

Where Does It Exist: Spatio-Temporal Video Grounding for Multi-Form Sentences

Zhu Zhang¹, Zhou Zhao^{1,*}, Yang Zhao¹, Qi Wang², Huasheng Liu², Lianli Gao³

¹College of Computer Science, Zhejiang University

²Alibaba Group, ³University of Electronic Science and Technology of China

{zhangzhu, zhaozhou, zhaoyang}@zju.edu.cn, {wq140362, fangkong.lhs}@alibaba-inc.com

Abstract

In this paper, we consider a novel task, *Spatio-Temporal Video Grounding for Multi-Form Sentences (STVG)*. Given an untrimmed video and a declarative/interrogative sentence depicting an object, STVG aims to localize the spatio-temporal tube of the queried object. STVG has two challenging settings: (1) We need to localize spatio-temporal object tubes from untrimmed videos, where the object may only exist in a very small segment of the video; (2) We deal with multi-form sentences, including the declarative sentences with explicit objects and interrogative sentences with unknown objects. Existing methods cannot tackle the STVG task due to the ineffective tube pre-generation and the lack of object relationship modeling. Thus, we then propose a novel *Spatio-Temporal Graph Reasoning Network (STGRN)* for this task. First, we build a spatio-temporal region graph to capture the region relationships with temporal object dynamics, which involves the implicit and explicit spatial subgraphs in each frame and the temporal dynamic subgraph across frames. We then incorporate textual clues into the graph and develop the multi-step cross-modal graph reasoning. Next, we introduce a spatio-temporal localizer with a dynamic selection method to directly retrieve the spatio-temporal tubes without tube pre-generation. Moreover, we contribute a large-scale video grounding dataset *VidSTG* based on video relation dataset *VidOR*. The extensive experiments demonstrate the effectiveness of our method. The *VidSTG* dataset is available at <https://github.com/Guaranteer/VidSTG-Dataset>.

1. Introduction

Grounding natural language in visual contents is a fundamental and vital task in the visual-language understanding field. Visual grounding aims to localize the object described by the given referring expression in an im-



Declarative Sentence: A little **boy** with a Christmas hat is catching a yellow toy.

Interrogative Sentence: **What** is caught by the squatting boy on the floor?

Figure 1. An example of STVG for multi-form sentences.

age, which has attracted much attention and made great progress [14, 24, 6, 40, 37]. Recently, researchers begin to explore video grounding, including temporal grounding and spatio-temporal grounding. Temporal sentence grounding [8, 11, 46, 35] determines the temporal boundaries of events corresponding to the given sentence, but does not localize the spatio-temporal tube (i.e., a sequence of bounding boxes) of the described object. Further, spatio-temporal grounding is to retrieve the object tubes according to textual descriptions, but existing strategies [48, 1, 36, 4] can only be applied to restricted scenarios, e.g. grounding in a frame of the video [48, 1] or grounding in trimmed videos [36, 4]. Moreover, due to the lack of bounding box annotations, researchers [48, 4] can only adopt a weakly-supervised setting, leading to suboptimal performance.

To break through above restrictions, we propose a novel task, *Spatio-Temporal Video Grounding for Multi-Form Sentences (STVG)*. Concretely, as illustrated in Figure 1, given an untrimmed video and a declarative/interrogative sentence depicting an object, STVG aims to localize the spatio-temporal tube of the queried object.

Compared with previous video grounding [48, 1, 36, 4], STVG has two novel and challenging points. First, we localize spatio-temporal object tubes from untrimmed videos. The objects may exist in a very small segment of the video and be hard to distinguish. And the sentences may only describe a short-term state of the queried object, e.g. the action "catching a toy" of the boy in Figure 1. So it is crucial to determine the temporal boundaries of object tubes by sufficient cross-modal understanding. Secondly, STVG deals with multi-form sentences, that it, not only grounds

*Zhou Zhao is the corresponding author.

the conventional declarative sentences with explicit objects, but also localizes the interrogative sentences with unknown objects, for example, the sentence "What is caught by the squatting boy on the floor?" in Figure 1. Due to the lack of the explicit characteristics of objects (e.g. the class "toy" and visual appearance "yellow"), grounding for interrogative sentences can only depend on relationships between the unknown object and other objects (e.g. the action relation "caught by the squatting boy" and spatial relation "on the floor"). Thus, sufficient relationship construction and cross-modal relation reasoning are crucial for the STVG task.

Existing video grounding methods [36, 4] often extract a set of spatio-temporal tubes from trimmed videos and then identify the target tube that matches the sentence. However, this framework may be unsuitable for STVG. On the one hand, the performance of this framework is heavily dependent on the quality of tube candidates. But it is difficult to pre-generate high-quality tubes without textual clues, since the sentences may describe a short-term state of objects in a very small segment, but the existing tube pre-generation framework [36, 4] can only produce the complete object tubes from trimmed videos. On the other hand, these methods only consider single tube modeling and ignore the relationships between objects. However, object relations are vital clues for the STVG task, especially for interrogative sentences that may only offer the interactions of the unknown objects with other objects. Thus, these approaches cannot deal with the complicated grounding of STVG.

To tackle above problems, we propose a novel Spatio-Temporal Graph Reasoning Network (STGRN) to capture region relationships with temporal object dynamics and directly localize the spatio-temporal tubes without tube pre-generation. Specifically, we first parse the video into a spatio-temporal region graph. Existing visual graph modeling approaches [39, 18] often build the spatial graph in an image, which cannot utilize the temporal dynamics information in videos to distinguish the subtle differences of object actions, e.g. distinguish "open the door" and "close the door". Different from them, our spatio-temporal region graph not only involves the implicit and explicit spatial subgraphs in each frame, but also includes a temporal dynamic subgraph across frames. The spatial subgraphs can capture the region-level relationships by implicit or explicit semantic interactions, and the temporal subgraph can incorporate the dynamics and transformation of objects across frames to further improve the relation understanding. Next, we fuse the textual clues into this spatio-temporal graph as the guidance, and develop the multi-step cross-modal graph reasoning. The multi-step process can support the multi-order relation modeling like "a man hug a baby wearing a red hat". After it, we introduce a spatio-temporal localizer to directly retrieve the spatio-temporal object tubes from the region level. Concretely, we first employ a temporal localizer to

determine the temporal boundaries of the tube, and then apply a spatial localizer with a dynamic selection method to ground the object in each frame and generate a smooth tube.

