Relational Graph Learning for
Grounded Video Description Generation

Wenqiao Zhang
Zhejiang University
wenqiaozhang@zju.edu.cn

Xin Eric Wang
University of California, Santa Cruz
xwang366@ucsc.edu

Siliang Tang*
Zhejiang Universety
siliang@zju.edu.cn

Haizhou Shi, Haocheng Shi
Zhejiang University
shihaizhou,hcshi@zju.edu.cn

Jun Xiao,Yueting Zhuang
Zhejiang University
junx,yzhuang@cs.zju.edu.cn

William Yang Wang
University of California, Santa
Barbara
william@cs.ucsb.edu

ABSTRACT

Grounded video description (GVD) encourages captioning models to attend to appropriate video regions (*e.g., objects*) dynamically and generate a description. Such a setting can help explain the decisions of captioning models and prevents the model from hallucinating object words in its description. However, such design mainly focuses on object word generation and thus may ignore fine-grained information and suffer from missing visual concepts. Moreover, relational words (*e.g., "jump left or right"*) are usual spatio-temporal inference results, *i.e.*, these words cannot be grounded on certain spatial regions. To tackle the above limitations, we design a novel relational graph learning framework for GVD, in which a language-refined scene graph representation is designed to explore fine-grained visual concepts. Furthermore, the refined graph can be regarded as relational inductive knowledge to assist captioning models in selecting the relevant information it needs to generate correct words. We validate the effectiveness of our model through automatic metrics and human evaluation, and the results indicate that our approach can generate more fine-grained and accurate description, and it solves the problem of object hallucination to some extent.

CCS CONCEPTS

• **Computing methodologies → Artificial intelligence**; **Scene understanding**; **Computer vision tasks**; **Natural language generation**;

KEYWORDS

grounded video description, language refined scene graph, object hallucination, fine-grained information

*Siliang Tang is the corresponding author.

1 INTRODUCTION

Recently, video captioning [11], the task of automatically generating a sequence of natural-language words to describe a video, has drawn increasing attention [19, 20, 33]. However, these models are known to have poor grounding performance, which leads to *objects hallucination* [23]. Captioning models generate descriptive objects that are not in the video because of similar semantic contexts or pre-extracted priors during the training stage. As a result, some models lack interpretability even if they have high captioning scores. Therefore, grounded video description (GVD) [53], which tries to improve the grounding performance of captioning models, has been proposed. In this approach, the captioning model learns to ground related video regions (*e.g., objects*) that are used as input to predict the next word. Such a setting can teach models to rely explicitly on the corresponding evidence in the video when generating descriptions.

Focusing on object word generation can only partially address the problem of object hallucination. Furthermore, this design may lead to some potential limitations. First, a GVD model encourages captioning models to focus on related regions of objects to generate correct words; thus, the fine details (*e.g., related objects and attributes of grounded object*) in the corresponding video may be ignored, and may generate a coarse-grained descriptive sentence. Second, relational words (relationships among or between salient objects), such as "*climb up or down*", are usual inference results from sequential frames. No specific spatial regions can ground these words, thus, using only grounding regions may lead to inappropriate word generation. Therefore, an optimization goal of GVD is to produce fine-grained information and a reasonable learning approach to describing a video with fine and correct details.

Inspired by recent breakthroughs in higher-level visual understanding tasks, scene graph ($\mathcal{SG}$) construction has become a popular area with remarkable advancements. [5, 6]. As an abstraction of objects and their complex relationships, $\mathcal{SG}$ can provide rich semantic information of the visual domain. Thus, $\mathcal{SG}$ is expected to deal with the problems mentioned above. On the one hand,
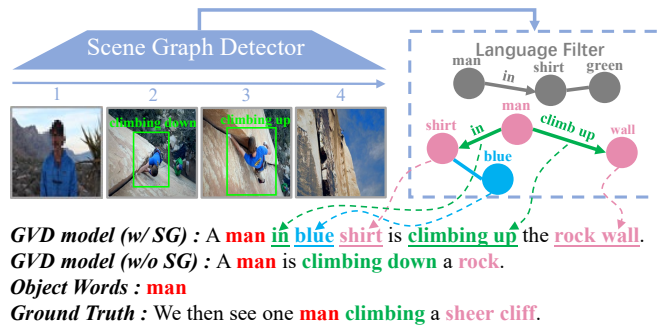
Figure 1: An example of how the language-refined graph facilitates GVD generation. The green bounding box is the attended region. Pink, green and blue words represent the objects, relationships and attributes.

$\mathcal{SG}$ provides complementary information to assist the captioning model in generating fine-grained phrases, such as "*man in blue shirt*" (Figure 1). On the other hand, $\mathcal{SG}$ with relational inductive knowledge guides the captioning model to generate appropriate relational words, such as "*climb up*", and ground the correct region in the third frame. While w/o $\mathcal{SG}$ grounding a similar region in the second frame but generate inappropriate relational words "*climb down*". However, directly using $\mathcal{SG}$ does not solve the problem of semantic inconsistency between predefined concept categories and target lexical words. For example, in Figure 1, ⟨ *man-in-shirt* ⟩ and ⟨ *shirt-green* ⟩ are the unconcerned concepts. We notice that a sentence (ground truth of video description) can also parse a language $\mathcal{SG}$ [1] that contains the key visual concepts. The language $\mathcal{SG}$ can be treated as linguistic guidance to refine the visual $\mathcal{SG}$. Exactly as shown in Figure 1, the unconcerned concepts with color gray are filtered and do not hallucinate objects in the description.h

Based on this insight, we propose the relational graph learning framework (RGL), which incorporates $\mathcal{SG}$ into the conventional encoder-decoder method for GVD generation. Specifically, given a video, the training pipeline has three parts: 1) *Relational Graph Encoder*. We first build the frame scene graph $\mathcal{SG}^{\mathcal{F}}$ and the language scene graph $\mathcal{SG}^{\mathcal{L}}$ from the video frames and its ground truth at the training stage. Then, to filter the unconcerned concepts, we regard the $\mathcal{SG}^{\mathcal{L}}$ as linguistic guidance to refine the $\mathcal{SG}^{\mathcal{F}}$. Thus, the refined graph representation $\mathcal{SG}^{\mathcal{R}}$ with relational inductive knowledge is obtained for improved GVD generation. 2) *Sentence Decoder*. We introduce a selection mechanism for a sentence generator that learns to decide the utilization of grounding regions and refined graph $\mathcal{SG}^{\mathcal{R}}$. Such a decoder can generate video descriptions reasonably on a fine-grained level. In addition, we derive a context generator with refined graph $\mathcal{SG}^{\mathcal{R}}$ to leverage and update the compressed semantic information and maintain coherence among sentences. 3) *Grounding Module*. We build a grounding and localizing mechanism, which not only encourages the model to ground the regions dynamically on the basis of the current semantic context to predict words but also localizes regions using the generated object words. Such a setting can boost the accuracy of the object word generation and address the problem of object hallucination. In summary, the major contributions of our study are as follows:

- We develop a language-refined $\mathcal{SG}$ representation that contains the key visual concepts for GVD generation.
- We propose a captioning model with a selection mechanism that selects the relevant information it needs on the basis of the current semantic context to describe a video reasonably.
- We demonstrate the superiority of RGL in generating fine-grained and accurate description via automatic metrics and human evaluation.
- This attempt is the first to combine GVD and visual $\mathcal{SG}$ organically. The proposed RGL can improve captioning quality, grounding performance and alleviate object hallucination simultaneously.

