-scale map and the sparse multi-scale map. C, G and X in the “TAN” column represent 2D-TAN with convolution layers, 2D-TAN with gated convolution layers and without using 2D-TAN. H , N , and K represent the size of each hidden layer, sliding window size, and the number of scales. A , κ , and L represent the number of anchors, kernel size, and the number of layers of gated convolutions at each scale, respectively.

4.5.1 Memory Usage and Speed

In this section, we compare the speed and memory cost of the dense single-scale map, the sparse single-scale map, the sparse multi-scale map, as well as other baselines. In the moment feature extraction module, the total number of moment candidates for dense single-scale map is $\sum_{i=1}^N i = O(N^2)$. Meanwhile, the total number of moment candidates for sparse single-scale map is

$$\begin{aligned}
 & A \cdot N + A \cdot \frac{N}{2} + \dots + A \cdot \frac{N}{2^{K-1}} \\
 &= (2 - \frac{1}{2^{K-1}})AN \\
 &= O(N),
 \end{aligned} \tag{6}$$

3. We use the finetuned hyperparameters ($H = 512$, $N = 512$, $A = 8$, $\kappa = 9$, and $L = 2$) for the ablation study on TACoS in Table 4 and Table 5 while others follow the default setting ($H = 512$, $N = 64$, $A = 16$, $\kappa = 17$, and $L = 2$).

TABLE 4: Ablation study on the feature extraction method, the type of 2D temporal map and 2D-TAN.

Row#	Feat	Map	TAN	Charades-STA				ActivityNet Captions						TACoS							
				Rank1@		Rank5@		Rank1@			Rank5@			Rank1@				Rank5@			
				0.5	0.7	0.5	0.7	0.3	0.5	0.7	0.3	0.5	0.7	0.1	0.3	0.5	0.7	0.1	0.3	0.5	0.7
1	Conv	DS	C	39.49	20.83	80.59	50.11	60.29	45.82	29.08	85.94	77.08	61.39	41.74	32.74	23.09	14.10	73.08	56.89	43.36	25.04
2	Conv	SS	C	39.95	21.83	80.97	54.09	60.10	44.45	26.94	86.84	78.36	59.44	44.89	37.19	28.57	18.30	73.11	62.26	49.86	30.84
3	Conv	MS	C	40.97	22.55	80.91	46.48	60.75	45.37	28.25	87.05	78.89	60.67	51.64	44.61	35.29	23.09	79.23	68.63	56.04	35.57
4	Conv	MS	G	41.10	23.25	81.53	48.55	61.04	46.16	29.21	87.30	78.80	60.85	52.39	45.61	35.77	23.44	79.26	69.11	57.31	36.09
5	Pool	MS	G	42.42	25.11	84.38	54.49	60.51	43.59	25.58	87.51	78.72	60.16	52.84	45.34	36.34	24.47	78.58	68.06	57.56	37.54




Fig. 4: Comparisons on inference speed and memory cost of three types of feature map. N is the number of video clips. **Left:** inference speed. **Right:** memory cost.




Fig. 5: Comparisons on inference speed and memory cost with other baselines on TACoS. **Left:** inference speed. **Right:** memory cost. MS-2D-TAN (small) is a MS-2D-TAN model with small size, where its hidden units are set to $H = 64$.

where the number of video clips N is much larger than the number of anchors A and the number of scales K in most cases. Therefore, we can reduce the computational cost of the moment feature extraction module from $O(N^2)$ to $O(N)$ by using sparse sampling.

However, only reducing the computational cost of moment feature extraction is not enough. In our preliminary work [21], we reconstruct a sparse single-scale map with a size of N^2 and feed it to 2D-TAN. Therefore, 2D-TAN still requires time and memory complexities of $O(N^2)$. Our proposed multi-scale extension can tackle this problem by separating it into multiple feature maps of size $A \cdot N$, $A \cdot \frac{N}{2}$, ..., $A \cdot \frac{N}{2^{K-1}}$, thus reducing the time and memory complexities of 2D-TAN from $O(N^2)$ to $O(N)$ [4].

We measure the computational cost of the three types of maps through experiments [5]. As shown in Fig. 4, the time and memory usage of the dense single-scale map

4. Since both the dilation and stride are equal across different scales in our implementation, we reformat the k -th sparse maps ($k > 0$) to dense maps with a stride of 2^k and do convolution operation with a stride of 1 and a dilation of 1. This is equal to the previous, however, using less memory cost.

5. Since the CuDNN package optimizes memory usage at run time, we turn it off for this comparison.

increases quadratically as the sequences become longer. The sparse single-scale map is better than dense single-scale map, however, its time and memory usage still increase quadratically. In contrast, the time and the memory usage of the sparse multi-scale map scales linearly with the sequence length. From both the theoretical and experimental results, we find that the sparse multi-scale modeling has faster speed and lower memory cost compared to the dense single-scale modeling and sparse single-scale modeling.

We also compare the speed and memory cost of MS-2D-TAN with several recent approaches, including SCDM [25], DRN [29], TGN [23] and CBP [45]. From Fig 5, we can observe that MS-2D-TAN* runs faster with less memory usage and achieves better performance than the aforementioned baselines. Further increasing the hidden state size of MS-2D-TAN achieves better performance (MS-2D-TAN *v.s.* MS-2D-TAN (Small)) but sacrifices speed and memory usage.