To facilitate this STVG task, we contribute a large-scale video grounding dataset VidSTG by adding the multi-form sentence annotations into video relation dataset VidOR.

Our main contributions can be summarized as follows:

- We propose a novel task STVG to explore the spatio-temporal video grounding for multi-form sentences.
- We develop a novel STGRN to tackle this STVG task, which builds a spatio-temporal graph to capture the region relationships with temporal object dynamics, and employs a spatio-temporal localizer to directly retrieve the spatio-temporal tubes without tube pre-generation.
- We contribute a large-scale video grounding dataset VidSTG as the benchmark of the STVG task.
- The extensive experiments show the effectiveness of our proposed STGRN method.

2. Related Work

2.1. Temporal Localization via Natural Language

Temporal natural language localization is to detect the video clip depicting the given sentence. Early approaches [11, 8, 12, 22, 23] are mainly based on the sliding window framework, which first samples abundant candidate clips and then ranks them. Recently, people begin to develop holistic and cross-modal methods [44, 2, 3, 33, 46, 35] to solve this problem. Chen, Zhang and Xu et al. [2, 3, 46, 21, 35] build frame-by-word interactions from visual and textual contents to aggregate the matching clues. Zhang et al. [44] deal with the structural and semantic misalignment challenges by an explicitly structured graph. Wang et al. [33] propose a reinforcement learning framework to adaptively observe frame sequences and associate video contents with sentences. Further, Mithun and Lin et al. [25, 20] devise the weakly-supervised temporal localization methods requiring only coarse video-level annotations for training. And besides the natural language query, Zhang et al. [47] attempt to localize the unseen temporal clip according to an image query. Although these methods have achieved promising performance, they still remain in the temporal grounding. We further explore spatio-temporal grounding in this paper.

2.2. Object Grounding in Images/Videos

Visual grounding [14, 24, 34, 41, 26, 42, 6, 49, 13, 40, 45, 37, 38] aims to localize the visual object described by the given referring expression. Early methods [14, 24, 34, 41, 26, 42] often extract object features




Figure 2. The Overall Architecture of the Spatio-Temporal Graph Reasoning Network (STGRN). We first extract the region features and learn the query representation. Next, we apply a spatio-temporal graph encoder to develop the multi-step cross-modal graph reasoning. After it, a spatio-temporal localizer with T spatio-temporal convolution layers directly retrieves the tubes from the region level.

by CNN, model the language expression through RNN and learn the object-language matching. Some recent approaches [40, 13] decompose the expression into multiple components and calculate matching scores of each module. And Deng and Zhuang et al. [6, 49] apply the co-attention mechanism to build cross-modal interactions. Further, Yang et al. [38, 37] explore the relationships between objects to improve the accuracy. As for video grounding, existing works [48, 1, 36, 4, 36, 4] can only be applied to restricted scenarios. Zhou and Balajee et al. [48, 1] only ground natural language in a frame of the video. With a sequence of transcriptions and temporal alignment video clips, Huang and Shi et al. [15, 31] ground nouns or pronouns in specific frames by the weakly-supervised MIL methods. And Chen and Yamaguchi et al. [36, 4] localize spatio-temporal object tubes from trimmed videos, but cannot directly deal with raw video streams. In this paper, we propose a novel STVG task to further explore the spatio-temporal video grounding for multi-form sentences.

3. The Proposed Method

Given a video $v \in V$ and a sentence $s \in S$ depicting an object, STVG is to retrieve its spatio-temporal tube. Figure 2 illustrates the overall framework.

3.1. Video and Text Encoder

We first apply a pre-trained Faster R-CNN [28] to extract a set of regions for each frame, where a video contains N frames and the t -th frame corresponds to K regions denoted by $\{r_i^t\}_{i=1}^K$. A region r_i^t is associated with its visual feature

$\mathbf{r}_i^t \in \mathbb{R}^{d_r}$ and bounding box vector $\mathbf{b}_i^t = [x_i^t, y_i^t, w_i^t, h_i^t]$, where (x_i^t, y_i^t) are the normalized coordinates of the center point and (w_i^t, h_i^t) are the normalized width and height. Besides, we obtain the frame features $\{\mathbf{f}^t\}_{t=1}^N$, i.e., the visual feature of the entire frame from Faster R-CNN.

For the sentence s with L words, we first input the word embeddings into a bi-directional GRU [5] to learn the word semantic features $\{\mathbf{s}_i\}_{i=1}^L \in \mathbb{R}^{d_s}$, where \mathbf{s}_i is the concatenation of the forward and backward hidden states of step i . As the STVG grounds objects described in the declarative or interrogative sentences, we need to extract the query representation from language context. First, we select the entity feature \mathbf{s}^e from $\{\mathbf{s}_i\}_{i=1}^L$, which represents the queried object, for example, \mathbf{s}_3 for "boy" in Figure 2. Note that for interrogative sentences, the feature of "who" or "what" is chosen as the entity feature. Next, an attention method aggregates the textual clues from language context by

$$\gamma_i = (\mathbf{W}_1^s \mathbf{s}^e)^\top \cdot (\mathbf{W}_2^s \mathbf{s}_i), \tilde{\gamma}_i = \frac{\exp(\gamma_i)}{\sum_{i=1}^L \exp(\gamma_i)}, \quad (1)$$

$$\mathbf{s}^a = \sum_{i=1}^L \tilde{\gamma}_i \mathbf{s}_i, \mathbf{s}^a = [\mathbf{s}^e; \mathbf{s}^a],$$

where \mathbf{W}_1^s and \mathbf{W}_2^s are projection matrices, \mathbf{s}^a is the entity-aware feature and \mathbf{s}_q is the query representation.

