

Relation-aware Video Reading Comprehension for Temporal Language Grounding

Jialin Gao^{1,2*} Xin Sun^{1,2*} MengMeng Xu³ Xi Zhou^{1,2} Bernard Ghanem³

¹Cooperative Medianet Innovation Center, Shanghai Jiao Tong University

²CloudWalk Technology Co., Ltd, China

³King Abdullah University of Science and Technology

{jialin_gao, huntersx}@sjtu.edu.cn, zhoxi@cloudwalk.cn

{mengmeng.xu,bernard.ghanem}@kaust.edu.sa

Abstract

Temporal language grounding in videos aims to localize the temporal span relevant to the given query sentence. Previous methods treat it either as a boundary regression task or a span extraction task. This paper will formulate temporal language grounding into video reading comprehension and propose a *Relation-aware Network (RaNet)* to address it. This framework aims to select a video moment choice from the predefined answer set with the aid of coarse-and-fine choice-query interaction and choice-choice relation construction. A choice-query interactor is proposed to match the visual and textual information simultaneously in sentence-moment and token-moment levels, leading to a coarse-and-fine cross-modal interaction. Moreover, a novel multi-choice relation constructor is introduced by leveraging graph convolution to capture the dependencies among video moment choices for the best choice selection. Extensive experiments on ActivityNet-Captions, TACoS, and CharadesSTA demonstrate the effectiveness of our solution. Codes will be available at <https://github.com/Huntersxsx/RaNet>.

1 Introduction

Recently, temporal language grounding in videos has become a heated topic in the computer vision, and natural language processing community (Gao et al., 2017; Krishna et al., 2017). This task requires a machine to localize a temporal moment semantically relevant to a given language query, as shown in Fig.1. It has also drawn great attention from industry due to its various applications such as video question answering (Huang et al., 2020; Lei et al., 2018), video content retrieval (Dong et al., 2019; Shao et al., 2018), and human-computer interaction (Zhu et al., 2020), etc.

A straightforward paradigm for this task is the proposing-and-ranking pipelines (Xu et al., 2019;

* Jialin Gao and Xin Sun are co-first authors with equal contributions, supervised by Prof. Xi Zhou in SJTU.

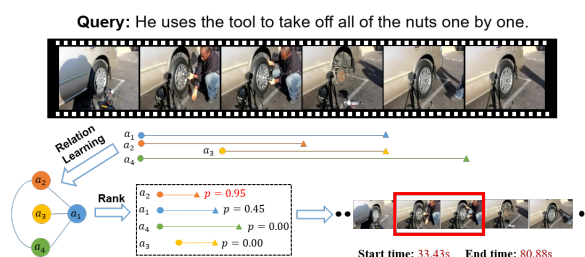


Figure 1: An illustration of temporal language grounding in videos based on the relation-aware network. Given a video and a query sentence, our approach aims to semantically align the query representation with a predefined answer set of video moment candidates (a_1, a_2, a_3 and a_4) and then mine the relationships between them to select the best-matched one.

Zhang et al., 2020b, 2019a). They first generate a number of video moment candidates and then rank them according to moment-query similarities. This requires a solution to achieve two key targets simultaneously, which are (1) *semantic visual-language interaction* and (2) *reliable candidate ranking*. The former ensures a satisfying cross-modal matching between video moments and the query, while the latter guarantees the distinction among candidates. For the first target, some previous works (Yuan et al., 2019; Zhang et al., 2020b; Chen and Jiang, 2019) resort to the visual clues by modeling moment-sentence or snippet-sentence relations. However, they overlook the linguistic clues from token-level, *i.e.*, token-moment relations, which contain fine-grained linguistic information. For the second target, previous solutions (Ge et al., 2019; Liu et al., 2018; Zhang et al., 2020a) generate ranking scores by considering different moment candidates separately or constructing moment-level relations in a simple way (Zhang et al., 2020b). Hence, they neglect the temporal and semantic dependencies among candidates. Without this information, it is difficult for previous approaches to distinguish these visually similar moment candidates correctly.

To this end, we propose a *Relation-aware Network (RaNet)* to address temporal language grounding. In our solution, we formulate this task as Video Reading Comprehension by regarding the video, query, and moment candidates as the text passage, question description, and multi-choice options, respectively. Unlike previous methods, we exploit a coarse-and-fine interaction, which captures not only sentence-moment relations but also token-moment relations. This interaction can allow our model to construct both sentence-aware and token-aware visual representation for each choice, which is expected to distinguish similar candidates in the visual modality. Moreover, we propose to leverage Graph Convolutional Networks (GCN) (Kipf and Welling, 2016) for mining the moment-moment relations between candidate choices based on their coarse-and-fine representations. With information exchange in GCNs, our RaNet can learn discriminative features for correctly ranking candidates regardless of their high relevance in visual content.

Similar to the system of multi-choice reading comprehension, our RaNet consists of five components: a modality-wise encoder for visual and textual encoding, a multi-choice generator for answer set generation, a choice-query interactor for cross-modality interaction, a multi-choice relation constructor for relationships mining and an answer ranker for the best-matched choice selection. Our contributions are summarized as three-fold: (1) We address temporal language grounding by a Relation-aware Network, which formulates this task as a video reading comprehension problem. (2) We exploit the visual and linguistic clues exhaustively, i.e., coarse-and-fine moment-query relations and moment-moment relations, to learn discriminative representations for distinguishing candidates. (3) The proposed RaNet outperforms other state-of-the-art methods on three widely-used challenging benchmarks: TACoS, CharadesSTA and ActivityNet-Captions, where we improve the grounding performance by a great margin (i.e., 33.54% v.s. 25.32% of 2D-TAN on TACoS dataset).