4.5.2 Effectiveness of Sparse Multi-Scale Map

In this section, we verify the effectiveness of our sparse multi-scale map. As shown in Table 4, comparing with the dense single-scale map (row 1), the sparse single-scale map (row 2) achieves competitive results using less number of moment candidates. [6] The sparse multi-scale map (row 4) further reduces the time and memory cost of the sparse single-scale map, while achieving better performance on TACoS and competitive results on Charades-STA and ActivityNet Captions. The reason is that the videos in TACoS have longer average duration than the other datasets, as shown in Fig. 3. The dense single-scale map and the sparse single-scale map cannot handle long videos well, since the receptive field is limited and context information from long range is missed. In contrast, the sparse multi-scale map has larger receptive field and involves more long range context and therefore outperforms the previous in TACoS dataset.

4.5.3 Effectiveness of Gated Convolution

In this section, we verify the effectiveness of the gated convolution in Table 4. Comparing row 3 and row 4, we find applying an additional gate on the convolution can improve the performance on TACoS dataset. The better performance of gated convolution is also consistent with its effectiveness in image inpainting [52]. The gated convolution is more flexible to adjust the weights of moments among different scales compared to the conventional convolution.

6. In the three datasets (see Fig. 3), most of the target moments are short in time. Our sparse sampling strategy retains short moment candidates while discarding most long candidates, which leverages the distribution prior of target moments. Therefore, it obtains better performance in most cases. (row 1 *v.s.* row 2)

TABLE 5: Performance comparison on the type of query on TACoS.

Row	Feat	Map	TAN	Simple Queries								Complex Queries							
				Rank1@				Rank5@				Rank1@				Rank5@			
				0.1	0.3	0.5	0.7	0.1	0.3	0.5	0.7	0.1	0.3	0.5	0.7	0.1	0.3	0.5	0.7
1	Conv	DS	X	29.53	22.49	15.09	5.41	63.81	51.20	35.77	13.75	28.92	16.06	8.43	2.81	68.67	49.80	28.92	8.84
2	Conv	DS	C	41.95	33.40	23.69	14.74	73.40	57.28	43.66	25.56	38.55	23.29	13.25	4.82	68.67	52.21	39.36	19.28
3	Conv	SS	C	45.52	38.09	29.37	19.00	73.32	62.69	50.67	31.45	36.95	24.90	16.87	7.63	71.49	55.82	41.37	19.68
4	Conv	MS	C	52.16	45.52	36.17	23.91	79.72	69.43	57.14	36.51	42.57	30.12	22.09	10.04	73.49	57.43	40.16	20.08
5	Conv	MS	G	53.30	46.46	36.73	24.15	79.98	69.96	58.18	36.91	38.15	32.13	21.69	11.24	72.29	57.43	44.98	24.10
6	Pool	MS	G	53.65	46.43	37.26	25.03	78.89	68.60	58.50	38.14	40.96	30.92	22.49	16.06	75.10	60.64	45.78	29.32

TABLE 6: Ablation study on receptive field size.

Row#	Receptive Field Size	Feat	Map	TAN	Hyperparameters							Rank1@				Rank5@			
					H	N	K	A	κ	L		0.1	0.3	0.5	0.7	0.1	0.3	0.5	0.7
1	17 × 17	Conv	DS	C	512	512	1	128	9	2		41.74	32.74	23.09	14.10	73.08	56.89	43.36	25.04
2							2	64	5	2		38.72	30.97	21.79	13.15	70.48	56.31	42.64	25.67
3	33 × 33	Conv	DS	C	512	512	1	128	17	2		43.11	34.49	25.79	16.65	70.86	58.21	45.96	29.47
4							3	32	5	2		38.84	31.67	22.27	14.72	72.86	58.89	43.84	25.69
5	65 × 65	Conv	DS	C	512	512	1	128	33	2		—	—	—	—	—	—	—	—
6							4	16	5	2		44.29	35.64	26.72	16.97	72.98	60.36	47.54	29.07
7	129 × 129	Conv	DS	C	512	512	1	128	65	2		—	—	—	—	—	—	—	—
8							5	8	5	2		45.89	39.29	31.32	20.24	73.36	61.83	50.66	32.37

4.5.4 Max Pooling v.s. Stacked Convolution

In this section, we compare the usage of max pooling and stacked convolution for moment feature extraction. As shown in Table 4, we observe that using max pooling has similar performance with stacked convolutions on ActivityNet Captions and TACoS, while getting better performance on Charades-STA. The max-pooling operation is fast in calculation, since it does not contain any parameters, thus it is suitable for computation intensive applications.