3.2. Spatio-Temporal Graph Encoder

Our STVG task requires to capture the object relationships and develop the cross-modal understanding, especially for interrogative sentences that may only offer the

interaction information of the unknown object with other objects. Thus, we build a spatio-temporal region encoder with T spatio-temporal convolution layers to capture the region relationships with temporal object dynamics and support the multi-step cross-modal graph reasoning.

3.2.1 Graph Construction

We first parse the video into a spatio-temporal region graph, which involves the implicit spatial subgraph $\mathcal{G}_{imp} = (\mathcal{V}, \mathcal{E}_{imp})$, explicit spatial subgraph $\mathcal{G}_{exp} = (\mathcal{V}, \mathcal{E}_{exp})$ in each frame and temporal dynamic subgraph $\mathcal{G}_{tem} = (\mathcal{V}, \mathcal{E}_{tem})$ across frames. Three subgraphs all treat regions as their vertexes \mathcal{V} but have different edges. Note that we add the self-loop of each vertex in each subgraph.

Implicit Spatial Graph. We regard the fully-connected region graph in each frame as the implicit spatial graph \mathcal{G}_{imp} , where \mathcal{E}_{imp} contains $K \times K$ undirected and unlabeled edges in each frame (including self-loops).

Explicit Spatial Graph. We extract the region triplet $\langle r_i^t, p_{ij}^t, r_j^t \rangle$ to construct the explicit spatial graph, where r_i^t and r_j^t are the i -th and j -th regions in frame t , and p_{ij}^t is the relation predicate between them. Each triplet can be regarded as an edge from i to j . Thus, explicit graph construction can be formulated as a relation classification task [39, 43]. Concretely, given the feature $[\mathbf{r}_i^t; \mathbf{b}_i^t]$ of region i , feature $[\mathbf{r}_j^t; \mathbf{b}_j^t]$ of region j and the united feature $[\mathbf{r}_{ij}^t; \mathbf{b}_{ij}^t]$ of the union bounding box of i and j (also extracted by Faster R-CNN), we first transform three features via different linear layers and then concatenate them into a classification layer to predict the relations. Similar to existing works [39, 18], we train such a classifier on the Visual Genome dataset [17], where we select top-50 frequent predicates in its training data and add an extra *no_relation* class for non-existent edges. We then predict the relationships between r_i^t and r_j^t . Eventually, the edges \mathcal{E}_{exp} have 3 directions (including i -to- j , j -to- i and i -to- i of self-loops) and 51 types of labels (top-50 classes plus the self-loop).

Temporal Dynamic Graph. While the spatial graphs model region-region interactions, our temporal graph is to capture the dynamics and transformation of objects across frames. So we expect to connect the regions containing the same object in different frames and then learn more expressive and discriminative object features. For the frame t , we connect its regions with adjacent $2M$ frames (M for forward frames and M for backward). Too distant frames cannot provide the real-time dynamics. Concretely, we first define the linking score $s(r_i^t, r_j^k)$ between r_i^t from frame t and r_j^k from frame k by

$$s(r_i^t, r_j^k) = \cos(\mathbf{r}_i^t, \mathbf{r}_j^k) + \frac{\epsilon}{|k-t|} \cdot \text{IoU}(r_i^t, r_j^k), \quad (2)$$

where $\cos(\cdot)$ is the cosine similarity of two features, $\text{IoU}(\cdot)$ is the intersection-over-union of two regions, and ϵ is the

balanced scalar. Here we simultaneously consider the appearance similarity and spatial overlap ratio of two regions. And the temporal distance $|k - t|$ of two frames is used to limit the IoU score, that is, for distant frames, the linking score is mainly determined by feature similarity rather than the spatial overlap. Next, for the region r_i^t , we select the region r_j^k with the maximal linking score from frame k to build an edge, and get $2M + 1$ edges for each region (including the self-loop). The unlabeled temporal edges \mathcal{E}_{tem} have 3 directions: forward, backward and self-loop.

3.2.2 Multi-Step Cross-Modal Graph Reasoning

After graph construction, we incorporate the textual clues into this graph and develop the multi-step cross-modal graph reasoning by T spatio-temporal convolution layers.

Cross-Modal Fusion. To capture the relationships associated with the sentence, we first use a cross-modal fusion that dynamically injects textual evidences into the spatio-temporal graph. Concretely, we first utilize an attention mechanism to aggregate the words features for each region. For a region r_i^t , we calculate the attention weights over word features $\{\mathbf{s}_i\}_{i=1}^L$, denoted by

$$\alpha_{ij}^t = \mathbf{w}_m^\top \tanh(\mathbf{W}_1^m \mathbf{r}_i^t + \mathbf{W}_2^m \mathbf{s}_j + \mathbf{b}^m),$$

$$\tilde{\alpha}_{ij}^t = \frac{\exp(\alpha_{ij}^t)}{\sum_{j=1}^L \exp(\alpha_{ij}^t)}, \quad \mathbf{q}_i^t = \sum_{j=1}^L \tilde{\alpha}_{ij}^t \mathbf{s}_j, \quad (3)$$

where \mathbf{W}_1^m , \mathbf{W}_2^m are projection matrices, \mathbf{b}^m is the bias and \mathbf{w}_m^\top is the row vector. And \mathbf{q}_i^t is the region-aware textual feature for each region i in frame t .

Next, we build the textual gate that takes language information as the guidance to weaken the text-irrelevant regions, given by

$$\mathbf{g}_i^t = \sigma(\mathbf{W}^g \mathbf{q}_i^t + \mathbf{b}^g), \quad \mathbf{v}_i^t = [\mathbf{r}_i^t \odot \mathbf{g}_i^t; \mathbf{q}_i^t], \quad (4)$$

where σ is the sigmoid function, \odot is the element-wise multiplication, $\mathbf{g}_i^t \in \mathbb{R}^{d_r}$ means the textual gate for region r_i^t . And we then concatenate the filtered region feature and textual feature to obtain the cross-modal region features $\{\{\mathbf{v}_i^t\}_{i=1}^K\}_{t=1}^N$. Next, we develop T spatio-temporal convolution layers for the multi-step graph reasoning.