2 RELATED WORK

Video Captioning: Video captioning is being actively studied in vision and language research. The prevailing video captioning techniques often incorporate the encoder-decoder pipeline inspired by the first successful sequence-to-sequence model S2VT [30]. Benefitting from the rapid development of deep learning, video captioning models have achieved remarkable advances using attention mechanism [27, 37, 43, 52], memory networks [3, 15, 22, 31], reinforcement learning [14, 21, 33] and generative adversarial networks [20, 40]. Although these encoder-decoder-based methods have reached impressive performance on automatic metrics, they often neglect how well the generated caption words (e.g., objects) are grounded in the video, making models less explainable and trustworthy.

Visual Grounding: Visual grounding models encourage captioning generators to link phrases with specific spatial regions of images or videos, thereby presenting a potential way to improve the explainability of models [7, 24, 35, 46, 49, 53]. The most common way of grounding models is to predict the next word using an attention mechanism, which is deployed over noun phrases, with supervised bounding boxes as input. However, these models often focus on object word generation and further produce a coarse-grained sentence to describe an image or a video. Moreover, grounding video regions seem inappropriate to predict relational words because these words do not usually correspond to a specific spatial region.

Scene Graph: Recently, $\mathcal{SG}$ construction have become popular research topics with significant advancements [9, 36, 38, 47, 48, 55] based on the Visual Genome [12] dataset. The $\mathcal{SG}$s contain structured semantic information and can represent scenes as directed graphs, where nodes are objects and edges are relationships and attributes. Using this inductive information is natural to improve the performance of vision-language tasks, e.g., image captioning [8, 39], VQA [26, 28]. However, directly feeding the $\mathcal{SG}$ to the captioning model may lead to a noncorrespondence problem between vision and language. Therefore, how to refine the relational knowledge is key to promote the vision-language field further.

Graph Convolutional Network: A graph convolutional network (GCN) [2] is a weighted average function that operates directly on a graph and induces the embedding vectors of nodes on the basis of the properties of their neighborhoods. Graph neural networks have achieved remarkable success in processing graph-structured data and have been widely adopted in many areas, such as semantic role labeling [17, 18], relation classification [32, 54], social network mining [4], text classification [41, 42], recommendation system [44],
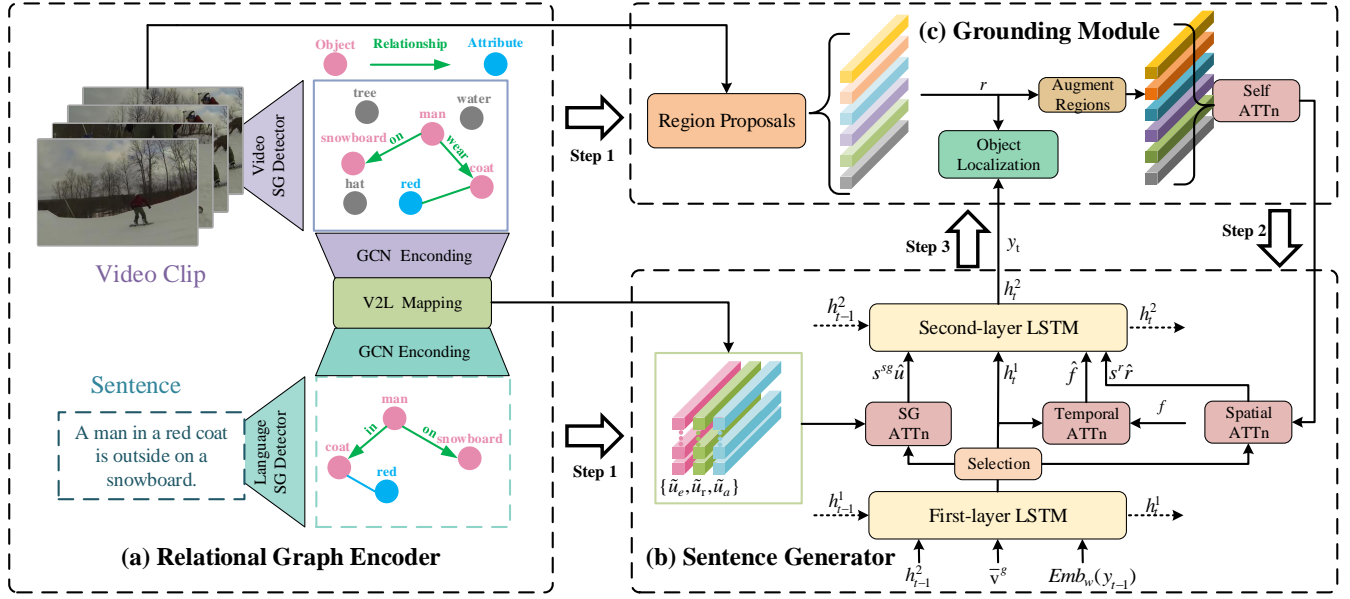
Figure 2: Overview of our method. It consists of three modules: (a) The relational graph encoder to produce the language-refined graph representation ũ from the source video and its ground truth. (b) The sentence generator uses a stacked two-layer LSTM with a selection mechanism for video description generation. (3) The grounding module has the mechanism of region grounding and object localization to predict the correct object words.

and scene understanding [38]. Given the effectiveness of GCN in handling graph-based data, we use a GCN to capture the contextual relations among the $\mathcal{SG}$s of a video for GVD generation.

3 METHOD

3.1 Task Description

Before presenting our method, we first introduce some basic notions and terminologies. Given a video $\boldsymbol{V}=\{\boldsymbol{V}_1,\cdots,\boldsymbol{V}_L\}$ consisting of L clips, the goal of grounded video description (GVD) is to generate the natural language description $\boldsymbol{S}=\{\boldsymbol{S}_1,\cdots,\boldsymbol{S}_L\}$, and localize the object words $\boldsymbol{O}=\{\boldsymbol{O}_1,\cdots,\boldsymbol{O}_K\}$ in frames, where K is the number of objects appearing in the description. We denote the model as G, and the sentence $\hat{\boldsymbol{S}}$, object grounding $\hat{\boldsymbol{O}}^g$ generated by $G(\boldsymbol{V})$, i.e., $\hat{\boldsymbol{S}}$, $\hat{\boldsymbol{O}}^g = G(\boldsymbol{V})$. For a model parameterized by θ, we define the loss for a training pair as $\mathcal{L}((\boldsymbol{S},\boldsymbol{O}^g), G(\boldsymbol{V};\Theta))$.