2 Related Work

Temporal Language Grounding. This task was introduced by (Anne Hendricks et al., 2017; Gao et al., 2017) to locate relevant moments given a language query. He et al. (He et al., 2019) and

Wang et al. (Wang et al., 2019) used reinforcement learning to solve this problem. Chen et al. (Chen et al., 2018) and Ghosh et al. (Ghosh et al., 2019) proposed to select the boundary frames based on visual-language interaction. Most of recent works (Xu et al., 2019; Yuan et al., 2019; Zhang et al., 2020b; Chen and Jiang, 2019; Zhang et al., 2019a) adopted the two-step pipeline to solve this problem. **Visual-Language Interaction.** It is vital for this task to semantically match query sentences and video. This cross-modality alignment was usually achieved by attention mechanism (Vaswani et al., 2017) and sequential modeling (Hochreiter and Schmidhuber, 1997; Medsker and Jain, 2001). Xu et al. (Xu et al., 2019) and Liu et al. (Liu et al., 2018) designed soft-attention modules while Hendricks et al. (Anne Hendricks et al., 2017) and Zhang et al. (Zhang et al., 2019b) chose the hard counterpart. Some works (Chen et al., 2018; Ghosh et al., 2019) attempt to use the property of RNN cells and others went beyond it by dynamic filters (Zhang et al., 2019a), Hadamard product (Zhang et al., 2020b), QANet (Lu et al., 2019) and circular matrices (Wu and Han, 2018). However, these alignments neglect the importance of token-aware visual feature in cross-modal correlating and distinguishing the similar candidates. **Machine Reading Comprehension.** Given the reference document or passage, Machine Reading Comprehension (MRC) requires the machine to answer questions about it (Zhang et al., 2020c). There are two types of the existing MRC variations related to the temporal language grounding in videos, i.e., *span extraction* and *multi-choice*. The former (Rajpurkar et al., 2016) extracts spans from the given passage and has been explored in temporal language grounding task by some previous works (Zhang et al., 2020a; Lu et al., 2019; Ghosh et al., 2019). The latter (Lai et al., 2017; Sun et al., 2019) aims to find the only correct option in the given candidate choices based on the given passage. We propose to formulate this task from the perspective of multi-choice reading comprehension. Based on this formulation, we focus on the visual-language alignment in a token-moment level. Compared with query-aware context representation in previous solutions, we aim to construct token-aware visual feature for each choice. Inspired by recent advanced attention module (Gao et al., 2020; Huang et al., 2019), we mine the relations between multi-choices in an effective and efficient way.

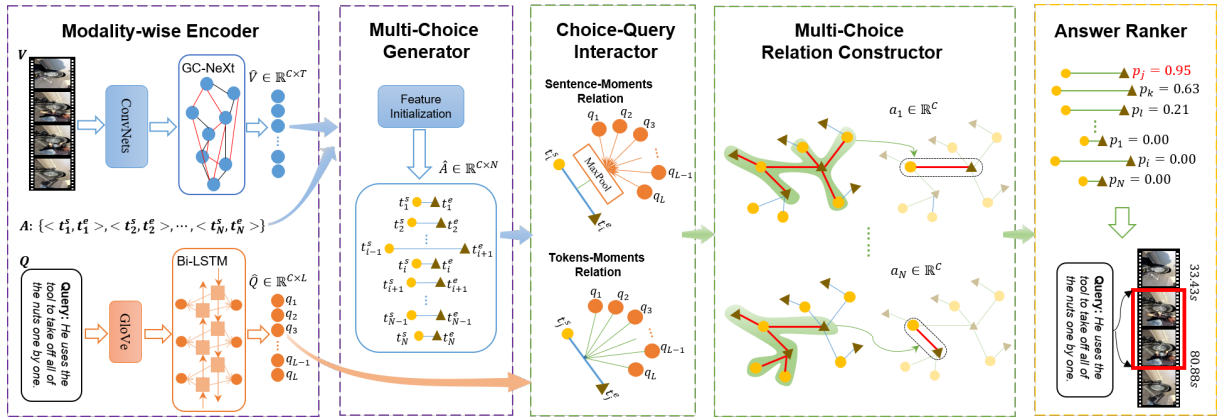


Figure 2: The overview of our proposed RaNet. It consists of modality-wise encoder, multi-choice generator, choice-query interactor, multi-choice relation constructor, and answer ranker. Video and language passages are first embedded in separately branches. Then we initialize the visual representation for each choice $\langle t_i^s, t_i^e \rangle$ from the video stream. Through choice-query interactor, each choice can capture the sentence-aware and token-aware representation from the query. Afterwards, the relation constructor takes advantage of GCNs to model relationships between choices. Finally, the answer ranker evaluates the probability of being selected for each choice based on the exchanged information from the former module.

3 Methodology

In this section, we first describe how to recast the temporal language grounding from the perspective of a multi-choice reading comprehension task, which is solved by the proposed Relation-aware Network (RaNet). Then, we introduce the detailed architecture of the RaNet, consisting of five components as shown in Fig.2. Finally, we illustrate the training and inference of our solution.