Spatial Graph Convolution. In each layer, we first develop the spatial graph convolution to capture visual relationships among regions in each frame. Concretely, with cross-modal region features $\{\{\mathbf{v}_i^t\}_{i=1}^K\}_{t=1}^N$, we first adopt the implicit graph convolution on \mathcal{G}_{imp} that is undirected and unlabeled, given by

$$\bar{\mathbf{v}}_i^t = \sum_{j \in \mathcal{N}_i(r_i^t)} \alpha_{ij}^{imp} \cdot (\mathbf{W}^{imp} \mathbf{v}_j^t),$$

$$\alpha_{ij}^{imp} = \frac{\exp((\mathbf{U}^{imp} [\mathbf{v}_i^t; \mathbf{b}_i^t])^\top \cdot (\mathbf{U}^{imp} [\mathbf{v}_j^t; \mathbf{b}_j^t]))}{\sum_{j \in \mathcal{N}_i(r_i^t)} \exp((\mathbf{U}^{imp} [\mathbf{v}_i^t; \mathbf{b}_i^t])^\top \cdot (\mathbf{U}^{imp} [\mathbf{v}_j^t; \mathbf{b}_j^t]))}, \quad (5)$$

where $\mathcal{N}_i(r_i^t)$ are the regions connected with r_i^t in \mathcal{G}_{imp} . The implicit graph convolution can be regarded as a variant of self-attention and we compute the coefficient α_{ij}^{imp} by combining the visual features and region locations.

Simultaneously, we develop the explicit graph convolution. Different from the original undirected GCN [16, 32], we consider the direction and label information of edges on the directed and labeled \mathcal{G}_{exp} , given by

$$\hat{\mathbf{v}}_i^t = \sum_{j \in \mathcal{N}_e(r_i^t)} \alpha_{lab(i,j)}^{exp} \cdot (\mathbf{W}_{dir(i,j)}^{exp} \mathbf{v}_j^t + \mathbf{b}_{lab(i,j)}^{exp}), \quad (6)$$

where $\mathbf{W}_{(\cdot)}^{exp}$ are optional matrices by the direction $dir(i, j)$ of edge (i, j) , $\mathbf{b}_{(\cdot)}^{exp}$ are optional bias by the label of edge (i, j) . Here, the edge has three directions (i -to- j , j -to- i , i -to- i) and 51 types. $\mathcal{N}_e(r_i^t)$ are the regions connected with r_i^t . Moreover, the relation coefficient $\alpha_{(\cdot)}^{exp}$ can also be chosen by the label of edge (i, j) . Different sentences describe different relations and their grounding is heavily dependent on the specific relation understanding. Thus, the coefficient of explicit edges can be decided by query presentation \mathbf{s}^q , given by

$$\alpha^{exp} = \text{Softmax}(\mathbf{W}^r \mathbf{s}^q + \mathbf{b}^r), \quad (7)$$

where $\alpha^{exp} \in \mathbb{R}^{51}$ corresponds to the coefficients of 51 types of relationships.

Temporal Graph Convolution. We next develop the temporal graph convolution on the directed and unlabeled graph \mathcal{G}_{tem} to capture the dynamics and transformation of objects across frames. We consider the forward, backward and self-loop edges for each region r_i^t with the cross-modal feature \mathbf{v}_i^t , denoted by

$$\begin{aligned} \tilde{\mathbf{v}}_i^t &= \sum_{j \in \mathcal{N}_i(r_i^t)} \alpha_{ij}^{tem} \cdot \mathbf{W}_{dir(i,j)}^{tem} \mathbf{v}_j, \\ \alpha_{ij}^{tem} &= \frac{\exp((\mathbf{U}^{tem} \mathbf{v}_i^t)^\top \cdot (\mathbf{V}_{dir(i,j)}^{tem} \mathbf{v}_j))}{\sum_{j \in \mathcal{N}_i(r_i^t)} \exp((\mathbf{U}^{tem} \mathbf{v}_i^t)^\top \cdot (\mathbf{V}_{dir(i,j)}^{tem} \mathbf{v}_j))}, \end{aligned} \quad (8)$$

where $\mathbf{W}_{(\cdot)}^{tem}$ and $\mathbf{V}_{(\cdot)}^{tem}$ are matrices and $dir(i, j)$ indicates the direction of edge (i, j) to select the corresponding projection matrix, where the temporal edges have three directions. And α_{ij}^{tem} is the semantic coefficient for each neighborhood region.

Next, we combine the outputs of spatial and temporal graph convolutions and obtain the result $\mathbf{v}_i^{t(1)}$ of the first spatio-temporal convolution layer by

$$\mathbf{v}_i^{t(1)} = \text{ReLU}(\tilde{\mathbf{v}}_i^t + \hat{\mathbf{v}}_i^t + \tilde{\mathbf{v}}_i^t + \mathbf{v}_i^t). \quad (9)$$

In order to support multi-order relation modeling, we perform the multi-step encoding by the spatio-temporal graph encoder with T spatio-temporal convolution layers and learn final relation-aware region features $\{\{\mathbf{m}_i^t\}_{i=1}^K\}_{t=1}^N$.

3.3. Spatio-Temporal Localizer

In this section, we devise a spatio-temporal localizer to determine the temporal tube boundaries and spatio-temporal tubes of objects from the region level.