Figure 2 shows the overall pipeline of our Relational Graph Learning (RGL) approach, which mainly consists of the *Relational Graph Encoder*, *Sentence Decoder* and *Grounding Module*. (1) Given a video and its corresponding ground truth, the *Relational Graph Encoder* module first transforms them into the $\mathcal{SG}$s. Then adapting the visual-language mapping is adopted to generate a language-refined graph representation $\tilde{\mathbf{u}}$ (Section 3.2). (2) The *Sentence Decoder* is encouraged to select information dynamically from $\tilde{\mathbf{u}}$ and ground regions $\mathbf{r}$ to generate the current word. (Section 3.3). (3) We compute the localization accuracy of object words on the ground truth sentence, and the language model dynamically attends region proposals $\mathbf{r}$ in the subsequent word prediction stage. (Section 3.4).For clarity, we use a GVD task as an example to illustrate our method.

3.2 Relational Graph Encoder

Generally, the $\mathcal{SG}$ defined in our task $\mathcal{SG} = (\mathcal{N}, \mathcal{E})$ and contains a set of nodes $\mathcal{N}$ and edges $\mathcal{E}$. As exemplified in Figure 2(a), the node set $\mathcal{N}$ contains three types of nodes: object node o, attribute node a, and relationship node r.

3.2.1 Scene Graph Detector. Given a video clip, we sample several frames to generate the *Frame Scene Graph $\mathcal{SG}^{\mathcal{F}}$*. In detail, for i^{th} frame, the corresponding $\mathcal{SG}_i^{\mathcal{F}}$ is extracted by the $\mathcal{SG}$ parser comprising an object detector Faster-RCNN [34], an attribute classifier [39], and a relationship classifier MOTIFS [47]. Hence, $\mathcal{SG}^{\mathcal{F}} = \{\mathcal{SG}_1^{\mathcal{F}},\cdots,\mathcal{SG}_Q^{\mathcal{F}}\}$, where Q is number of sampled frames. To generate the *Language Scene Graph $\mathcal{SG}^{\mathcal{L}}$*, we use the $\mathcal{SG}$ generator [1] that parses the sentence to a syntactic dependency tree. Then, the rule-based method [25] is developed to transform the tree into an $\mathcal{SG}$.

3.2.2 Scene Graph Encoder. To represent the nodes of the $\mathcal{SG}$, we denote the nodes $\{o, a, r\}$ by the label embeddings as $\{\mathbf{e}^o, \mathbf{e}^a, \mathbf{e}^r\}$ $\in \mathbb{R}^e$, corresponding respectively to objects, attributes and relationships both in $\mathcal{SG}^{\mathcal{F}}$ and $\mathcal{SG}^{\mathcal{L}}$. In particular, different from using the simple label embeddings $\mathbf{e}^o$ and $\mathbf{e}^r$ in $\mathcal{SG}^{\mathcal{F}}$, $\mathbf{e}^o$ is the key to connect domains of vision and language. Thus, we introduce the Multi-modal Factorized Bilinear Pooling (MFB) [45] to fuse the region features and label embeddings to augment object representation $\mathbf{e}^o$, which is known to be effective in multi-modal tasks [13, 16].

Node embedding : To encode the $\mathcal{SG}$ nodes at a unified representation $\mathbf{u} = \{\mathbf{u}^o, \mathbf{u}^a, \mathbf{u}^r\} \in \mathbb{R}^u$, we introduce the Graph Convolutional Network (*GCN*) [18], which can embed the graph structure

into vector representations. Thus, we use the GCN that encodes three kinds of node embeddings by considering their neighborhood information. All the *GCNs* are defined with the same structure but independent parameters.

Objects Encoding : An object o_j in $\mathcal{SG}$, it can play different roles ("*subject*" or "*object*") due to different edge directions, *i.e.*, two triples with different relationships, $\langle o_i - r_{i,j} - o_j \rangle$ and $\langle o_j - r_{j,k} - o_k \rangle$. Such associated objects by the cascaded encoding scheme can represent the global information of objects for a frame or sentence. Therefore we compute $\mathbf{u}^{o_j}$ via explicit role modeling:

$$\mathbf{u}^{o_j} = \frac{1}{N_{o_j}} \Big[\sum_{\mathbf{e}^{o_j} \in sub} G_s(\mathbf{e}^{o_j}, \mathbf{e}^{r_{j,k}}, \mathbf{e}^{o_k}) \\ + \sum_{\mathbf{e}^{o_j} \in obj} G_o(\mathbf{e}^{o_i}, \mathbf{e}^{r_{i,j}}, \mathbf{e}^{o_j}) \Big] \tag{1}$$

where $N_{o_j} = \big|sub(o_j)\big| + \big|obj(o_j)\big|$ is the total number of the relationship triplets in $\mathcal{SG}$ that object o_j has. G_s and G_o are the convolutional operation for objects as a "*subject*" or an "*object*".
Attributes Encoding : An o_i in $\mathcal{SG}$ usually includes several attributes $\{a_1^i, \cdots, a_{T_{a_i}}^i\}$, where T_{a_i} is the total number of attributes. Therefore, the unified $\mathbf{u}^{a_i}$ can be computed as:

$$\mathbf{u}^{a_i} = \frac{1}{T_{a_i}} \sum_{i \in T_{a_i}} G_a(\mathbf{e}^{o_i}, \mathbf{e}^{a_i}) \tag{2}$$

where G_a is the convolutional operation for object o_i and its attributes.
Relationships Encoding : The relationship between two salient objects o_i and o_j is given by the triplet $\langle o_i - r_{i,j} - o_j \rangle$. Similarly, the unified relationship encoding $\mathbf{u}^{r_{i,j}}$ is produced as follows:

$$\mathbf{u}^{r_{i,j}} = G_{i,j}^r(\mathbf{e}^{o_i}, \mathbf{e}^{r_{i,j}}, \mathbf{e}^{o_j}) \tag{3}$$

where $G_{i,j}^r$ is the convolutional operation for relational object o_i and o_j.

3.2.3 Visual-Language Mapping. To refine the visual $\mathcal{SG}^{\mathcal{F}}$, we adapt the visual domain to the language domain using the across-modality mapping function. In our model, we translate $\mathbf{u}^{\mathcal{F}}$ (video) on the basis of linguistic guidance $\mathbf{u}^{\mathcal{L}}$ (sentence) to a refined $\mathcal{SG}$ representation $\tilde{\mathbf{u}} = \{\tilde{\mathbf{u}}^o, \tilde{\mathbf{u}}^a, \tilde{\mathbf{u}}^r\}$ using an effective multi-layer perceptron. The visual-language mapping loss $\mathcal{L}(M)$ is computed by mean squared error (MSE) loss as follows:

$$\mathcal{L}(M) = MSE(\mathbf{u}^{\mathcal{F}}, \mathbf{u}^{\mathcal{L}}) \tag{4}$$

Such language-refined $\mathcal{SG}$ $\tilde{\mathbf{u}}$ is then fed into the sentence decoder, providing relational inductive knowledge to generate GVD.