3.1 Problem Definition

The goal of this task is to answer where is the semantically corresponding video moment given a language query in an untrimmed video. Referring to the forms of MRC, we treat the video \mathbf{V} as a text passage, the query sentence \mathbf{Q} as a question description and provide a set of video moment candidates as a list of answer options \mathbf{A} . Based on the given triplet $(\mathbf{V}, \mathbf{Q}, \mathbf{A})$, temporal language grounding in videos is equivalent to cross-modal MRC, termed video reading comprehension.

For each query-video pair, we have one natural language sentence and an associated ground-truth video moment with the start g^s and end g^e boundary. Each language sentence is represented by $\mathbf{Q} = \{q_i\}_{i=1}^L$, where L is the number of tokens. The untrimmed video is represented as a sequential snippets $\mathbf{V} = \{v_1, v_2, \dots, v_{n_v}\} \in \mathcal{R}^{n_v \times C}$ by a pretrained video understanding network, such as C3D (Tran et al., 2015), I3D (Carreira and Zisserman, 2017), etc..

In temporal language grounding, the answer should be a consecutive subsequence (namely time span) of the video passage. For any video moment candidate (i, j) , it can be treated as a possible answer if it meets the condition of $0 < i < j < n_v$. Hence, we follow the fixed-interval sampling strategy in previous work (Zhang et al., 2020b) and construct a set of video moment candidates as the answer list $A = \{a_1, \dots, a_N\}$ with N valid candidates. After these notations, we can recast the temporal language grounding task from the perspective of multi-choice reading comprehension as:

$$\arg \max_i P(a_i | (V, Q, A)). \quad (1)$$

However, different from the traditional multi-choice reading comprehension, previous solutions in temporal language grounding also compare their performance in terms of top- K most matching candidates for each query sentence. For fair comparison, it requires our approach to scores K candidate moments $\{(p_i, t_i^s, t_i^e)\}_{i=1}^K$, where p_i, t_i^s, t_i^e represent the probability of selection, the start, end time of answer a_i , respectively. Without additional mention, the video moment and answer/choice are interchangeable in this paper.

3.2 Architecture

As shown in Figure 2, we describe the details of each component in our framework as followings: **Modality-wise Encoder**. This module aims to separately encode the content of language query sen-




Figure 3: (a). Illustration of choice generator and feature initialization. Each block (i, j) is a valid choice when $i < j$, denoted with blue. The Ψ combines the boundary feature for feature initialization of $(1, 3)$, the dark blue square. (b). An example of all graph edges connected to one choice in our Multi-Choice Relation Constructor. Moments with the same start or end index (dark green) are connected with the illustrative choice (red). (c). The information propagation between two unconnected moment choices. For other moments (dark green) that are not connected with target moment (red) but have overlaps, relations can be implicit captured with two loops, namely 2 graph attention layers.

tence and video. Each branch aggregates the intra-modality context for each snippet and token.

· *Video Encoding.* We first apply a simple temporal 1D convolution operation to map the input feature sequence to the desired dimension, which is followed by an average pooling layer to reshape the sequence into a desired length T . To enrich the multi-hop interaction, we use a graph convolution block called GC-NeXt block (Xu et al., 2020), which aggregates the context from both temporal and semantic neighbors of each snippet v_i and has been proved effective in Temporal Action Localization task. Finally, We get the encoded visual feature as $\hat{\mathbf{V}} \in \mathcal{R}^{C \times T}$

· *Language Encoding.* Each word q_i of \mathbf{Q} is represented with the embedding vector from GloVe 6B 300d (Pennington et al., 2014) to get $\mathbf{Q} \in \mathcal{R}^{300 \times L}$. Then we sequentially feed the initialized embeddings into a three-layer Bi-LSTM network to capture semantic information and the temporal context. We take the last layer’s hidden states as the language representation $\hat{\mathbf{Q}} \in \mathcal{R}^{C \times L}$ for cross-modality fusion with video representation $\hat{\mathbf{V}}$. In addition, the effect of different word embeddings, is also compared in 4.4.4.

The encoded visual and textual features can be formulated as follows:

$$\begin{aligned} \hat{\mathbf{V}} &= \text{VisualEncoder}(\mathbf{V}) \\ \hat{\mathbf{Q}} &= \text{LanguageEncoder}(\mathbf{Q}) \end{aligned} \quad (2)$$

Multi-Choice Generator. As shown in Fig.3 (a),

the vertical and horizontal axes represent the start and end index of visual sequence. The blocks in the same row have the same start indices, and those in the same column have the same end indices. The white blocks indicate all the invalid choices in the left-bottom, where the start boundaries exceed the end boundaries. Therefore, we have the multi-choice options as $A = \{(t_i^s, t_i^e)\}_{i=1}^N$. To capture the visual-language interaction, we should initialize the visual feature for the answer set A so that it can be integrated with the textual features from the language encoder. To ensure the boundary-aware ability inspired by (Wang et al., 2020), the initialization method Ψ combines boundary information, i.e., $\hat{v}_{t_i^s}$ and $\hat{v}_{t_i^e}$ in $\hat{\mathbf{V}}$, to construct the moment-level feature representation for each choice a_i . The initialized feature representation can be written as:

$$\mathbf{F}_A = \Psi(\hat{\mathbf{V}}, A), \quad (3)$$

where Ψ is the concatenation of $\hat{v}_{t_i^s}$ and $\hat{v}_{t_i^e}$, A is the answer set and $\mathbf{F}_A \in \mathcal{R}^{C \times N}$. We also explore the effect of different Ψ on grounding performance in 4.4.3.