4.5.5 Performance on Different Types of Query

In this section, we investigate how our proposed model contributes to different types of queries. Following previous works [65], [66], we define complex queries as sentences that include temporal keywords, such as “before”, “while”, “then”, “after”, “continue” and “again”. All the rest sentences are regarded as simple queries. This rule separates sentences into 3,752 simple queries and 249 complex queries on TACoS. The results are shown in Table 5. Comparing row 1-2, we can observe that using 2D-TAN benefits both simple queries and complex queries. Comparing the row 2-4, using our multi-scale extension can further improve the performance, where the largest improvement comes from simple queries. Comparing the row 4-6, we can observe that the gated convolution can further improve MS-2D-TAN on simple queries and maintain comparable performance on complex queries. Comparing row 5-6, using max pooling achieves better performance on complex queries and similar performance on simple queries. These experiments verify that our multi-scale extension can improve localizing simple queries as well as complex queries with temporal relations. Besides, tuning on the feature extraction method and the convolution type of 2D-TAN may obtain better performance on a specific dataset.

4.5.6 Effectiveness of Receptive Field

To investigate whether the model improvements comes from the large receptive field or the multi-scale context

modeling, we conduct an ablation study in Table 6. To increase receptive fields, the dense single-scale map uses larger kernel size κ , while sparse single-scale map uses larger number of scales K . As shown in row 1-4, within the same receptive field, using a dense single-scale map is slightly better than using a sparse multi-scale map. However, when the receptive field is large enough (row 3 and row 4), the dense single-scale map is not feasible since the time and memory cost increase a lot. In contrast, the sparse multi-scale map can further benefit from the larger receptive fields and achieves a better performance.

4.5.7 Number of Scales

In this section, we evaluate the performance under different number of scales K . We vary the number of scales K from 1 to 7 in our MS-2D-TAN model. The results are shown in Table 7 (row 17 – 23). We observe that, increasing K from 1 to 5 brings improvements (25.49 v.s. 30.42 v.s. 34.09 v.s. 35.04 v.s. 35.77 in Rank5@0.5). This observation is also consistent with the theoretical upper bound, as listed in Table 7 (Row 3 – 7). Here, the upper bound represents the performance of an ideal model that can provide a correct prediction on all the sampled video clips. The upper bound is smaller than 100% since the sampling of video moments may not cover all possible results. We also observe that further increasing the scale K from 5 to 7 does not benefit the overall performance. Since the upper bounds (Row 7 – 9) are already saturated, further increasing moment candidates may introduce additional data imbalance problem.

4.5.8 Number of Anchors

In this section, we evaluate the performance under different number of anchors A . As mentioned in Sec. 4.5.7, the upper bounds are vital to the overall performance. However, it is not the only factor. With the same upper bounds (row 2, 7, 10), we also vary the anchor size and the number of scales, as shown in Table 4 (row 11, 12, 14 and 21). From the experiments, we find that using $K = 5$ and $A = 8$ (row 21) achieves the best performance in the TACoS evaluation

TABLE 7: Ablation study on hyperparameters.

Row#	Method	Hyperparameters						Rank1@				Rank5@				#Params ($\times 10^6$)
		H	N	K	A	κ	L	0.1	0.3	0.5	0.7	0.1	0.3	0.5	0.7	
1	Upper Bound	—	—	1	128	—	—	100.00	99.40	98.95	96.38	100.00	99.40	98.95	96.38	—
2		—	—	4	16	—	—	100.00	99.40	98.95	96.38	100.00	99.40	98.95	96.38	—
3		—	—	1	8	—	—	91.90	68.53	54.16	41.61	91.90	68.53	54.16	41.61	—
4		—	—	2	8	—	—	96.68	85.80	72.81	64.31	96.68	85.80	72.81	64.31	—
5		—	—	3	8	—	—	99.23	94.65	89.35	80.93	99.23	94.65	89.35	80.93	—
6		—	—	4	8	—	—	100.00	98.10	95.65	92.58	100.00	98.10	95.65	92.58	—
7		—	—	5	8	—	—	100.00	99.40	98.95	96.38	100.00	99.40	98.95	96.38	—
8		—	—	6	8	—	—	100.00	100.00	99.75	98.20	100.00	100.00	99.75	98.20	—
9		—	—	7	8	—	—	100.00	100.00	100.00	98.90	100.00	100.00	100.00	98.90	—
10		—	—	6	4	—	—	100.00	99.40	98.95	96.35	100.00	99.40	98.95	96.35	—
11	MS-2D-TAN	512	512	4	16	9	2	47.79	39.74	30.64	20.32	76.33	64.53	52.11	33.57	19.87
12		512	512	6	4	9	1	45.04	37.24	27.54	16.72	69.26	58.84	48.26	33.84	14.30
13		512	512	5	8	1	1	21.04	12.87	6.17	1.87	55.59	39.02	21.94	9.32	21.64
14		512	512	5	8	9	1	49.81	41.71	33.07	20.89	76.03	66.63	55.39	36.22	12.67
15		512	512	5	8	5	4	49.54	42.59	33.49	23.64	76.03	64.68	53.16	36.64	15.17
16		512	512	5	8	17	1	50.61	43.14	35.17	24.32	76.31	66.33	56.54	39.79	39.99
17		512	512	1	8	9	2	40.44	32.52	25.49	17.70	59.39	45.69	35.14	23.84	53.42
18		512	512	2	8	9	2	43.14	36.84	30.42	22.44	67.78	55.49	44.26	32.37	98.33
19		512	512	3	8	9	2	48.81	42.29	34.09	23.64	73.21	63.01	52.39	34.99	143.25
20		512	512	4	8	9	2	50.44	44.04	35.04	23.37	74.33	64.23	53.16	35.44	188.16
21		512	512	5	8	9	2	52.39	45.61	35.77	23.44	79.26	69.11	57.31	36.09	233.08
22		512	512	6	8	9	2	51.14	42.69	32.64	21.04	77.93	65.93	54.06	36.52	278.00
23		512	512	7	8	9	2	49.86	42.24	32.97	21.47	80.03	67.26	54.79	35.02	322.92
24		512	256	5	8	9	2	50.69	43.96	35.34	24.34	77.61	66.36	54.81	37.07	233.08
25		512	128	5	8	9	2	48.56	41.64	33.34	22.19	76.36	65.83	53.79	35.62	233.08
26		128	512	3	8	9	2	44.64	38.24	31.54	22.32	69.81	58.96	49.54	33.87	9.52
27		162	512	3	8	9	2	44.21	38.72	31.22	22.44	72.26	62.88	51.81	35.49	14.99
28	CBP	—	—	1	32	—	—	<i>27.31</i>	<i>24.79</i>	<i>19.10</i>	—	<i>43.64</i>	<i>37.40</i>	<i>25.59</i>	—	15.34