3.3 Sentence Decoder

3.3.1 Single Sentence Generator. We develop a stacked two-layer LSTM network using a selection mechanism to generate video description. The first layer of the LSTM network is for encoding the video features. Specifically, we concatenate the previous hidden state $\mathbf{h}_{t-1}^2$, global video representation $\bar{\mathbf{v}}^g$, and embed the previously generated word $\mathbf{y}_{t-1}$ as the source input to the first-layer LSTM:

$$\mathbf{h}_t^1 = LSTM_{1st}(\mathbf{h}_{t-1}^1, [\mathbf{h}_{t-1}^2, \bar{\mathbf{v}}^g, Emb_w(\mathbf{y}_{t-1})]) \tag{5}$$

Figure 3: Overview of the context generator. The final hidden states in sentence generator are fed into context generator, whose output and $\mathcal{SG}$ representation $\tilde{\mathbf{u}}$ are return to initialize the next sentence generator.

where $[\cdot, \cdot]$ indicates the row-wise concatenation. Emb_w is the word embedding matrix. $\bar{\mathbf{v}}^g$ contains the pooling result of $\mathbf{v}^g$ and the positional information of current video clip.

The second layer of the LSTM network is for sentence generation. To improve the robustness of the sentence generator, we enrich the language input as four parts instead of simply utilizing $\mathbf{h}_t^1$ as the source input. We first extract the frame features $\mathbf{f} = \{\mathbf{f}_1, \cdots, \mathbf{f}_Q\}$ and design a temporal attention [37, 50, 51] to determine the weight for each frame. The attended frame representation is given by $\hat{\mathbf{f}}$.

We augment region proposals $\mathbf{r} = \{\mathbf{r}_1, \cdots, \mathbf{r}_N\}$ using its positional encoding and region-class similarity matrix. To model the relations between regions further, we deploy a self-attention [29] layer for $\mathbf{r}$, thereby allowing information to pass across all regions in the sample frames. Thus, the augmented region representation is given by $\tilde{\mathbf{r}}$.
Selection Mechanism : The hidden state $\mathbf{h}_t^1$ is determined to select dynamically *how much* and *what* from the $\mathcal{SG}$ and grounding regions. (1) *how much*, a soft attention is developed on $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{r}}$, and the corresponding weight score s^r and s^{sg} are determined to select dynamically *how much* from $\mathcal{SG}$ and grounding regions. (2) *what*, then we compute the spatial attention a^r for $\{\tilde{\mathbf{r}}_1, \cdots, \tilde{\mathbf{r}}_N\}$ and soft attention a^{sg} for $\{\tilde{\mathbf{u}}^o, \tilde{\mathbf{u}}^a, \tilde{\mathbf{u}}^r\}$ to generate attended representation $\hat{\mathbf{r}}$ and $\hat{\mathbf{u}}$. a^r and a^{sg} are determined to select *what* in $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{r}}$. Lastly, we concatenate $\mathbf{h}_{t-1}^1$, $\hat{\mathbf{f}}$, $s^{sg}\hat{\mathbf{u}}$ and $s^r\tilde{\mathbf{r}}$ as the input to the second layer of the LSTM network:

$$\mathbf{h}_t^2 = LSTM_{2nd}(\mathbf{h}_{t-1}^2, [\mathbf{h}_t^1, \hat{\mathbf{f}}, s^{sg}\hat{\mathbf{u}}, s^r\hat{\mathbf{r}}]) \tag{6}$$

Using the notation $\mathbf{y}_{1:T}$, we refer to a sequence of words $\{\mathbf{y}_1, \cdots, \mathbf{y}_T\}$. For each step t, the conditional distribution over possible words is given by:

$$p(\mathbf{y}_t|\mathbf{y}_{1:t-1}, \mathbf{r}, \mathbf{f}, \tilde{\mathbf{u}}) = Softmax(\mathbf{W}_{hy}^{\top}\mathbf{h}_t^2) \tag{7}$$

where $\mathbf{W}_{hy} \in \mathbb{R}^{m \times v}$, $\mathbf{h}_t^2$ is the hidden state from the second layer of the LSTM network .

3.3.2 Context Generator. The clips belong to the same video and are thus contextually co-dependent, *i.e.*, the previous sentences generated by the single sentence generator are semantically relevant to generate the next sentence. To exploit contextual dependencies among the sentences, we build a superordinate context generator that using a LSTM runs over the embedding of each sentence (Figure 3), and is therefore asynchronous to the sentence generator. Each time the state of sentence generator is updated, its output is

utilized to initialize the sentence generator of the first layer LSTM:

$$\mathbf{x}_0^i = I_{NI}([LSTM_c(\mathbf{h}_M^{2,i-1}, \mathbf{h}_{t-1}^c), \tilde{\mathbf{u}}^i]) \tag{8}$$

where $\mathbf{x}_0^i$ is the initial input of i^{th} sentence generator. I_{NI} is a dense layer with an activation function, $\mathbf{h}_M^{2,i-1}$ is the last hidden state of the i-1^{th} sentence, and $\mathbf{h}_{t-1}^c$ is the previous hidden state of context LSTM.

In summary, combining visual-language mapping loss $\mathcal{L}(M)$, the cross-entropy loss $\mathcal{L}(S)$ for description generation can be computed by:

$$\mathcal{L}(S) = -\sum_{i=0}^{M} log(p(\mathbf{y}_t|\mathbf{y}_{0:t-1})) + \lambda_M \mathcal{L}(M) \tag{9}$$

3.4 Grounding Module

Region Grounding : In this section, we aim to evaluate how well the captioning model grounds visual objects. To assist the language model in attending to the correct regions, following [53], we develop the region attention loss $\mathcal{L}(R)$: we denote the indicators of positive/negative regions as $\gamma = \{\gamma_1, \cdots, \gamma_N\}$ in each time step, where $\gamma_i = 1$ if the region γ_i has over 0.5 IoU with a ground truth box and otherwise 0. In combination with the treating attention a^r (Section 3.3.1), the region attention loss function is defined as:

$$\mathcal{L}(R) = -\sum_{1}^{N} \gamma_i log a_i^r \tag{10}$$

Object Localizaiton : Given an object word $\mathbf{w}$ with a specific class label, we aim to localize the related region proposals. We first define the region-class similarity function with the treating attention weights a^r as below:

$$p^s(\mathbf{r}, a^r) = \text{Softmax}(\mathbf{W}_s^\top \mathbf{r} + a^r) \tag{11}$$

where $W_s \in \mathbb{R}^{d \times N}$ is a simple object classifier to estimate the class probability distribution.

Thus we use the $p^s(\mathbf{r}, a^r)$ to calculate the confidence score for $\mathbf{w}$, combing the supervised attention loss $\mathcal{L}(R)$, the grounding loss function $\mathcal{L}(G)$ for word $\mathbf{w}$ is defined as:

$$\mathcal{L}(G) = -\lambda_L \sum_{i=1}^{N} \gamma_i log p^s(\mathbf{r}_i) + \lambda_R \mathcal{L}(R) \tag{12}$$

3.5 Training Algorithm

Algorithm 1 details the pseudocode of our RGL algorithm for GVD generation. First, we pre-train the language GCN by reconstructing the sentences from the latent vector that is encoded from the language $\mathcal{SG}$ with the GCN (sentence $\rightarrow \mathcal{SG} \rightarrow$ latent vector $\rightarrow$ sentence $\rightarrow$ object grounding). Then, we keep the language GCN fixed, and learn the visual GCN by mapping the visual $\mathcal{SG}$ and the language $\mathcal{SG}$ in the latent space. Basically, the encoded latent vector from the language $\mathcal{SG}$ is used as supervised signals to learn the visual GCN and assist the GVD generation. Through the learning procedure, we found that most of the unconcerned visual concepts are filtered.