Choice-Query Interactor. As shown in Figure 2, this module explores the inter-modality context for visual-language interaction. Unlike previous methods (Zhang et al., 2020b,a; Zeng et al., 2020), we propose a *coarse-and-fine* cross-modal interaction. We integrate the initialized features \mathbf{F}_A with the query both in the sentence-level and token-level. The former can be obtained by a simple Hadamard product and an normalization as:

$$\mathbf{F}_1 = \|\varphi(\hat{\mathbf{Q}}) \odot \text{Conv}(\mathbf{F}_A)\|_F, \quad (4)$$

where φ is the aggregation function for a global representation of $\hat{\mathbf{Q}}$ and we set it as the max-pooling, \odot is element-wise multiplication, and $\|\cdot\|_F$ indicates the Frobenius normalization.

To ensure the token-aware visual feature for each choice a_i , we adopt attention mechanism to learn the token-moment relation between each choice and the query. Firstly, we adopt a 1D convolution layer to project the visual and textual features to the common space and then calculate their semantic similarities, which depict the relationships $\mathbf{R} \in \mathcal{R}^{N \times T}$ between N candidates and L tokens. Secondly, we generate query-related feature for each candidate based on the relationships \mathbf{R} . Finally, we integrate these two features of candidates

for token-aware visual representation.

$$\begin{aligned} \mathbf{R} &= \text{Conv}(\mathbf{F}_A)^T \otimes \text{Conv}(\hat{\mathbf{Q}}) \\ \mathbf{F}_2 &= \text{Conv}(\mathbf{F}_A) \odot (\text{Conv}(\hat{\mathbf{Q}}) \otimes \mathbf{R}^T), \end{aligned} \quad (5)$$

where T denotes the matrix transpose, \odot and \otimes are element-wise and matrix multiplications, respectively. We add the sentence-aware \mathbf{F}_1 and token-aware features \mathbf{F}_2 as the output of this module $\hat{\mathbf{F}}_A$. **Multi-choice Relation Constructor.** In order to explore the relation between multi-choices, we propose this module to aggregate the information from the overlapping moment candidates by GCNs. Previous methods MAN (Zhang et al., 2019a), 2D-TAN (Zhang et al., 2020b) also considered moment-wise temporal relations, while both of them have two drawbacks: expensive computations and the noise from unnecessary relations. Inspired by CCNet (Huang et al., 2019), which proposed a sparsely-connected graph attention module to collect contextual information in horizontal and vertical directions, we propose a Graph Attention layer (GAT) to constrain the relation between moment candidates that have high temporal overlaps with each other.

Concretely, we take each answer candidate $a_i = (t_i^s, t_i^e)$ as a graph node, and there is a graph edge connecting two candidate choices a_i, a_j if they share the same start/end time spot, e.g., $t_i^s = t_j^s$ or $t_i^e = t_j^e$. An example is shown in Figure 3 (b), where neighbors of the target moment choice (the red one) is denoted as dark green in a criss-cross shape. As shown in Figure 3 (c), our model is also able to achieve the information propagation between two unconnected moment choices. For other moments (dark green) that are not connected with the target moment (red) but have overlapped, their relations can be implicitly captured with two loops, namely two graph attention layers. We can guarantee the message passing between the dark green moment and cyan moments in the first loop. And then, in the second loop, we can construct relations between cyan moments and target moment, where the information from the dark green moment is finally propagated to the red moment.

Given the choice-query $\hat{\mathbf{F}}_A \in \mathbb{R}^{C \times N}$, there are N nodes and approximately $2TN$ edges in the graph. A GAT layer inspired by (Huang et al., 2019) is applied on the graph: for each moment, we compute attention weights of its neighbours in a criss-cross path, and average the features with the weights. The output of the GAT layer can be

formulated as:

$$\hat{\mathbf{F}}_A^* = \text{Conv}(\text{GAT}(\text{Conv}(\hat{\mathbf{F}}_A), \hat{A})) \quad (6)$$

where \hat{A} is the adjacency matrix of the graph to determine the connections between two moment choices, defined by the predefined answer set A . **Answer Ranker.** Since we have captured the relationship between multi-choices by GCNs, we adopt this answer ranker to predict the ranking score of each answer candidate a_i for selecting the best-matched one. This ranker takes the query-aware feature $\hat{\mathbf{F}}_A$ and relation-aware feature $\hat{\mathbf{F}}_A^*$ as input and concatenate them (denoted as \parallel) to aggregate more contextual information. After that, we employ a convolution layer to generate the probability P_A of being selected for a_i in the predefined answer set A . The output can be computed as:

$$P_A = \sigma(\text{Conv}(\hat{\mathbf{F}}_A^* \parallel \hat{\mathbf{F}}_A)), \quad (7)$$

where σ represents the *sigmoid* activation function.