Temporal Localizer. We first introduce the temporal localizer, which estimates a set of candidate clips and adjust their boundaries to obtain the temporal grounding [46]. Specifically, we first aggregate the relation-aware region graph into the frame level by an attention mechanism. With the query representation \mathbf{s}^q , the region features of each frame are attended by

$$\begin{aligned} \beta_i^t &= \mathbf{w}_f^\top \tanh(\mathbf{W}_1^f \mathbf{m}_i^t + \mathbf{W}_2^f \mathbf{s}^q + \mathbf{b}^f), \\ \tilde{\beta}_i^t &= \frac{\exp(\beta_i^t)}{\sum_{i=1}^K \exp(\beta_i^t)}, \quad \mathbf{m}^t = \sum_{i=1}^K \tilde{\beta}_i^t \mathbf{m}_i^t, \end{aligned} \quad (10)$$

where \mathbf{m}^t represents the relation-aware feature of frame t . We then concatenate these features with their corresponding global frame features $\{\mathbf{f}^t\}_{t=1}^N$, and apply another BiGRU to learn final frame features $\{\mathbf{h}^t\}_{t=1}^N$. Next, we define multi-scale candidate clips at each time step t as $B^t = \{(s_i^t, e_i^t)\}_{i=1}^P$, where $(s_i^t, e_i^t) = (t - w_i/2, t + w_i/2)$ are the start and end boundaries of the i -th clips, w_i is the temporal length of i -th clip and P is the clip number. After it, we estimate all candidate clips by a linear layer with the sigmoid nonlinearity and simultaneously produce the offsets of their boundaries, given by

$$\mathbf{C}^t = \sigma(\mathbf{W}^c [\mathbf{h}^t; \mathbf{s}^q] + \mathbf{b}^c), \quad \delta^t = \mathbf{W}^o [\mathbf{h}^t; \mathbf{s}^q] + \mathbf{b}^o \quad (11)$$

where $\mathbf{C}^t \in \mathbb{R}^P$ corresponds to confidence scores of P candidates at step t and $\delta^t \in \mathbb{R}^{2P}$ are the offsets of P clips.

The temporal localizer has two losses: the alignment loss for the clip selection and a regression loss for boundary adjustments. Concretely, for alignment loss, we first compute the temporal IoU score \hat{C}_i^t of each clip with the ground truth. And the alignment loss is denoted by

$$\mathcal{L}_{align} = -\frac{1}{NP} \sum_{t=1}^N \sum_{i=1}^P (1 - \hat{C}_i^t) \cdot \log(1 - C_i^t) + \hat{C}_i^t \cdot \log(C_i^t), \quad (12)$$

where we use the temporal IoU score \hat{C}_i^t rather than 0/1 score to further distinguish high-score clips. Next, we fine-tune the boundaries of the best clip with highest C_i^t , which has the boundaries (s, e) and offsets (δ_s, δ_e) . We first compute the offsets of this clip from ground truth boundaries (\hat{s}, \hat{e}) by $\hat{\delta}_s = s - \hat{s}$ and $\hat{\delta}_e = e - \hat{e}$ and define the regression loss by

$$\mathcal{L}_{reg} = \text{R}(\delta_s - \hat{\delta}_s) + \text{R}(\delta_e - \hat{\delta}_e), \quad (13)$$

where R represents the smooth L1 function.

Spatial Localizer. With the temporal grounding, we next localize the target regions in each frame. For the t -th

frame with region features $\{\mathbf{m}_i^t\}_{i=1}^K$, we directly estimate the matching scores of each region by integrating the query representation \mathbf{s}^q and final frame feature \mathbf{h}^t , denoted by

$$S_i^t = \sigma(\mathbf{W}^c[\mathbf{m}_i^t; \mathbf{s}^q; \mathbf{h}^t] + \mathbf{b}^c), \quad (14)$$

where S_i^t is the matching score of region i of frame t . Similar to temporal alignment loss, the spatial loss first compute the spatial IoU score \hat{S}_i^t for each region with the ground truth region, where the frames outside the temporal ground truth are omitted. And the spatial loss is denoted by

$$\mathcal{L}_{exp} = -\frac{1}{K|\mathcal{S}_t|} \sum_{t \in \mathcal{S}_t} \sum_{i=1}^K (1 - \hat{S}_i^t) \cdot \log(1 - S_i^t) + \hat{S}_i^t \cdot \log(S_i^t), \quad (15)$$

where \mathcal{S}_t is the set of frames in the temporal ground truth.

Eventually, we devise a multi-task loss to train our proposed STGRN in an end-to-end manner, given by

$$\mathcal{L}_{STGRN} = \lambda_1 \mathcal{L}_{align} + \lambda_2 \mathcal{L}_{reg} + \lambda_3 \mathcal{L}_{exp}, \quad (16)$$

where λ_1 , λ_2 and λ_3 are the hyper-parameters to control the balance of three losses.

3.4. Dynamic Selection Method

During inference, we first retrieve the temporal boundaries (T_s, T_e) of the tube from the temporal localizer and then determine the grounded region for each frame by the spatial localizer. A greedy method directly selects the regions with highest matching scores S_i^t . However, such generated tubes may not be very smooth. The bounding boxes between adjacent frames may have too large displacements. Thus, to make the trajectory smoother, we introduce a dynamic selection method. Concretely, we first define the linking score $s(r_i^t, r_j^{t+1})$ between regions of successive frames t and $t + 1$ by

$$s(r_i^t, r_j^{t+1}) = S_i^t + S_j^{t+1} + \theta \cdot \text{IoU}(r_i^t, r_j^{t+1}), \quad (17)$$

where S_i^t and S_j^{t+1} are matching scores of regions r_i^t and r_j^{t+1} , and θ is the balanced scalar which is set to 0.2. Next, we generate the final spatio-temporal tube Y with the maximal energy E given by

$$E(Y) = \frac{1}{|T_e - T_s|} \sum_{t=T_s}^{T_e-1} s(r_i^t, r_j^{t+1}), \quad (18)$$

where (T_s, T_e) are the temporal boundaries and we solve this optimization problem using a Vitervi algorithm [10].

4. Dataset

As a novel task, STVG lacks a suitable dataset as the benchmark. Therefore, we contribute a large-scale spatio-temporal video grounding dataset VidSTG by augmenting the sentence annotations on VidOR [30].

Table 1. Dataset Statistics about the Number of Declarative and Interrogative Sentences.