Algorithm 1: Relational Graph Learning Algorithm

Input: Training pairs $(V, (S, O^g))$.
Output: $(\hat{S}, \hat{O}^g)$.

1 **Initialization**: Load the pre-trained $\mathcal{SG}$ generator, and generate $\mathcal{SG}^{\mathcal{F}}$ and $\mathcal{SG}^{\mathcal{L}}$;

2 **repeat**

3   **if** *Language GCN have not trained* **then**

4   **repeat**

5   Randomly sample a minibatch;

6             $u^{\mathcal{L}} \Longleftarrow \mathcal{SG}^{\mathcal{L}}$ by using Equation (1)(2)(3);

7             $\bar{S} \Longleftarrow u^{\mathcal{L}}$ by using Equation (9);

8             $\bar{O}^g \Longleftarrow \bar{S}$ by using Equation (10) (12);

9   Update Θ, and minimizing $\mathcal{L}(S), \mathcal{L}(R), \mathcal{L}(G)$;

10   **until** *Convergence;*;

11   Fix the Language GCN;

12   Randomly sample a minibatch;

13     $u^{\mathcal{F}} \Longleftarrow \mathcal{SG}^{\mathcal{F}}$ by using Equation (1)(2)(3);

14   Non-linearly map $u^{\mathcal{F}}$ into $\tilde{u}$ by using Equation (4);

15     $\hat{S} \Longleftarrow \tilde{u}$ by using Equation (9);

16     $\hat{O}^g \Longleftarrow \hat{S}$ by using Equation (10) (12);

17   Update Θ, and minimizing $\mathcal{L}(M), \mathcal{L}(S), \mathcal{L}(R), \mathcal{L}(G)$;

18 **until** *Convergence*;

---

4 EXPERIMENTS

4.1 Dataset and Setting

Dataset : We benchmark the RGL method on the ActivityNet-Entities dataset[1] and compare it with other baselines and state-of-the-art models. Moreover, the ActivityNet-Entities collect NPs with bounding box annotations in the frame region level, every NP only in one frame inside each clip. It defines the 4-D tuple (V, D, E, B) that represents the number of videos, descriptions, objects, and bounding boxes, and the official split has *(10k ,35k, 432 ,105k)/,(2.5k, 8.6k, 427, 26.5k)/,(2.5k ,8.5k, 421, 26.1k)* for train/val/test data, respectively.

Settings : For video description, we tokenized the texts on white space, and the sentences are "cut" at a maximum length of 20 words. All the Arabic numerals are converted to the English word. We add a special *Unknown token* to replace the words out of the vocabulary list. The vocabulary has 4,904 words, and each word represents as a 512-dimensional vector. The hidden state size of the LSTM is 1024. The embedding size of nodes in the scene graph and the unified scene graph representation is 1000. The dimension of other learnable matrices are $\mathbf{W}_{hy} \in 512 \times 4904$, $\mathbf{W}_{hy} \in 2048 \times 1024$. We train our model with cross-entropy objectives and use the ADAM optimizer [10] with a learning rate of 1.25e-4, the batch size is fixed to 64 when training all models. All experiments were run on the 2080Ti GPUs. The training is limited to 40 epochs and the model with the best validation CIDEr score is selected for testing.

---

[1] https://github.com/facebookresearch/ActivityNet-Entities

4.2 Metrics

Captioning Evaluation: We use the performance evaluation tool [2] provided by the 2018 ActivityNet Captions Challenge, which includes *BLEU@1*, *BLEU@4*, *METEOR*, *CIDEr*, and *SPICE* to evaluate the results of video captioning.

Grounding Evaluation: Following the grounding evaluation from GVD [53] on the generated sentences, we define the number of object words as A, the number of correctly predicted object words as B, and the number of correctly predicted and localized words as C. A region prediction is considered correct if the object word is correctly predicted and also correctly localized (*i.e.*, IoU with ground truth box >0.5). Thus, we compute two versions of precision Fl_{ALL} and Fl_{LOC} to evaluate the object localization accuracy for attention.

$$Fl_{ALL} = \frac{C}{A}, \qquad Fl_{LOC} = \frac{C}{B} \qquad (13)$$

During model training, we restrict the grounding region candidates within the target frame (w/ GT box), i.e., only consider the N proposals on the frame f with the GT box. We also compute the localization accuracy *GRD*. and attention correctness *ATT*. at each annotated object word.

Hallucination Evaluation: Some related object words have not been annotated in the ActivityNet-Entities dataset. Thus, we relabel five related object words for 100 video clips (randomly selected) on the basis of visual content and its ground truth. According to [23], we compute $CHAIR_i$ and $CHAIR_s$, which indicate the fraction of object instances that are hallucinated per-instance and per-sentence, respectively. Besides, we also record the recall of object word prediction $RECALL_o$.

Human Evaluation: Since our model can describe a video in more fine-grained details, and some of these details (e.g., Figure 4.a) don't exist in the ground truth. Thus, evaluating the grounding performance of object/relational/attribute words seems unfair. To verify the grounding performance of our model further, we evaluate the description quality through human judgment. On the one hand, the human evaluation validates the grounded performance of object/relational/attribute words by $Relevant_{obj}$, $Relevant_{rel}$ and $Relevant_{att}$. On the other hand, we also allow humans to evaluate the coherence of the generated descriptions subjectively using $Performance_{coh}$.

4.3 Baseline and SoTA

We compare the proposed RGL algorithm with the existing SoTA method GVD [53] [3] and baseline ST-LSTM on the ActivityNet-Entities dataset. We also conduct an ablation study to investigate the contributions of individual components in RGL. In our experiment, we train the following baselines of ST-LSTM and four variants of RGL:

- **ST-LSTM**, which is a captioning model with a two-layer LSTM and spatiotemporal attention and a simplified version of our model.
- **RGL (w/o SG)**, which only uses video representation to generate GVD.

- **RGL (w/o CG)**, which generates captions without a context generator and merely uses a stack LSTM as its sentence generator.
- **RGL (w/o RA)**, which generates captions without regional attention to study the importance of regional supervision.
- **RGL (w/o OG)**, which generates captions sequentially without object word grounding and is similar to a standard video captioning algorithm.

4.4 Experimental Results

Captioning Results : Table 1 shows the overall qualitative results of our model and SoTA on the ActivityNet-Entities dataset. Generally, RGL achieves the best performance on all the metrics in comparison with SoTA. We obtain an improvement of 0.7 on SPICE because our method learns the refined representation of $\mathcal{SG}$, which provides relational knowledge and positional semantic prior, to improve this score. It is noteworthy that the RGL (w/o OG) achieves almost all the best captioning scores, this is reasonable, because without grounding operation, the captioning model may pay more attention to the description generation. Moreover, compared with other variants of our method, the performance of RGL (w/o SG) is relatively poor, thus validating the effectiveness of refined $\mathcal{SG}$ for video description generation.