3.3 Training and Inference

Following (Zhang et al., 2020b), we first calculate the Intersection-over-Union (IoU) between the answer set A and ground-truth annotation (g^s, g^e) and then rescale them by two thresholds θ_{min} and θ_{max} , which can be written as:

$$g_i = \begin{cases} 0 & \theta \leq \theta_{min} \\ \frac{\theta_i - \theta_{min}}{\theta_{max} - \theta_{min}} & \theta_{min} < \theta < \theta_{max} \\ 1 & \theta \geq \theta_{max} \end{cases} \quad (8)$$

where g_i and θ_i are the supervision label and corresponding IoU between a_i and ground-truth respectively. Hence, the total training loss function of our RaNet is:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (g_i \log p_i + (1 - g_i) \log(1 - p_i)), \quad (9)$$

where p_i is the output score in P_A for the answer choice a_i and N is the number of multi-choices. In the inference stage, we rank all the answer options in A according to the probability in P_A .

4 Experiments

To evaluate the effectiveness of the proposed approach, we conduct extensive experiments on three public challenging datasets: TACoS (Regneri et al., 2013), ActivityNet Captions (Krishna et al., 2017) and Charades-STA (Sigurdsson et al., 2016).

Methods	Rank1@		Rank5@	
	0.3	0.5	0.3	0.5
MCN	-	5.58	-	10.33
CTRL	18.32	13.30	36.69	25.42
ACRN	19.52	14.62	34.97	24.88
ROLE	15.38	9.94	31.17	20.13
TGN	21.77	18.9	39.06	31.02
ABLR	19.50	9.40	-	-
SM-R1	20.25	15.95	38.47	27.84
CMIN	24.64	18.05	38.46	27.02
QSPN	20.15	15.23	36.72	25.30
ACL-K	24.17	20.01	42.15	30.66
2D-TAN	<u>37.29</u>	<u>25.32</u>	<u>57.81</u>	<u>45.04</u>
DRN	-	23.17	-	33.36
DEBUG	23.45	11.72	-	-
VSLNet	29.61	24.27	-	-
Ours	43.34	33.54	67.33	55.09

Table 1: Performance comparison on TACoS. All results are reported in percentage (%).

4.1 Dataset

TACoS. It consists of 127 videos, which contain different activities that happened in the kitchen. We follow the convention in (Gao et al., 2017), where the training, validation, and testing contain 10,146, 4,589, and 4,083 query-video pairs.

Charade-STA. It is extended by (Gao et al., 2017) with language descriptions leading to 12,408 and 3,720 query-video pairs for training and testing.

ActivityNet Captions. It is introduced into the temporal language grounding task recently. Following the setting in CMIN (Lin et al., 2020), we use val_1 as validation set and val_2 as testing set, which have 37,417, 17,505, and 17,031 query-video pairs for training, validation, and testing, respectively.

4.2 Implementation Details

Evaluation metric. Following Gao et al. (Gao et al., 2017), we compute the $Rank\ k@μ$ for a fair comparison. It denotes the percentage of testing samples that have at least one correct answer in the top-K choices. A selected choice a_i is correct when its IoU θ_i with the ground-truth is larger than the threshold $μ$; otherwise, the choice is wrong. Specifically, we set $k \in \{1, 5\}$ and $μ \in \{0.3, 0.5\}$ for TACoS and $μ \in \{0.5, 0.7\}$ for the other two.

Feature Extractor. We follow the (Zhang et al., 2019a; Lin et al., 2020) and adopt the same extrac-

tor, e.g., VGG (Simonyan and Zisserman, 2014) feature for Charades and C3D (Tran et al., 2015) for other two. We also use the I3D (Carreira and Zisserman, 2017) feature to make comparison with (Ghosh et al., 2019; Zhang et al., 2020a; Zeng et al., 2020) on Charades. For word embedding, we use the pre-trained GloVe 6B 300d (Pennington et al., 2014) as previous solutions (Ge et al., 2019; Chen et al., 2018).

Architecture settings. In all experiments, we set the hidden units of Bi-LSTM as 256 and the number of reshaped snippet T is defined as 128 for TACoS, 64 for ActivityNet Captions and 16 for Charades-STA. The dimension C of channels is 512. We adopt 2 GAT layers for all benchmarks and position embedding is used in ActivityNet Caption as (Zeng et al., 2020).

Training settings. We adopt the Adam optimizer with learning rate of 1×10^{-3} , batch size of 32, and training epoch of 15. Following (Zhang et al., 2020b), thresholds θ_{min} and θ_{max} are set to 0.5 and 1.0 for Charades-STA and ActivityNet Captions, while 0.3 and 0.7 for TACoS.

4.3 Comparison with State-of-the-arts

Our RaNet is compared with recent published state-of-the-art methods: VLSNet (Zhang et al., 2020a), 2D-TAN (Zhang et al., 2020b), DRN (Zeng et al., 2020), CMIN (Lin et al., 2020), DEBUG (Lu et al., 2019), QSPN (Xu et al., 2019), MAN (Zhang et al., 2019a), ExCL (Ghosh et al., 2019), CTRL (Gao et al., 2017), etc.. The top-2 performance values are highlighted by bold and underline, respectively.

TACoS. Table 1 summarizes performance comparison of different methods on the test split. Our RaNet outperforms all the competitive methods with clear margins and reports the highest scores for all IoU thresholds. Compared with the previous best method 2D-TAN, our model achieves 6% absolute improvement at least across all evaluation settings in terms of $Rank\ 1@μ$, i.e., 8.22% for $μ = 0.5$. For evaluation metric of $Rank\ 5@μ$, we even reach around 10% absolute improvement. It is worth noting that we exceed VLSNet by 9.27% and 13.73% in terms of $Rank\ 1@μ = 0.5$, $μ = 0.3$ respectively, which also formulates this task from the perspective of MRC.