	#Declar. Sent.	#Inter. Sent.	All
Training	36,202	44,482	80,684
Validation	3,996	4,960	8,956
Testing	4,610	5,693	10,303
All	44,808	55,135	99,943

4.1. Dataset Annotation

VidOR [30] is the existing largest object relation dataset, containing 10,000 videos and fine-grained annotations for objects and their relations. Specifically, VidOR annotates 80 categories of objects with dense bounding boxes and annotates 50 categories of relation predicates among objects (8 spatial relations and 42 action relations). Specifically, VidOR denotes a relation as a triplet $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ and each triplet is associated with the temporal boundaries and spatio-temporal tubes of *subject* and *object*. Based on VidOR, we can select the suitable triplets, and describe the *subject* or *object* with multi-form sentences. Taking VidOR as the basic dataset has many advantages. On the one hand, we can avoid labor-intensive annotations for bounding boxes. On the other hand, the relationships in the triplets can be simply incorporated into the annotated sentences.

We first split and clean the VidOR data, and then annotate the rest video-triplet pairs with multi-form sentences. The cleaning process is introduced in the supplementary material. For each video-triplet pair, we choose the *subject* or *object* as the queried object, and then describe its appearance, relationships with other objects and visual environments. For interrogative annotations, the appearance of queried objects is ignored. We discard video-triplet pairs that are too hard to give a precise description. And a video-triplet pair may correspond to multiple sentences.

4.2. Dataset Statistics

After annotation, there are 99,943 sentence descriptions about 79 types of queried objects for 44,808 video-triplet pairs, shown in Table 1. The average duration of videos is 28.01s and the average temporal length of object tubes is 9.68s. The average lengths of declarative and interrogative sentences are 11.12 and 8.98, respectively. Further, we provide the distribution of 79 types of queried objects and some annotation examples in the supplementary material.

5. Experiments

5.1. Experimental Settings

Implementation Details. In STGRN, we first sample 5 frames per second and downsample the frame number of overlong videos to 200. We then pretrain the Faster R-CNN

Table 2. Performance Evaluation Results on the VidSTG Dataset.

Method	Declarative Sentence Grounding				Interrogative Sentence Grounding			
	m_tIoU	m_vIoU	vIoU@0.3	vIoU@0.5	m_tIoU	m_vIoU	vIoU@0.3	vIoU@0.5
Random	5.18%	0.69%	0.04%	0.01%	5.35%	0.60%	0.02%	0.01%
GroundeR + TALL		9.78%	11.04%	4.09%		9.32%	11.39%	3.24%
STPR + TALL	34.63%	10.40%	12.38%	4.27%	33.73%	9.98%	11.74%	4.36%
WSSTG + TALL		11.36%	14.63%	5.91%		10.65%	13.90%	5.32%
GroundeR + L-Net		11.89%	15.32%	5.45%		11.05%	14.28%	5.11%
STPR + L-Net	40.86%	12.93%	16.27%	5.68%	39.79%	11.94%	14.73%	5.27%
WSSTG + L-Net		14.45%	18.00%	7.89%		13.36%	17.39%	7.06%
STGRN (Greedy)	48.47%	18.99%	23.63%	13.48%	46.98%	17.46%	20.02%	11.92%
STGRN		19.75%	25.77%	14.60%		18.32%	21.10%	12.83%
GroundeR + Tem. Gt	-	28.80%	43.20%	22.74%	-	26.11%	38.37%	18.34%
STPR + Tem. Gt	-	29.72%	44.78%	23.83%	-	26.97%	39.89%	20.07%
WSSTG + Tem. Gt	-	33.32%	50.01%	29.98%	-	30.05%	44.54%	25.76%
STGRN + Tem. Gt	-	38.04%	54.47%	34.80%	-	35.70%	47.79%	31.41%

on MSCOCO [19] to extract 20 region proposals for each frame (i.e. $K = 20$). The region feature dimension d_r is 1,024 and we map it 256 before graph modeling. For sentences, we use a pretrained Glove word2vec [27] to extract 300-d word embeddings. As for the hyper-parameters, we set M to 5, ϵ to 0.8, θ to 0.2 and set $\lambda_1, \lambda_2, \lambda_3$ to 1.0, 0.001 and 1.0, respectively. The layer number T of the spatio-temporal graph encoder is set to 2. For the temporal localizer, we set P to 8 and define 8 window widths [8, 16, 32, 64, 96, 128, 164, 196]. We set the dimension of almost parameter matrices and bias to 256, including the $\mathbf{W}_{(\cdot)}^{exp}, \mathbf{b}_{(\cdot)}^{exp}$ in the explicit graph convolution, \mathbf{W}^f and \mathbf{b}^f in the temporal localizer and so on. And the BiGRU networks have 128-d hidden states for each direction. During training, we apply an Adam optimizer [7] to minimize the multi-task loss \mathcal{L}_{STGRN} , where the initial learning rate is set to 0.001 and the batch size is 16.

Evaluation Criteria. We employ the **m_tIoU**, **m_vIoU** and **vIoU@R** as evaluation criteria [9, 4]. The **m_tIoU** is the average temporal IoU between the selected clips and ground truth clips. And we define S_U as the set of frames contained in the selected or ground truth clips, and S_I as the set of frames in both selected and ground truth clips. We calculate vIoU by $vIoU = \frac{1}{|S_U|} \sum_{t \in S_I} IoU(r^t, \hat{r}^t)$, where r^t and \hat{r}^t are selected and ground truth regions of frame t . The **m_vIoU** is the average vIoU of samples and **vIoU@R** is the proportion of samples which vIoU > R.

Baseline. Since no existing strategy can be directly applies to STVG, we extend the existing visual grounding method **GroundeR** [29] and video grounding approaches **STPR** [36] and **WSSTG** [4] as the baselines. Considering these methods all lack temporal grounding, we first apply the temporal sentence localization methods **TALL** [8] and **L-Net** [3] to obtain a clip and then retrieve the tubes from the trimmed clip by GroundeR, STPR and WSSTG.

The GroundeR is a frame-level approach, which originally grounds natural language in a still image. We apply it for each frame of the clip and generate a tube. The STPR and WSSTG are both tube-level methods and adopt the tube pre-generation framework. Specifically, the original STPR [36] only grounds persons from multiple videos, we extend it to multi-type object grounding in a single clip. The original WSSTG [4] employs a weakly-supervised setting, we extend it by applying a supervised triplet loss [37] to select candidate tubes. So we obtain 6 combined baselines **GroundeR+TALL**, **STPR+TALL** and so on. We also provide the temporal ground truth to form 3 baselines. We show more baseline details in the supplementary material.