metrics. Using more maps with fewer anchors for each map may not provide sufficient context information at each scale (row 12 *v.s.* row 14). Meanwhile, using less maps with more anchors at each map involves less long range context at each scale due to the smaller receptive field. (row 11 *v.s.* row 21).

4.5.9 Kernel and Layer Setting

In this section, we evaluate different kernel size κ and number of layers L . With the same number of scales ($S = 5$), we vary the kernel size and layers at each scale, as shown in Table 7 (row 13 – 16 and row 21). There are several observations. First, we compare the model with $\kappa = 1, L = 1$ (row 13) and $\kappa = 9, L = 1$ (row 14). Row 13 predicts scores for each moment independently and row 14 jointly considers other neighboring moments in the 2D temporal map. The significant improvement from these two rows verify the importance of modeling context information from adjacent moments. Further enlarging the receptive field can achieve better performance (row 15, 16, and 21). Comparing row 15, 16 and 21 with the same receptive field size, we observe that the performance benefits more from larger kernel with fewer layers (row 16 and row 21) compared to small kernel with more layers (row 15).

4.5.10 Effectiveness of Window Size

In this section, we compare the performance with different sliding window sizes (N) during training in Table 7 (row 24, 25 and 21). We observe that larger window size can benefit the overall performance, since more context moments are observed during training. By using 512 clips extracted by C3D features (row 21), we can cover approximate 96.42% video length in average. And therefore it achieves the best

performance compared with the row 25 (covering 87.57%) and row 24 (covering 67.47%).

4.5.11 Effectiveness of Hidden State Size

In this section, we evaluate the performance under different hidden state size H , as shown in (Row 26, 27, and 21). We observe that with larger number of hidden size, our model can achieve better performance. Even with about $\frac{2}{3}$ parameters of the previous best approach CBP [45], our model can achieve superior performance (Row 27 *v.s.* 28).

4.5.12 Number of Candidate Moments

Proposal based methods can benefit from a larger number of moments. In this section, we verify whether the superior performance of our proposed MS-2D-TAN comes from the larger number of moments. The number of moment candidates is a vital factor in moment localization models. We first tune this factor in our MS-2D-TAN approach, and show its impacts on final performance. Next, we compare the previous best approach CBP [45] with respect to this factor. The upper bound performance of proposal based methods is decided by the selection of moment candidates. Therefore, we evaluate our MS-2D-TAN with the same moment selection of previous state-of-the-art proposal based methods. For a fair comparison, we compare CBP with our MS-2D-TAN on TACoS. In CBP (Row 27), each clip corresponds to 32 anchors, while in the MS-2D-TAN with $A = 8$ and $K = 3$ (Row 26), the corresponding anchors are a subset of the CBP anchors (16 anchors). With approximately the same number of parameters and half less number of anchors, our model outperforms the previous best approach by a large margin.