Grounding Results : Table 1 also reveals that our method effectively improves the accuracy of GRD., ATT., $F1_{ALL}$ and $F1_{LOC}$. According to our observation, RGL almost obtains the best performance for all of the grounding metrics. Since the $\mathcal{SG}$ representation can be regarded as relational inductive knowledge, the captioning model can dynamically select relevant information from grounding correct regions and refined $\mathcal{SG}$ to generate correct words.

Hallucination Results : Table 2 presents object hallucination on the test set. We note an interesting phenomenon that the $\mathcal{SG}$ based methods tend to perform considerably better on the $CHAIR_i$, $CHAIR_s$ and $RECALL_o$ metrics than methods without $\mathcal{SG}$ by a large margin. This finding proves that the $\mathcal{SG}$ contains refined visual concepts that can help the captioning model to capture more correct objects in a video. On this basis, object hallucination is decreased. Therefore, the refined language $\mathcal{SG}$ can assist the region grounding operation in generating correct object words.

Human Evaluation : As commonly known, the text-matching-based metrics are not perfect, and some descriptions with lower scores actually depict the videos more accurately (See related cases in a later section of Qualitative Examples), i.e., some captioning models can describe a video in detail but with low captioning scores. Thus, human evaluation is conducted. We show five experienced workers the descriptions generated by RGL and RGL (w/o SG) and asked them which one is more coherent and captures more related visual concepts. For each pairwise comparison, a total of 100 video clips (same videos in hallucination evaluation) are randomly extracted for comparison. The results of the comparisons are shown in Table 3. RGL considerably outperforms RGL (w/o SG) in terms of $Relevant_{obj}$, $Relevant_{rel}$ and $Relevant_{att}$ by a large margin. These results indicate that the model with $\mathcal{SG}$ can generate more relevant video description in fine detail. The results also prove the effectiveness of RGL for GVD generation. Meanwhile, the sentence produced by RGL achieves a higher score in terms of coherent

Method	RSG.	Captioning Evaluation					Grounding Evaluation			
		BLEU@1	BLEU@4	METEOR	CIDEr	SPICE	GRD.	ATT.	$F1_{ALL}$	$F1_{LOC}$
GVD (w/o SA) [53]*		23.1	2.16	10.8	44.9	14.9	22.3	16.1	3.73	11.7
GVD [53]*		23.6	2.35	**11.0**	45.5	14.7	44.9	35.7	7.10	23.8
ST-LSTM[†]		22.7	2.13	10.3	42.5	11.6	-	-	-	-
RGL (w/o SG)[†]		23.4	2.30	10.7	45.9	14.5	45.0	35.9	7.11	23.4
RGL (w/o RA)[†]	✓	25.0	2.46	10.7	45.1	15.1	44.3	35.1	5.29	14.8
RGL (w/o OG)[†]	✓	25.4	**2.60**	**11.1**	**47.4**	**15.4**	**45.8**	**36.7**	-	-
RGL (w/o CG)[†]	✓	25.2	**2.60**	10.8	46.9	**15.4**	45.3	36.2	7.69	26.2
RGL[†]	✓	**25.5**	2.59	11.0	47.2	**15.4**	45.6	36.6	**7.70**	**26.4**

Table 1: Captioning results and grounding results on ANet-Entities test set. BLEU-1/4, METEOR, CIDEr and SPICE are used as captioning metrics. Grd., ATT, $F1_{ALL}$ and $F1_{LOC}$ are used as grounding metrics. RSG. indicates the language refined scene graph is used or not for GVD generation. ∗ indicates the results are obtained from the original papers, [†] sentences obtained directly from the author. Larger value indicates better performance. All accuracies are in %. Top one score on each metric are in bold. Acronym notations of each method see in Sec 4.3.

Method	RSG.	Hallucination Evaluation		
		$CHAIR_i$	$CHAIR_s$	$RECALL_o$
ST-LSTM[†]		0.686	0.894	0.147
RGL (w/o SG)[†]		0.572	0.824	0.198
RGL (w/o RA)[†]	✓	0.659	0.867	0.164
RGL (w/o OG)[†]	✓	0.609	0.836	0.206
RGL (w/o CG)[†]	✓	0.531	0.807	0.228
RGL[†]	✓	**0.513**	**0.795**	**0.234**

Table 2: Hallucination results on ANet-Entities test set. $CHAIR_i$, $CHAIR_s$ and ACC_o are used as hallucination metrics. Lower value of $CHAIR_i$, $CHAIR_s$ indicates better performance, while the large value of ACC_o is better.

Metric	Human Evaluation		
	RGL is Better	RGL is Worse	Equal
$Relevant_{obj}$	0.40 (+9%)	0.31	0.29
$Relevant_{rel}$	0.43 (+14%)	0.29	0.28
$Relevant_{att}$	0.39 (+16%)	0.23	0.38
$Performance_{coh}$	0.30 (+6%)	0.24	0.46

Table 3: Human evaluation of our model and its variant RGL (w/o SG). $Relevant_{obj}$, $Relevant_{rel}$ and $Relevant_{att}$ indicate the relevant visual concepts of objects, relationships and attributes in generated sentence. $Performance_{coh}$ is to evaluate the coherence in sentence.

evaluation. This result is obtained probably due to the structured semantic information of $\mathcal{SG}$.

Analysis of Scene Graph Refinement : To justify the contribution of $\mathcal{SG}$ refinement to GVD generation, we investigate the importance of the $\mathcal{SG}$ with language guidance using three different $\mathcal{SG}$. As we can see in Figure 4, the best performance is obtained when the language $\mathcal{SG}$ is used directly. This result is reasonable because the language $\mathcal{SG}$ contains key language concepts that can naturally reconstruct its video description. In addition, Refined $\mathcal{SG}$ considerably outperforms visual $\mathcal{SG}$ by a large margin on all important metrics, especially the CHAIRi (-9.0). The experimental results indicate the importance of $\mathcal{SG}$ refinement in GVD generation.

Qualitative Examples : To evaluate the quality of video description, we conduct qualitative analysis to compare the RGL and RGL (w/o SG) in Figure 6. We present the $\mathcal{SG}$ for each video, with the annotated bounding boxes as object words. From these exemplary results, RGL produces a more fine-grained caption compared with RGL (w/o SG). As shown in case (a), RGL generates the related objects "*woman, cat, shirt*", their relationships "*in, grab*", and attribute "*blue*" in accordance with the SG, even if the object words only have "*lady*". Moreover, the relational phrase "*next to*" in the $\mathcal{SG}$ is converted to "*grab*" through language refinement. However, the sentence generated by RGL (w/o SG) only gives the object word

Figure 4: The major captioning (CIDEr), grounding (GRD.) and hallucination ($CHAIR_i$) results of using different scene graphs. Visual SG: using video scene graph. Language SG: using the language scene graph. Refined SG: using the language refined scene graph. To facilitate the comparison, the value of $CHAIR_i$ is multiplied by 100.