5.2. Experiment Results

Table 2 shows the overall experiment results of all methods, where **STGRN(Greedy)** uses the greedy region selection for the tube generation rather than the dynamic method. The **Random** selects the temporal clip and spatial regions randomly. **Tem. Gt** means that the temporal ground truth is provided. We can find several interesting points:

- The GroundeR+{·} methods independently ground sentences in every frame and achieve worse performance than STPR+{·} and WSSTG+{·} methods, validating the temporal object dynamics across frames are vital for spatio-temporal video grounding.
- The model performance on interrogative sentences is obviously lower than declarative sentences, which shows the interrogative sentences with unknown objects are more difficult to ground.
- For temporal grounding, our STGRN achieves a better performance than the frame-level localization method TALL and L-Net, demonstrating the spatio-temporal

Table 3. Ablation Results on the VidSTG Dataset.

\mathcal{G}_{imp}	\mathcal{G}_{exp}	\mathcal{G}_{tem}	m.tIoU	m.vIoU	vIoU@0.3
✓			44.81%	17.13%	21.08%
	✓		45.56%	17.58%	21.49%
		✓	45.12%	17.53%	21.91%
✓	✓		45.99%	17.72%	22.07%
✓		✓	46.70%	18.07%	22.23%
	✓	✓	47.12%	18.28%	22.52%
✓	✓	✓	47.64%	18.96%	23.19%

region modeling is effective to determine the temporal boundaries of object tubes.

- For spatio-temporal grounding, our STGRN outperforms all baselines on both declarative and interrogative sentences with or without temporal ground truth, which suggests our cross-modal spatio-temporal graph reasoning can effectively capture the object relationships with temporal dynamics and our spatio-temporal localizer can retrieve the object tubes precisely.
- Our STGRN with the dynamic selection method outperforms the STGRN(Greedy) with the greedy method, showing the dynamic smoothness is beneficial to generate high-quality tubes.

5.3. Ablation Study

In this section, we conduct the ablation studies on the spatio-temporal region graph that is the key component of our STGRN. Concretely, the spatio-temporal graph includes the implicit spatial subgraph \mathcal{G}_{imp} , explicit spatial subgraph \mathcal{G}_{exp} and temporal dynamic subgraph \mathcal{G}_{tem} . We selectively discard them to generate ablation models and report all ablation results in Table 3, where we do not distinguish the declarative and interrogative sentences. From these results, we can find that the full model outperforms all ablation models, validating each subgraph is helpful for spatio-temporal video grounding. If only a subgraph is applied, the model with \mathcal{G}_{exp} achieves the best performance, demonstrating the explicit modeling is the most important to capture object relationships. And if two subgraphs are used, the model with \mathcal{G}_{exp} and \mathcal{G}_{tem} outperforms other models, which suggests the spatio-temporal modeling play a crucial role in relation understanding and high-quality video grounding.

Moreover, the layer number T is the essential hyperparameter of the spatio-temporal graph. We investigate the effect of T by varying it from 1 to 5. Figure 3 shows the experimental results on the criteria m.tIoU and m.vIoU for both declarative and interrogative sentences. From the results, we can find our STGRN has the best performance when T is set to 2. The one-layer graph cannot sufficiently capture the object relationships and temporal dynamics. And too many layers may result in region over-smoothing, that is,




Figure 3. Effect of the Number of Spatio-Temporal Graph Convolution Layers.

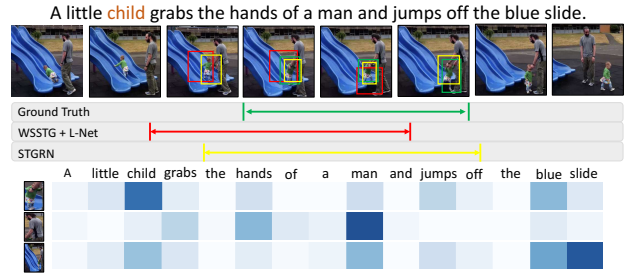


Figure 4. A Example of the Spatio-Temporal Grounding Result.

each region feature tends to be identical. The performance changes on different criteria and sentence types are basically consistent, demonstrating the stable influence of T .

5.4. Qualitative Analysis

We display a typical example in Figure 4. The sentence describes two parallel actions "grab the hands" and "jump off the slide" of the boy in a short-term segment, requiring the accurate spatio-temporal grounding. By intuitive comparison, our STRGN gives a more precise temporal segment and generates a more reasonable spatio-temporal tube than the baseline WSSTG+L-net. Moreover, the attention method in the cross-modal fusion module builds a bridge between visual and textual contents, and we visualize the weights of several key regions over the sentences here. We can see that semantic related region-word pairs have a larger weight, e.g., the region of the boy and the word "child".

6. Conclusion

In this paper, we propose a novel spatio-temporal video grounding task STVG and contribute a large-scale dataset VidSTG. We then design a STGRN to capture region relationships with temporal object dynamics and directly localize the spatio-temporal tubes from the region level.

Acknowledgments This work was supported by Zhejiang Natural Science Foundation LR19F020006 and the National Natural Science Foundation of China under Grant No.61836002, No.U1611461 and No.61751209, Sponsored by China Knowledge Centre for Engineering Sciences and Technology and Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies.