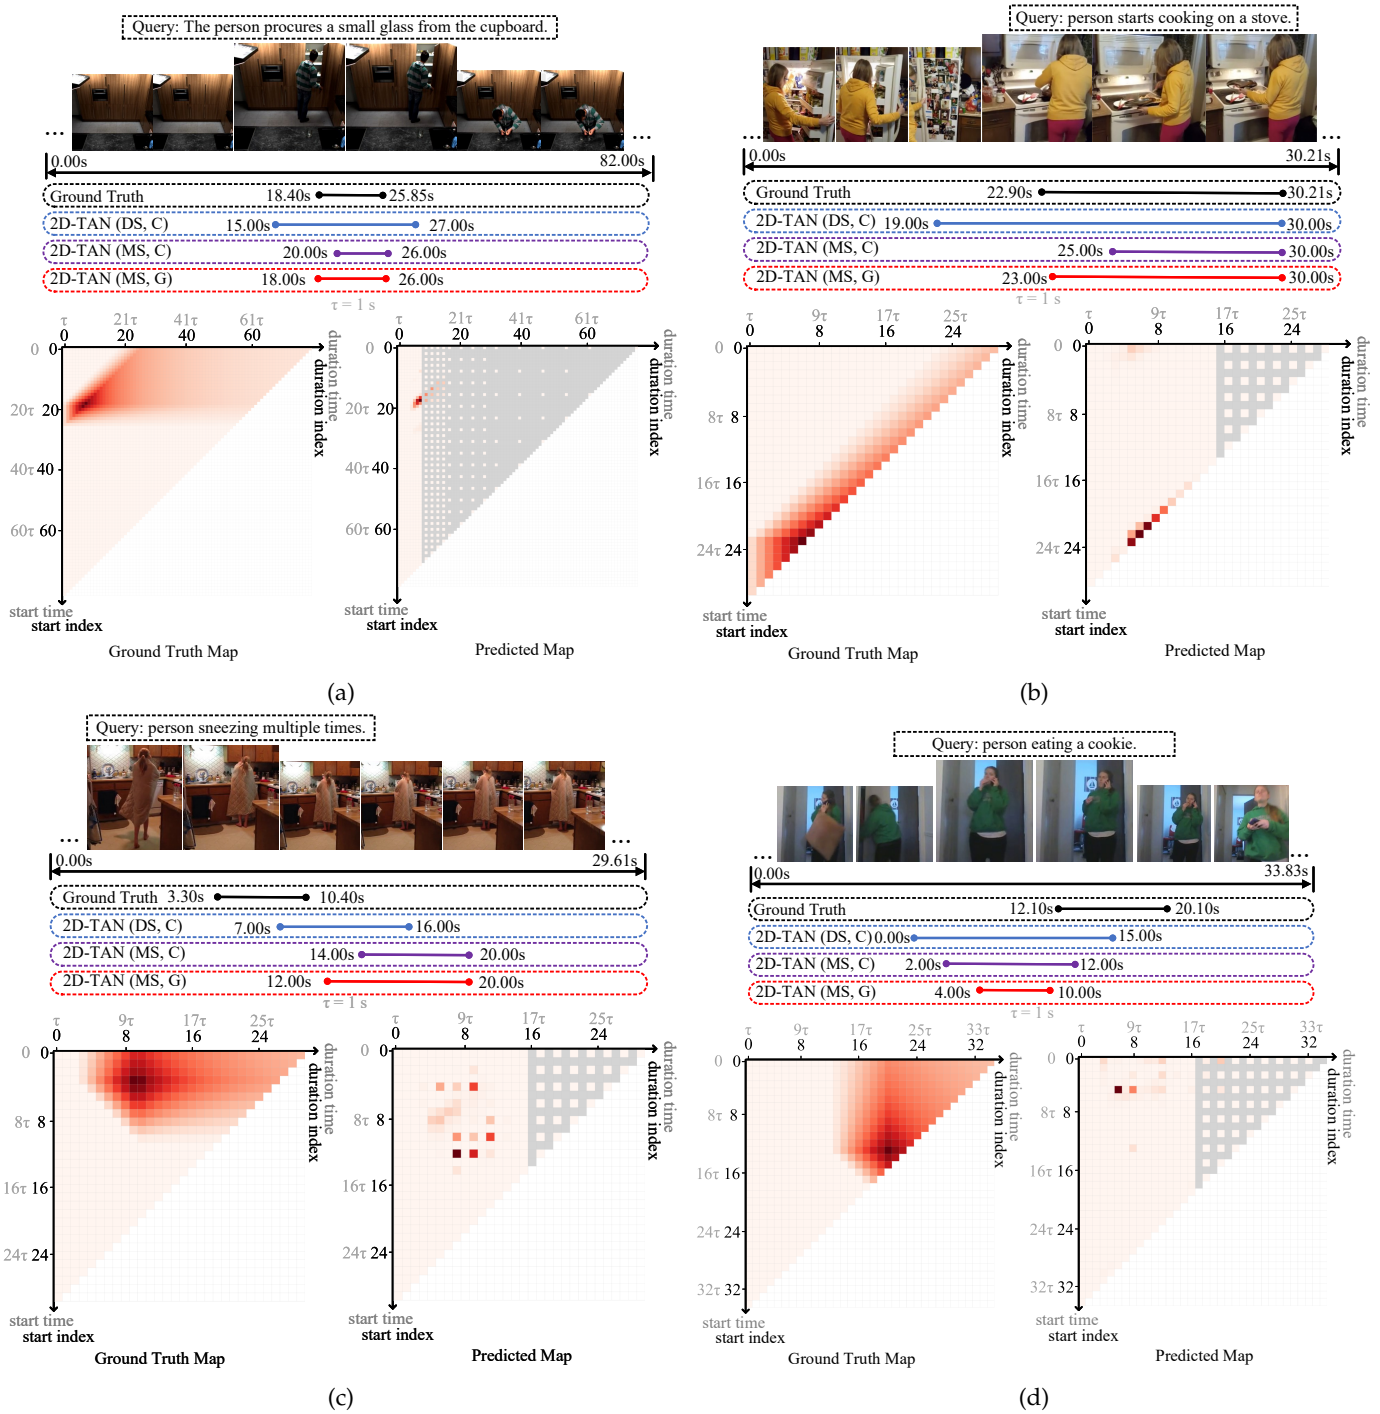


Fig. 6: Prediction examples of our model and the baselines. The map is predicted by our full model.