"*woman*" to describe the video in a coarse-grained manner, even if it has a higher captioning score and a higher similarity with the ground truth. This result further shows that RGL can generate a more fine-grained and accurate description.

To study how RGL generates a video description, we visualize the selection process during word generation (Figure 5). The histogram
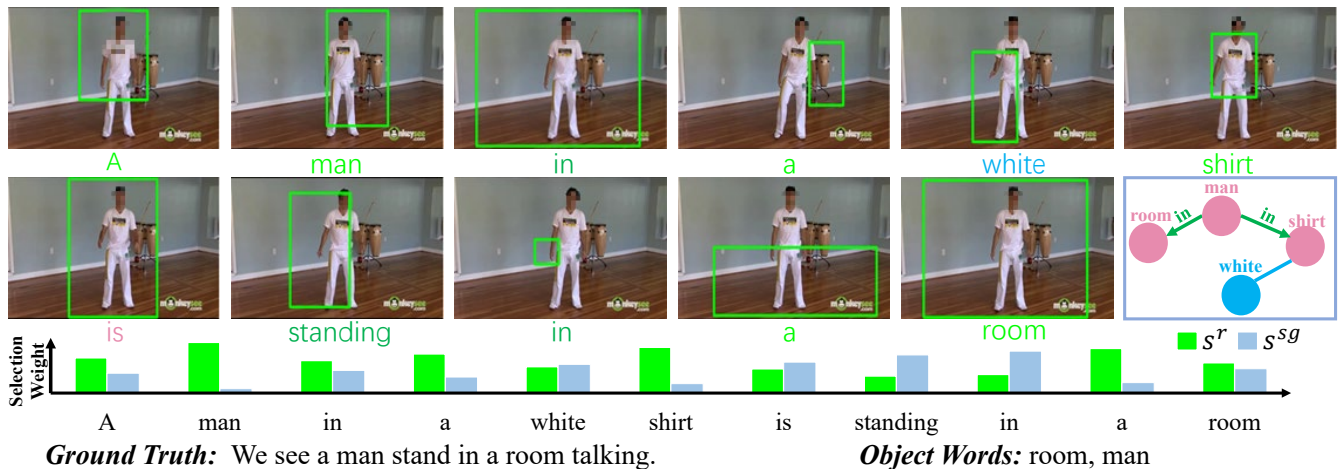
Figure 5: Visualization of words generation. Green texts represent the higher weight of grounding regions. Pink(objects), green(relationships) and blue(attributes) texts indicate that $\mathcal{SG}$ is more important. Green box corresponds to the region with the highest attention a^r. Histogram presents the selection weight of s^{sg} and s^r.
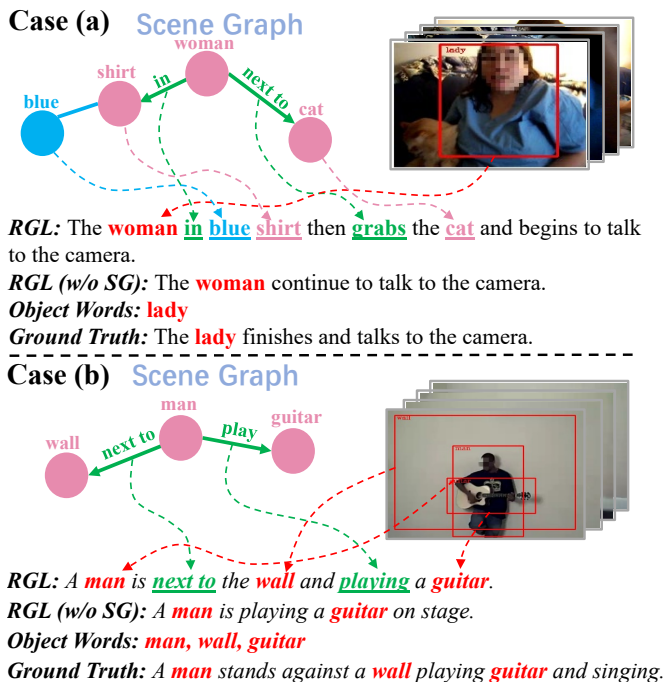

Figure 6: Qualitative examples from RGL and RGL (w/o SG). For each figure, the $\mathcal{SG}$ is pruned to avoid clutter. Three word colors correspond to objects, relationships and attributes in detected $\mathcal{SG}$.

presents that the RGL selects *how much* from $\tilde{r}$ (grounding regions) and $\tilde{u}$ ($\mathcal{SG}$) to generate each word. The words in green/(pink, green, or blue) color show the RGL select *what* from grounding regions and $\mathcal{SG}$. On the one hand, our model correctly attends to video

regions while generating the object words "*room*" and "*man*", the extra word "*shirt*" is also grounded correctly. On the other hand, "*in*" and "*white*", the words in "⟨*man-in-shirt*⟩ and "⟨*shirt-white*⟩", the $\mathcal{SG}$ is more important than grounding regions. Thus, the fine-grained phrase ¡°man in white shirt¡± is generated.

5 CONCLUSION

We propose a novel RGL framework to train a GVD generation model. Moreover, we develop a language-refined $\mathcal{SG}$-based method that contains additional visual concepts to describe a video in fine detail. A novel language model with a selection mechanism is also designed. This model can dynamically select the required information from the grounding regions and refined $\mathcal{SG}$ to generate descriptions in a reasonable manner. Our experimental results show the effectiveness of our method via multiple qualitative evaluations. We hope the RGL method can complement the existing literature on video description and benefit further studies on vision and language.

6 ACKNOWLEDGEMENTS

REFERENCES

[1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*. Springer, 382–398.

[2] Rianne van den Berg, Thomas N Kipf, and Max Welling. 2017. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263* (2017).

[3] Shizhe Chen, Jia Chen, Qin Jin, and Alexander Hauptmann. 2017. Video captioning with guidance of multimodal latent topics. In *Proceedings of the 25th ACM international conference on Multimedia*. ACM, 1838–1846.

[4] Tyler Derr, Yao Ma, and Jiliang Tang. 2018. Signed graph convolutional networks. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 929–934.

[5] Henghui Ding, Xudong Jiang, Ai Qun Liu, Nadia Magnenat Thalmann, and Gang Wang. 2019. Boundary-aware feature propagation for scene segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*. 6819–6829.

[6] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. 2019. Semantic correlation promoted shape-variant context for segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8885–8894.

[7] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847* (2016).

[8] Jiuxiang Gu, Shafiq Joty, Jianfei Cai, Handong Zhao, Xu Yang, and Gang Wang. 2019. Unpaired Image Captioning via Scene Graph Alignments. *arXiv preprint arXiv:1903.10658* (2019).

[9] Jiuxiang Gu, Handong Zhao, Zhe Lin, Sheng Li, Jianfei Cai, and Mingyang Ling. 2019. Scene graph generation with external knowledge and image reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1969–1978.

[10] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[11] Atsuhiro Kojima, Takeshi Tamura, and Kunio Fukunaga. 2002. Natural language description of human activities from video images based on concept hierarchy of actions. *International Journal of Computer Vision* 50, 2 (2002), 171–184.