Charades-STA. We evaluate our method both on VGG and I3D features used in previous works for fair comparison. Our approach reaches the highest score in terms of $Rank\ 1$ no matter which

Methods	Rank1@		Rank5@	
	0.5	0.7	0.5	0.7
VGG				
MCN	17.46	8.01	48.22	26.73
CTRL	23.63	8.89	58.92	29.52
ABLR	24.36	9.01	-	-
QSPN	35.60	15.80	79.40	45.40
ACL-K	30.48	12.20	64.84	35.13
DEBUG	37.39	17.69	-	-
MAN	41.24	20.54	83.21	51.85
2D-TAN	39.70	23.31	80.32	51.26
DRN	<u>42.90</u>	<u>23.68</u>	87.80	54.87
Ours	43.87	26.83	<u>86.67</u>	<u>54.22</u>
I3D				
ExCL	44.10	22.40	-	-
VSLNet	<u>54.19</u>	<u>35.22</u>	-	-
DRN	53.09	31.50	89.06	60.05
Ours	60.40	39.65	89.57	64.54

Table 2: Performance comparison on Charades-STA. All results are reported in percentage (%).

kind of feature is adopted as illustrated in Table 2. For VGG feature, we improve the performance from 23.68% in DRN to 26.83% in terms of $Rank\ I@μ = 0.7$. By adopting the stronger I3D feature, our method also exceeds VSLNet in terms of $Rank\ I@μ = \{0.5, 0.7\}$ (*i.e.*, 60.40% vs. 54.19% and 39.65% vs. 35.22%).

ActivityNet-Captions. In Table 3, we compare our model with other competitive methods. Our model achieves the highest scores over all IoU thresholds in the evaluation except the result of $Rank\ 5@μ = 0.5$. Particularly, our model outperforms the previous best method (*i.e.*, 2D-TAN) by around 1.29% absolute improvement, in terms of $Rank\ I@μ = 0.7$. Due to the same sampling strategy for moment candidates, this improvement is mostly attributed to the token-aware visual representation and the relationships mining between multi-choices.

4.4 Ablation Study

4.4.1 Effectiveness of Network Components

We perform complete and in-depth studies on the effectiveness of our choice-query interactor and multi-choice relation constructor based on the TACoS and Charades-STA datasets. On each dataset, we conduct five comparison experiments

Methods	Rank1@		Rank5@	
	0.5	0.7	0.5	0.7
MCN	21.36	6.43	53.23	29.70
CTRL	29.01	10.34	59.17	37.54
ACRN	31.67	11.25	60.34	38.57
TGN	27.93	-	44.20	-
QSPN	33.26	13.43	62.39	40.78
ExCL	42.7	24.1	-	-
CMIN	44.62	24.48	69.66	52.96
ABLR	36.79	-	-	-
DEBUG	39.72	-	-	-
2D-TAN	44.05	<u>27.38</u>	76.65	<u>62.26</u>
DRN	<u>45.45</u>	24.39	77.97	50.30
VSLNet	43.22	26.16	-	-
Ours	45.59	28.67	<u>75.93</u>	62.97

Table 3: Performance comparison on ActivityNet Captions. All results are reported in percentage (%).

for evaluation. First, we remove F_2 and R to explore the *RaNet-base* model, compared with only using F_2 . Then, we integrate the interaction and relation modules into our third and fourth experiments respectively. Finally, we show the best performance achieved by our proposed approach. Table 4 summarizes the grounding results in terms of $Rank\ I@μ \in \{0.3, 0.5, 0.7\}$. Without the interaction and relation modules, our framework can achieve 40.99% and 28.54% for $μ = 0.3$ and 0.5 respectively. It outperforms the previous best method 2D-TAN, indicating the power of our modality-wise encoder. When we consider the token-aware visual representation, our framework can bring significant improvement on both datasets. Improvements can also be observed when adding relation module into our framework. These results demonstrate the effect of our RaNet on temporal language grounding.

4.4.2 Improvement on different IoUs

To have a better understanding of our approach, we illustrate the performance gain achieved on three datasets in terms of different $μ \in (0, 1)$ with previous best method, 2D-TAN, as shown in Figure 4. This figure visualizes the detailed comparison between our model and 2D-TAN, which shows that our approach can continuously improve the performance, especially for higher IoUs (*i.e.*, $μ > 0.7$). It is observed that the value of relative improvement increases along with the increasing IoU on TACoS and ActivityNet Captions datasets.




Figure 4: Detailed comparison across different IoUs on three benchmarks in terms of Rank 1

Datasets	Components			Rank1@	
	F ₁	F ₂	R	0.3	0.5
TACoS	✓	✗	✗	40.99	28.54
	✗	✓	✗	41.26	29.22
	✓	✓	✗	41.51	29.64
	✓	✗	✓	42.26	32.04
	✓	✓	✓	43.34	33.54
Charades-STA	F ₁	F ₂	R	0.5	0.7
	✓	✗	✗	43.06	24.70
	✗	✓	✗	42.72	24.33
	✓	✓	✗	42.10	24.78
	✓	✗	✓	43.60	25.30
✓	✓	✓	43.87	26.83	

Table 4: Effectiveness of each component in our proposed approach on TACoS and Charades-STA, measured by Rank 1@ $\mu \in \{0.3, 0.5, 0.7\}$. VGG features are used in Charades-STA. ✓ and ✗ denote the net with and without that component, respectively.