4.6 Qualitative Analysis

We provide some qualitative examples to validate the effectiveness of our MS-2D-TAN. We first demonstrate two success cases in Fig. 6 (a - b). The sparse multi-scale map exploits richer context information with a larger receptive field compared to the dense single-scale map. By combining with gated convolution, the location of the desired moment can be more precise. We also visualize the predicted scores of our model. Comparing with the ground truth, we can observe that the predicted map can correctly localize the target moment, while the surrounding moments also have

high predicted scores. This demonstrates that our model can learn discriminative features from the neighboring moments.

We further present two failure cases. In Fig. 6 (c), all models failed, since it's hard to localize "sneezing" without hearing the sound. In Fig. 6 (d), all models failed to localize "eating a cookie". Due to the lack of audio modality and object category information, MS-2D-TAN is not able to correctly localize these queries shown in Figure 6. In the future, we would like to explore the multi-modal information and study how to interact multi-modal 2D temporal maps with

sentence queries.

5 CONCLUSION

In this paper, we study the problem of moment localization with natural language and present a novel Multi-Scale 2D Temporal Adjacency Networks (MS-2D-TAN) method. The core idea is to retrieve a moment on the multi-scale two-dimensional temporal map, which considers adjacent moment candidates as the temporal context. MS-2D-TAN is capable of encoding adjacent temporal context, while learning discriminative feature for matching video moments with the sentence query. Our model is simple in design and achieves competitive performance in comparison with other state-of-the-art methods on three benchmark datasets.


6 ACKNOWLEDGEMENT

This work is supported in part by NSF awards IIS-1704337, IIS-1722847, and IIS-1813709, as well as our corporate sponsors.