References

- [1] Arun Balajee Vasudevan, Dengxin Dai, and Luc Van Gool. Object referring in videos with language and human gaze. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4129–4138, 2018.
- [2] Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. Temporally grounding natural sentence in video. In *EMNLP*, pages 162–171. ACL, 2018.
- [3] Jingyuan Chen, Lin Ma, Xinpeng Chen, Zequn Jie, and Jiebo Luo. Localizing natural language in videos. In *AAAI*, 2019.
- [4] Zhenfang Chen, Lin Ma, Wenhan Luo, and Kwan-Yee K Wong. Weakly-supervised spatio-temporally grounding natural sentence in video. 2019.
- [5] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS*, 2014.
- [6] Chaorui Deng, Qi Wu, Qingyao Wu, Fuyuan Hu, Fan Lyu, and Mingkui Tan. Visual grounding via accumulated attention. In *CVPR*, pages 7746–7755, 2018.
- [7] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- [8] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. TALL: temporal activity localization via language query. In *ICCV*, pages 5277–5285. IEEE, 2017.
- [9] Jiyang Gao, Zhenheng Yang, Chen Sun, Kan Chen, and Ram Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. 2017.
- [10] Georgia Gkioxari and Jitendra Malik. Finding action tubes. In *CVPR*, pages 759–768, 2015.
- [11] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, pages 5803–5812, 2017.
- [12] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with temporal language. In *EMNLP*, pages 1380–1390. ACL, 2018.
- [13] Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. Modeling relationships in referential expressions with compositional modular networks. In *CVPR*, pages 1115–1124, 2017.
- [14] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In *CVPR*, pages 4555–4564, 2016.
- [15] De-An Huang, Shyamal Buch, Lucio Dery, Animesh Garg, Li Fei-Fei, and Juan Carlos Niebles. Finding "it": Weakly-supervised reference-aware visual grounding in instructional videos. In *CVPR*, June 2018.
- [16] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2016.
- [17] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017.
- [18] Linjie Li, Zhe Gan, Yu Cheng, and Jingjing Liu. Relation-aware graph attention network for visual question answering. In *ICCV*, 2019.
- [19] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014.
- [20] Zhijie Lin, Zhou Zhao, Zhu Zhang, Qi Wang, and Huasheng Liu. Weakly-supervised video moment retrieval via semantic completion network. *AAAI*, 2020.
- [21] Zhijie Lin, Zhou Zhao, Zhu Zhang, Zijian Zhang, and Deng Cai. Moment retrieval via cross-modal interaction networks with query reconstruction. *IEEE Transactions on Image Processing*, 29:3750–3762, 2020.
- [22] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. Attentive moment retrieval in videos. In *SIGIR*, pages 15–24. ACM, 2018.
- [23] Meng Liu, Xiang Wang, Liqiang Nie, Qi Tian, Baoquan Chen, and Tat-Seng Chua. Cross-modal moment localization in videos. In *MM*, pages 843–851. ACM, 2018.
- [24] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *CVPR*, pages 11–20, 2016.
- [25] Niluthpol Chowdhury Mithun, Sujoy Paul, and Amit K Roy-Chowdhury. Weakly supervised video moment retrieval from text queries. In *CVPR*, pages 11592–11601, 2019.
- [26] Varun K Nagaraja, Vlad I Morariu, and Larry S Davis. Modeling context between objects for referring expression understanding. In *ECCV*, pages 792–807. Springer, 2016.
- [27] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015.
- [29] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. Grounding of textual phrases in images by reconstruction. In *ECCV*, pages 817–834. Springer, 2016.
- [30] Xindi Shang, Donglin Di, Junbin Xiao, Yu Cao, Xun Yang, and Tat-Seng Chua. Annotating objects and relations in user-generated videos. In *ICMR*, pages 279–287. ACM, 2019.
- [31] Jing Shi, Jia Xu, Boqing Gong, and Chenliang Xu. Not all frames are equal: Weakly-supervised video grounding with contextual similarity and visual clustering losses. In *CVPR*.
- [32] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2017.
- [33] Weining Wang, Yan Huang, and Liang Wang. Language-driven temporal activity localization: A semantic matching reinforcement learning model. In *CVPR*, pages 334–343, 2019.
- [34] Fanyi Xiao, Leonid Sigal, and Yong Jae Lee. Weakly-supervised visual grounding of phrases with linguistic structures. In *CVPR*, pages 5945–5954, 2017.

- [35] Huijuan Xu, Kun He, L Sigal, S Sclaroff, and K Saenko. Multilevel language and vision integration for text-to-clip retrieval. In *AAAI*, volume 2, page 7, 2019.
- [36] Masataka Yamaguchi, Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Spatio-temporal person retrieval via natural language queries. In *ICCV*, pages 1453–1462, 2017.
- [37] Sibeil Yang, Guanbin Li, and Yizhou Yu. Cross-modal relationship inference for grounding referring expressions. In *CVPR*, pages 4145–4154, 2019.
- [38] Sibeil Yang, Guanbin Li, and Yizhou Yu. Dynamic graph attention for referring expression comprehension. In *ICCV*, pages 4644–4653, 2019.
- [39] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *ECCV*, pages 684–699, 2018.
- [40] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. Mattnet: Modular attention network for referring expression comprehension. In *CVPR*, pages 1307–1315, 2018.
- [41] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *ECCV*, pages 69–85. Springer, 2016.
- [42] Licheng Yu, Hao Tan, Mohit Bansal, and Tamara L Berg. A joint speaker-listener-reinforcer model for referring expressions. In *CVPR*, pages 7282–7290, 2017.
- [43] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *CVPR*, pages 5831–5840, 2018.
- [44] Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S Davis. Man: Moment alignment network for natural language moment retrieval via iterative graph adjustment. In *CVPR*, pages 1247–1257, 2019.
- [45] Hanwang Zhang, Yulei Niu, and Shih-Fu Chang. Grounding referring expressions in images by variational context. In *CVPR*, pages 4158–4166, 2018.
- [46] Zhu Zhang, Zhijie Lin, Zhou Zhao, and Zhenxin Xiao. Cross-modal interaction networks for query-based moment retrieval in videos. In *SIGIR*, 2019.
- [47] Zhu Zhang, Zhou Zhao, Zhijie Lin, Jingkuan Song, and Deng Cai. Localizing unseen activities in video via image query. In *IJCAI*, 2019.
- [48] Luowei Zhou, Nathan Louis, and Jason J Corso. Weakly-supervised video object grounding from text by loss weighting and object interaction. *arXiv preprint arXiv:1805.02834*, 2018.
- [49] Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, and Anton van den Hengel. Parallel attention: A unified framework for visual object discovery through dialogs and queries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4252–4261, 2018.