4.4.3 Feature Initialization Functions

We conduct experiments to reveal the effect of different feature initialization functions. For a moment candidate (t_i^s, t_i^e) , it has the corresponding feature sequence Y from \hat{V} denoted as $Y = \{\hat{v}_k\}_{k=t_i^s}^{t_i^e}$. We explore four types of operators (*i.e.*, pooling, sampling, concatenation and addition) in the multi-choice generator. The first two consider all the information of the region within the temporal span of the candidate. Pooling operator focuses on the statistic characteristic and the sampling serves as weight average operator. On the contrary, the last two only consider the boundary information ($v_{t_i^s}$ and $v_{t_i^e}$) of a moment candidate, which expect the cross-modal interaction to be boundary sensitive. Table 5 reports the performance of different operators on TACoS dataset. It is observed that concatenation operator achieves

Ψ	Rank1@		Rank5@	
	0.3	0.5	0.3	0.5
Pooling	38.84	29.29	63.31	49.86
Sampling	41.33	30.82	65.58	54.69
Addition	42.69	31.59	64.98	54.36
Concatenation	43.34	33.54	67.33	55.09

Table 5: Effectiveness of different operators used in Multi-Choice Generator on TACoS, measured by Rank 1@ $\mu \in \{0.3, 0.5\}$ and Rank 5@ $\mu \in \{0.3, 0.5\}$.

the highest score across all the evaluation criterion, which indicates boundary sensitive operators have better performance than the statistical operators.

4.4.4 Word Embeddings Comparison

To further explore the effect of different textual features, we also conduct experiments on four pre-trained word embeddings (*i.e.*, GloVe 6B, GloVe 42B, GloVe 840B and BERT_{Base}). GloVe (Pennington et al., 2014) is an unsupervised learning algorithm for obtaining vector representations for words, which has some common word vectors trained on different corpora of varying sizes. BERT (Devlin et al., 2019) is a language representation model considering bidirectional context, which achieved SOTA performance on many NLP tasks. All the GloVe vectors have 300 dimensions whereas BERT_{Base} is a 768-dimensional vector. Table 6 compares the performance of these four pre-trained word embeddings on TACoS dataset. From the results we can see that better word embeddings (*i.e.* BERT) tend to have better performance, indicating us pay more attention to textual features encoding. All the models in our paper use concatenation feature initialization functions and GloVe 6B word vectors if not specified.

Methods	Rank1@			
	0.1	0.3	0.5	0.7
GloVe 6B	54.26	43.34	33.54	18.57
GloVe 42B	54.74	44.21	34.37	20.24
GloVe 840B	53.11	44.51	34.87	19.65
BERT _{Base}	57.34	46.26	34.72	21.54

Table 6: Comparison of different word embeddings on TACoS, measured by $Rank\ 1@μ \in \{0.1, 0.3, 0.5, 0.7\}$.

		TACoS	Charades	ActivityNet
Params	2D-TAN	60.93M	60.93M	91.59M
	RaNet-base	61.52M	59.95M	90.64M
	RaNet	12.80M	12.80M	10.99M
FLOPs	2D-TAN	2133.26G	104.64G	997.30G
	RaNet-base	2137.68G	104.72G	999.54G
	RaNet	175.36G	4.0G	43.92G

Table 7: Parameters and FLOPs of our RaNet with the previous best method 2D-TAN, which also considers moment-level relations. M and G represent 10^6 and 10^9 respectively.

4.4.5 Efficiency of Our RaNet

Both fully-connected graph neural networks and stacked convolution layers result in high computation complexity and occupy a huge number of GPU memory. With the help of a sparsely-connected graph attention module used in our Multi-choice Relation Constructor, we can capture moment-wise relations from global dependencies in a more efficient and effective way. Table 7 shows the parameters and FLOPs of our model and 2D-TAN, which uses several convolution layers to capture context of adjacent moment candidates. We can see that RaNet is more lightweight with only 11 M parameters compared with 92 M of 2D-TAN on ActivityNet. Compared with RaNet, RaNet-base replaces the relation constructor with the same 2D convolutional layers as 2D-TAN. Hence, their comparison on FLOPs further indicates the efficiency of our relation constructor against simple convolution layers.

4.4.6 Qualitative Analysis

We further show some examples in Figure 5 from ActivityNet Captions dataset. From this comparison, we can find that predictions of our approach are closer to ground truth than our baseline model, which is the one removing F_2 and R in Table 4. Considering the same setting for the moment candidate, it also demonstrates the effect of our pro-

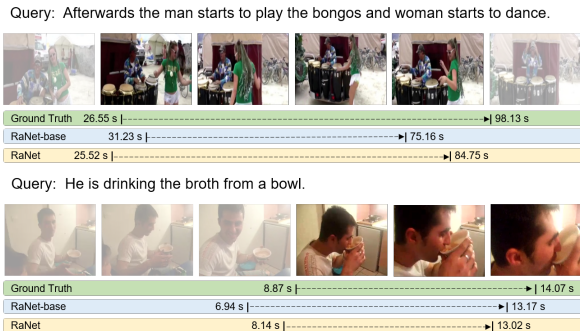


Figure 5: The qualitative results of RaNet and RaNet-base on the ActivityNet Captions dataset.

posed modules. With the interaction and relations construction modules, our approach can select the choice of video moments matching the query sentence best. In turn, it reflects that capturing the token-aware visual representation for moment candidates and relations among candidates facilitate the net scoring candidates better.