REFERENCES

- [1] Y. Zhao, Y. Xiong, L. Wang, Z. Wu, X. Tang, and D. Lin, "Temporal action detection with structured segment networks," in *ICCV*, 2017.
- [2] P. Chen, C. Gan, G. Shen, W. Huang, R. Zeng, and M. Tan, "Relation attention for temporal action localization," *TMM*, 2019.
- [3] R. Zeng, C. Gan, P. Chen, W. Huang, Q. Wu, and M. Tan, "Breaking winner-takes-all: Iterative-winners-out networks for weakly supervised temporal action localization," *TIP*, 2019.
- [4] M. Xu, C. Zhao, D. S. Rojas, A. Thabet, and B. Ghanem, "G-tad: Sub-graph localization for temporal action detection," in *CVPR*, 2020.
- [5] C. Lin, C. Xu, D. Luo, Y. Wang, Y. Tai, C. Wang, J. Li, F. Huang, and Y. Fu, "Learning salient boundary feature for anchor-free temporal action localization," in *CVPR*, 2021.
- [6] M. Hasan, J. Choi, J. Neumann, A. K. Roy-Chowdhury, and L. S. Davis, "Learning temporal regularity in video sequences," in *CVPR*, 2016.
- [7] Y. Song, J. Vallmitjana, A. Stent, and A. Jaimes, "TVSum: Summarizing web videos using titles," in *CVPR*, 2015.
- [8] W.-S. Chu, Y. Song, and A. Jaimes, "Video co-summarization: Video summarization by visual co-occurrence," in *CVPR*, 2015.
- [9] J. Gao, C. Sun, Z. Yang, and R. Nevatia, "TALL: Temporal activity localization via language query," in *ICCV*, 2017.
- [10] L. A. Hendricks, O. Wang, E. Shechtman, J. Sivic, T. Darrell, and B. Russell, "Localizing moments in video with natural language," in *ICCV*, 2017.
- [11] V. Escorcia, M. Soldan, J. Sivic, B. Ghanem, and B. Russell, "Temporal localization of moments in video collections with natural language," *arXiv preprint arXiv:1907.12763*, 2019.
- [12] Y. Zhao, Z. Zhao, Z. Zhang, and Z. Lin, "Cascaded prediction network via segment tree for temporal video grounding," in *CVPR*, 2021.
- [13] Z. Lin, Z. Zhao, Z. Zhang, Z. Zhang, and D. Cai, "Moment retrieval via cross-modal interaction networks with query reconstruction," *TIP*, 2020.
- [14] J. Lei, L. Yu, M. Bansal, and T. L. Berg, "TVQA: Localized, compositional video question answering," in *EMNLP*, 2018.
- [15] D. Shao, Y. Xiong, Y. Zhao, Q. Huang, Y. Qiao, and D. Lin, "Find and focus: Retrieve and localize video events with natural language queries," in *ECCV*, 2018.
- [16] X. Duan, W. Huang, C. Gan, J. Wang, W. Zhu, and J. Huang, "Weakly supervised dense event captioning in videos," in *NeurIPS*, 2018.
- [17] S. Gella, M. Lewis, and M. Rohrbach, "A dataset for telling the stories of social media videos," in *EMNLP*, 2018.
- [18] R. Ge, J. Gao, K. Chen, and R. Nevatia, "MAC: Mining activity concepts for language-based temporal localization," in *WACV*, 2019.
- [19] M. Liu, X. Wang, L. Nie, X. He, B. Chen, and T.-S. Chua, "Attentive moment retrieval in videos," in *SIGIR*, 2018.
- [20] X. Song and Y. Han, "VAL: Visual-attention action localizer," in *PCM*, 2018.
- [21] S. Zhang, H. Peng, J. Fu, and J. Luo, "Learning 2d temporal adjacent networks for moment localization with natural language," in *AAAI*, 2020.
- [22] T. Lin, X. Zhao, and Z. Shou, "Single shot temporal action detection," in *ACMMM*, 2017.
- [23] J. Chen, X. Chen, L. Ma, Z. Jie, and T.-S. Chua, "Temporally grounding natural sentence in video," in *EMNLP*, 2018.
- [24] Z. Zhang, Z. Lin, Z. Zhao, and Z. Xiao, "Cross-modal interaction networks for query-based moment retrieval in videos," in *SIGIR*, 2019.
- [25] Y. Yuan, L. Ma, J. Wang, W. Liu, and W. Zhu, "Semantic conditioned dynamic modulation for temporal sentence grounding in videos," in *NeurIPS*, 2019.
- [26] H. Zhang, A. Sun, W. Jing, and J. T. Zhou, "Span-based localizing network for natural language video localization," in *ACL*, 2020.
- [27] Y. Yuan, T. Mei, and W. Zhu, "To find where you talk: Temporal sentence localization in video with attention based location regression," in *AAAI*, 2019.
- [28] S. Ghosh, A. Agarwal, Z. Parekh, and A. Hauptmann, "ExCL: Extractive clip localization using natural language descriptions," in *NAACL*, 2019.
- [29] R. Zeng, H. Xu, W. Huang, P. Chen, M. Tan, and C. Gan, "Dense regression network for video grounding," *CVPR*, 2020.
- [30] M. Liu, X. Wang, L. Nie, Q. Tian, B. Chen, and T.-S. Chua, "Cross-modal moment localization in videos," in *ACMMM*, 2018.
- [31] J. Mun, M. Cho, and B. Han, "Local-global video-text interactions for temporal grounding," in *CVPR*, 2020.
- [32] C. Rodriguez, E. Marrese-Taylor, F. S. Saleh, H. Li, and S. Gould, "Proposal-free temporal moment localization of a natural-language query in video using guided attention," in *WACV*, March 2020.
- [33] L. Chen, C. Lu, S. Tang, J. Xiao, D. Zhang, C. Tan, and X. Li, "Rethinking the bottom-up framework for query-based video localization," in *AAAI*, 2020.
- [34] C. Lu, L. Chen, C. Tan, X. Li, and J. Xiao, "Debug: A dense bottom-up grounding approach for natural language video localization," in *EMNLP*, 2019.
- [35] B. Jiang, X. Huang, C. Yang, and J. Yuan, "Cross-modal video moment retrieval with spatial and language-temporal attention," in *ICMR*, 2019.
- [36] D. He, X. Zhao, J. Huang, F. Li, X. Liu, and S. Wen, "Read, watch, and move: Reinforcement learning for temporally grounding natural language descriptions in videos," in *AAAI*, 2019.
- [37] W. Wang, Y. Huang, and L. Wang, "Language-driven temporal activity localization: A semantic matching reinforcement learning model," in *CVPR*, 2019.
- [38] M. Hahn, A. Kadav, J. M. Rehg, and H. P. Graf, "Tripping through time: Efficient localization of activities in videos," in *CVPR Workshop*, 2019.
- [39] J. Wu, G. Li, S. Liu, and L. Lin, "Tree-structured policy based progressive reinforcement learning for temporally language grounding in video," in *AAAI*, 2020.
- [40] H. Xu, K. He, B. A. Plummer, L. Sigal, S. Sclaroff, and K. Saenko, "Multilevel language and vision integration for text-to-clip retrieval," in *AAAI*, 2019.
- [41] S. Chen and Y.-G. Jiang, "Semantic proposal for activity localization in videos via sentence query," in *AAAI*, 2019.
- [42] D. Zhang, X. Dai, X. Wang, Y.-F. Wang, and L. S. Davis, "MAN: Moment alignment network for natural language moment retrieval via iterative graph adjustment," in *CVPR*, 2019.
- [43] N. Dehmamy, A.-L. Barabási, and R. Yu, "Understanding the representation power of graph neural networks in learning graph topology," in *NeurIPS*, 2019.
- [44] B. Liu, S. Yeung, E. Chou, D.-A. Huang, L. Fei-Fei, and J. Carlos Nibbles, "Temporal modular networks for retrieving complex compositional activities in videos," in *ECCV*, 2018.
- [45] J. Wang, L. Ma, and W. Jiang, "Temporally grounding language queries in videos by contextual boundary-aware prediction," in *AAAI*, 2020.
- [46] I. Beltagy, M. E. Peters, and A. Cohan, "Longformer: The long-document transformer," *arXiv:2004.05150*, 2020.
- [47] S. Sukhbaatar, E. Grave, P. Bojanowski, and A. Joulin, "Adaptive attention span in transformers," *ACL*, 2019.