[12] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision* 123, 1 (2017), 32–73.

[13] Juzheng Li, Hang Su, Jun Zhu, and Bo Zhang. 2018. Essay-Anchor Attentive Multi-Modal Bilinear Pooling for Textbook Question Answering. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 1–6.

[14] Lijun Li and Boqing Gong. 2019. End-to-end video captioning with multitask reinforcement learning. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 339–348.

[15] Wei Li, Dashan Guo, and Xiangzhong Fang. 2018. Multimodal architecture for video captioning with memory networks and an attention mechanism. *Pattern Recognition Letters* 105 (2018), 23–29.

[16] Jinlai Liu, Zehuan Yuan, and Changhu Wang. 2018. Towards good practices for multi-modal fusion in large-scale video classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 0–0.

[17] Diego Marcheggiani, Joost Bastings, and Ivan Titov. 2018. Exploiting semantics in neural machine translation with graph convolutional networks. *arXiv preprint arXiv:1804.08313* (2018).

[18] Diego Marcheggiani and Ivan Titov. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. *arXiv preprint arXiv:1703.04826* (2017).

[19] Jonghwan Mun, Linjie Yang, Zhou Ren, Ning Xu, and Bohyung Han. 2019. Streamlined Dense Video Captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6588–6597.

[20] Jae Sung Park, Marcus Rohrbach, Trevor Darrell, and Anna Rohrbach. 2019. Adversarial inference for multi-sentence video description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6598–6608.

[21] Ramakanth Pasunuru and Mohit Bansal. 2017. Reinforced video captioning with entailment rewards. *arXiv preprint arXiv:1708.02300* (2017).

[22] Wenjie Pei, Jiyuan Zhang, Xiangrong Wang, Lei Ke, Xiaoyong Shen, and Yu-Wing Tai. 2019. Memory-attended recurrent network for video captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8347–8356.

[23] Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object Hallucination in Image Captioning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 4035–4045.

[24] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. 2016. Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision*. Springer, 817–834.

[25] Sebastian Schuster, Ranjay Krishna, Angel Chang, Li Fei-Fei, and Christopher D Manning. 2015. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *Proceedings of the fourth workshop on vision and language*. 70–80.

[26] Jiaxin Shi, Hanwang Zhang, and Juanzi Li. 2019. Explainable and explicit visual reasoning over scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8376–8384.

[27] Jingkuan Song, Zhao Guo, Lianli Gao, Wu Liu, Dongxiang Zhang, and Heng Tao Shen. 2017. Hierarchical lstm with adjusted temporal attention for video captioning. *arXiv preprint arXiv:1706.01231* (2017).

[28] Damien Teney, Lingqiao Liu, and Anton van den Hengel. 2017. Graph-structured representations for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1–9.

[29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.

[30] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to sequence-video to text. In *Proceedings of the IEEE international conference on computer vision*. 4534–4542.

[31] Junbo Wang, Wei Wang, Yan Huang, Liang Wang, and Tieniu Tan. 2018. M3: Multimodal memory modelling for video captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7512–7520.

[32] Linlin Wang, Zhu Cao, Gerard De Melo, and Zhiyuan Liu. 2016. Relation classification via multi-level attention cnns. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1298–1307.

[33] Xin Wang, Wenhu Chen, Jiawei Wu, Yuan-Fang Wang, and William Yang Wang. 2018. Video captioning via hierarchical reinforcement learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4213–4222.

[34] Xiaolong Wang, Abhinav Shrivastava, and Abhinav Gupta. 2017. A-fast-rcnn: Hard positive generation via adversary for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2606–2615.

[35] Fanyi Xiao, Leonid Sigal, and Yong Jae Lee. 2017. Weakly-supervised visual grounding of phrases with linguistic structures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5945–5954.

[36] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. 2017. Scene graph generation by iterative message passing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5410–5419.

[37] Chenggang Yan, Yunbin Tu, Xingzheng Wang, Yongbing Zhang, Xinhong Hao, Yongdong Zhang, and Qionghai Dai. 2019. STAT: Spatial-temporal attention mechanism for video captioning. *IEEE transactions on multimedia* 22, 1 (2019), 229–241.

[38] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. 2018. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 670–685.

[39] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. 2019. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 10685–10694.

[40] Yang Yang, Jie Zhou, Jiangbo Ai, Yi Bin, Alan Hanjalic, Heng Tao Shen, and Yanli Ji. 2018. Video captioning by adversarial LSTM. *IEEE Transactions on Image Processing* 27, 11 (2018), 5600–5611.

[41] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Clinical text classification with rule-based features and knowledge-guided convolutional neural networks. *BMC medical informatics and decision making* 19, 3 (2019), 71.

[42] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7370–7377.

[43] Li Yao, Atousa Torabi, Kyunghyun Cho, Nicolas Ballas, Christopher Pal, Hugo Larochelle, and Aaron Courville. 2015. Describing videos by exploiting temporal structure. In *Proceedings of the IEEE international conference on computer vision*. 4507–4515.

[44] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 974–983.

[45] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. 2017. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In *Proceedings of the IEEE international conference on computer vision*. 1821–1830.

[46] Mihai Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. 2016. Spatio-temporal attention models for grounded video captioning. In *asian conference on computer vision*. Springer, 104–119.

[47] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. 2018. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5831–5840.

[48] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. 2017. Visual translation embedding network for visual relation detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5532–5540.

[49] Wenqiao Zhang, Haochen Shi, Siliang Tang, Jun Xiao, Qiang Yu, and Yueting Zhuang. 2021. Consensus graph representation learning for better grounded image captioning. In *Proc 35 AAAI Conf on Artificial Intelligence*.

[50] Wenqiao Zhang, Siliang Tang, Yanpeng Cao, Shiliang Pu, Fei Wu, and Yueting Zhuang. 2019. Frame augmented alternating attention network for video question answering. *IEEE Transactions on Multimedia* 22, 4 (2019), 1032–1041.

[51] Wenqiao Zhang, Siliang Tang, Yanpeng Cao, Jun Xiao, Shiliang Pu, Fei Wu, and Yueting Zhuang. 2020. Photo Stream Question Answer. In *Proceedings of the 28th ACM International Conference on Multimedia*. 3966–3975.

[52] Wenqiao Zhang, Siliang Tang, Jiajie Su, Jun Xiao, and Yueting Zhuang. 2021. Tell and guess: cooperative learning for natural image caption generation with hierarchical refined attention. *Multimedia Tools and Applications* 80, 11 (2021), 16267–16282.

[53] Luowei Zhou, Yannis Kalantidis, Xinlei Chen, Jason J Corso, and Marcus Rohrbach. 2019. Grounded video description. In *Proceedings of the IEEE Conference on*

[54] Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. Attention-based bidirectional long short-term memory networks for relation classification. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 2: Short papers)*. 207–212.

[55] Bohan Zhuang, Lingqiao Liu, Chunhua Shen, and Ian Reid. 2017. Towards context-aware interaction recognition for visual relationship detection. In *Proceedings of the IEEE International Conference on Computer Vision*. 589–598.

Computer Vision and Pattern Recognition. 6578–6587.