5 Conclusion

In this paper, we propose a novel Relation-aware Network to address the problem of temporal language grounding in videos. We first formulate this task from the perspective of multi-choice reading comprehension. Then we propose to interact the visual and textual modalities in a coarse-and-fine fashion for token-aware and sentence-aware representation of each choice. Further, a GAT layer is introduced to mine the exhaustive relations between multi-choices for better ranking. Our model is efficient and outperforms the state-of-the-art methods on three benchmarks, i.e., ActivityNet-Captions, TACoS, and Charades-STA.

Acknowledgements

First of all, I would like to give my heartfelt thanks to all the people who have ever helped me in this paper. The support from CloudWalk Technology Co., Ltd is gratefully acknowledged. This work was also supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research through the Visual Computing Center (VCC) funding.

References

- Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. 2017. [Localizing moments in video with natural language.](#)

- In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812.
- Joao Carreira and Andrew Zisserman. 2017. [Quo vadis, action recognition? a new model and the kinetics dataset](#). In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308.
- 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.
- 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*, volume 33, 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jianfeng Dong, Xirong Li, Chaoxi Xu, Shouling Ji, Yuan He, Gang Yang, and Xun Wang. 2019. [Dual encoding for zero-example video retrieval](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9346–9355.
- Jialin Gao, Zhixiang Shi, Guanshuo Wang, Jiani Li, Yufeng Yuan, Shiming Ge, and Xi Zhou. 2020. [Accurate temporal action proposal generation with relation-aware pyramid network](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10810–10817.
- 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 5267–5275.
- R. Ge, J. Gao, K. Chen, and R. Nevatia. 2019. [Mac: Mining activity concepts for language-based temporal localization](#). In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 245–253.
- Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander G Hauptmann. 2019. [Excl: Extractive clip localization using natural language descriptions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1984–1990.
- 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*, volume 33, pages 8393–8400.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural computation*, 9(8):1735–1780.
- Deng Huang, Peihao Chen, Runhao Zeng, Qing Du, Mingkui Tan, and Chuang Gan. 2020. [Location-aware graph convolutional networks for video question answering](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11021–11028.
- Zilong Huang, Xinggong Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. 2019. [Ccnets: Criss-cross attention for semantic segmentation](#). In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 603–612.
- Thomas N Kipf and Max Welling. 2016. [Semi-supervised classification with graph convolutional networks](#). *arXiv preprint arXiv:1609.02907*.
- 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.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [Race: Large-scale reading comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. 2018. [Tvqa: Localized, compositional video question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1369–1379.
- Zhijie Lin, Zhou Zhao, Zhu Zhang, Zijian Zhang, and Deng Cai. 2020. [Moment retrieval via cross-modal interaction networks with query reconstruction](#). *IEEE Transactions on Image Processing*, 29:3750–3762.
- Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. 2018. [Attentive moment retrieval in videos](#). In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 15–24.
- Chujie Lu, Long Chen, Chilie 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 5147–5156.

- Larry R Medsker and LC Jain. 2001. [Recurrent neural networks](#). *Design and Applications*, 5.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. 2013. [Grounding action descriptions in videos](#). *Transactions of the Association for Computational Linguistics*, 1:25–36.
- 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 (ECCV)*, pages 200–216.
- Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. 2016. [Hollywood in homes: Crowdsourcing data collection for activity understanding](#). In *ECCV (1)*, pages 510–526.
- Karen Simonyan and Andrew Zisserman. 2014. [Very deep convolutional networks for large-scale image recognition](#). *arXiv preprint arXiv:1409.1556*.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019. [Dream: A challenge data set and models for dialogue-based reading comprehension](#). *Transactions of the Association for Computational Linguistics*, 7:217–231.
- 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, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).
- Jingwen Wang, Lin Ma, and Wenhao Jiang. 2020. [Temporally grounding language queries in videos by contextual boundary-aware prediction](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12168–12175.
- Weining Wang, Yan Huang, and Liang Wang. 2019. [Language-driven temporal activity localization: A semantic matching reinforcement learning model](#). In *Proceedings of the IEEE/CVF 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 *IJCAI*, volume 3, page 8.
- 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*, volume 33, pages 9062–9069.
- Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. 2020. [G-tad: Sub-graph localization for temporal action detection](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- 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*, volume 33, pages 9159–9166.
- Runhao Zeng, Haoming Xu, Wenbing Huang, Peihao Chen, Minghui Tan, and Chuang Gan. 2020. [Dense regression network for video grounding](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 1247–1257.
- Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. 2020a. [Span-based localizing network for natural language video localization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6543–6554, Online. Association for Computational Linguistics.
- 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*, volume 34, pages 12870–12877.
- Songyang Zhang, Jinsong Su, and Jiebo Luo. 2019b. [Exploiting temporal relationships in video moment localization with natural language](#). *arXiv e-prints*, pages arXiv–1908.
- Zhuosheng Zhang, Hai Zhao, and Rui Wang. 2020c. [Machine reading comprehension: The role of contextualized language models and beyond](#). *arXiv preprint arXiv:2005.06249*.
- Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. 2020. [Vision-language navigation with self-supervised auxiliary reasoning tasks](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10012–10022.