- [48] T. Lin, X. Liu, X. Li, E. Ding, and S. Wen, "BMN: Boundary-matching network for temporal action proposal generation," in *ICCV*, 2019.
- [49] C. Lin, J. Li, Y. Wang, Y. Tai, D. Luo, Z. Cui, C. Wang, J. Li, F. Huang, and R. Ji, "Fast learning of temporal action proposal via dense boundary generator," in *AAAI*, 2020.
- [50] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *EMNLP*, 2014.
- [51] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [52] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Free-form image inpainting with gated convolution," in *ICCV*, 2019.
- [53] R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles, "Dense-captioning events in videos," in *ICCV*, 2017.
- [54] M. Regneri, M. Rohrbach, D. Wetzels, S. Thater, B. Schiele, and M. Pinkal, "Grounding action descriptions in videos," *TACL*, 2013.
- [55] M. Rohrbach, M. Regneri, M. Andriluka, S. Amin, M. Pinkal, and B. Schiele, "Script data for attribute-based recognition of composite activities," in *ECCV*, 2012.
- [56] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICPR*, 2015.
- [57] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *ICCV*, 2015.
- [58] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *CVPR*, 2017.
- [59] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *CVPR*, 2009.
- [60] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *CVPR*, 2014.
- [61] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta, "Hollywood in homes: Crowdsourcing data collection for activity understanding," in *ECCV*, 2016.
- [62] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICPR*, 2015.
- [63] A. Wu and Y. Han, "Multi-modal circulant fusion for video-to-language and backward," in *IJCAI*, 2018.
- [64] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *NeurIPS*, 2015.
- [65] L. A. Hendricks, O. Wang, E. Shechtman, J. Sivic, T. Darrell, and B. Russell, "Localizing moments in video with temporal language," in *EMNLP*, 2018.
- [66] S. Zhang, J. Su, and J. Luo, "Exploiting temporal relationships in video moment localization with natural language," in *ACMMM*, 2019.




Songyang Zhang received the BS degree from Southeast University, China, in 2015 and MS degree from Zhejiang University, China, in 2018. He is currently working toward the PhD degree at University of Rochester. His research interest includes moment localization with natural language, unsupervised grammar induction, skeleton-based action recognition, etc.



Houwen Peng is a researcher working on computer vision and deep learning at Microsoft Research as of 2018. Before that he was a senior engineer at Qualcomm AI Research. He received Ph.D. from NLPRI, Institution of Automation, Chinese Academy of Sciences in 2016. From 2015 to 2016, he worked as a visiting research scholar at Temple University. His research interest includes neural architecture search and design, video object tracking, segmentation and detection, video moment localization, saliency detection, etc.



Jianlong Fu is currently a Senior Research Manager with the Multimedia Search and Mining Group, Microsoft Research Asia (MSRA). He is now leading Multimedia Search and Mining Group, focusing on computer vision, image graphics, vision and language, especially on fine-grained image/video recognition and detection, multimedia content editing, personal media experience of browsing, searching, sharing and storytelling. He has shipped core technologies to a number of Microsoft products, including Windows, Office, Bing Multimedia Search, Azure Media Service, Xiaoice, etc.



Yijuan(Lucy) Lu is a Principal Scientist at Microsoft Azure AI, where she worked on OCR, object detection and video understanding in the recent two years. Prior to joining Microsoft, she was an associate professor in the Department of Computer Science at Texas State University. Her major publications appear in leading publication venues in multimedia and computer vision research. She was the First Place Winner in many challenging retrieval competitions in Eurographics for many years. She received 2015 Texas State Presidential Distinction Award and 2014 College Achievement Award. She also received the Best Paper award from ICME 2013 and ICIMCS 2012. She has obtained many competitive external grants from NSF, US Army, US Department of Defense and Texas Department of Transportation.



Jiebo Luo (S93, M96, SM99, F09) is a Professor of Computer Science at the University of Rochester which he joined in 2011 after a prolific career of fifteen years at Kodak Research Laboratories. He has authored nearly 500 technical papers and holds over 90 U.S. patents. His research interests include computer vision, NLP, machine learning, data mining, computational social science, and digital health. He has been involved in numerous technical conferences, including serving as a program co-chair of ACM Multimedia 2010, IEEE CVPR 2012, ACM ICMR 2016, and IEEE ICIP 2017, as well as a general co-chair of ACM Multimedia 2018. He has served on the editorial boards of the IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), IEEE Transactions on Multimedia (TMM), IEEE Transactions on Circuits and Systems for Video Technology (TCSVT), IEEE Transactions on Big Data (TBD), ACM Transactions on Intelligent Systems and Technology (TIST), Pattern Recognition, Knowledge and Information Systems (KAIS), Machine Vision and Applications, and Journal of Electronic Imaging. He is the current Editor-in-Chief of the IEEE Transactions on Multimedia. Professor Luo is also a Fellow of ACM, AAAI, SPIE, and IAPR.