

Consensus Graph Representation Learning for Better Grounded Image Captioning

Wenqiao Zhang^{1*}, Haochen Shi^{1*}, Siliang Tang^{1†}, Jun Xiao¹, Qiang Yu², Yueting Zhuang¹

¹ Zhejiang University,

² Citycloud Technology

wenqiaozhang,hcshi,siliang.junx,yzhuang@zju.edu.cn, yq@citycloud.com.cn

Abstract

The contemporary visual captioning models frequently hallucinate objects that are not actually in a scene, due to the visual misclassification or over-reliance on priors that resulting in the semantic inconsistency between the visual information and the target lexical words. The most common way is to encourage the captioning model to dynamically link generated object words or phrases to appropriate regions of the image, i.e., the grounded image captioning (GIC). However, GIC utilizes an auxiliary task (grounding objects) that has not solved the key issue of object hallucination, i.e., the semantic inconsistency. In this paper, we take a novel perspective on the issue above: exploiting the semantic coherency between the visual and language modalities. Specifically, we propose the Consensus Rraph Representation Learning framework (CGRL) for GIC that incorporates a consensus representation into the grounded captioning pipeline. The consensus is learned by aligning the visual graph (*e.g.*, *scene graph*) to the language graph that consider both the nodes and edges in a graph. With the aligned consensus, the captioning model can capture both the correct linguistic characteristics and visual relevance, and then grounding appropriate image regions further. We validate the effectiveness of our model, with a significant decline in object hallucination (-9% CHAIR_i) on the Flickr30k Entities dataset. Besides, our CGRL also evaluated by several automatic metrics and human evaluation, the results indicate that the proposed approach can simultaneously improve the performance of image captioning (+2.9 Cider) and grounding (+2.3 F1_{LOC}).

Introduction

More recently, image captioning models have achieved impressive or even super-human performance on many benchmark datasets (He et al. 2019; Shuster et al. 2019; Deshpande et al. 2019). However, further quantitative analyses show that they are likely to generate hallucinated captions (Zhou et al. 2019; Ma et al. 2019), such as hallucinated objects words. Previous studies (Rohrbach et al. 2018) believed that this caption hallucination problem is caused by biased or inappropriate visual-textual correlations learned from the datasets, i.e., the semantic inconsistency between

the visual and language domain. Therefore, Grounded Image Captioning (GIC) is proposed to tackle this problem by introducing a new auxiliary task that requires the captioning model to ground object words back to the corresponding image regions during the caption generation. The auxiliary grounding task provides extra labels between visual and textual modals that can be used to remove biases and to reform the correct correlations between the two modalities.

However, GIC may not be the true savior of this hallucination problem. First, grounding object words is still far from solving the problem since the model can still hallucinate attributes of the objects and also relationships among the objects. Of course, we can introduce more grounding tasks to alleviate these new problems, but it comes at a tremendous cost and may induce new biases that are even harder to detect. Second, the correct correlations can hardly be fully constructed by grounding the annotation, since the image and the annotated caption are not always consistent (?). As is commonly known, such inconsistency happens frequently in real life, however, we humans have the reasoning ability to summarize or to infer the consensus knowledge between currently imperfect information and existing experience from the inconsistent environment. This ability enables us to perform better than machines in high-level reasoning and would be the most precious capacity for modern AI. Therefore, it is more important to enhance the reasoning capacities of models rather than just create more annotations.

Based on this insight, we propose a novel learning framework to mimic the human inference procedure, i.e., consensus graph representation learning framework (CGRL). The CGRL can leverage structural knowledge (*e.g.* *scene graph*, *SG*) from both vision and text, and further generate grounded captions based on consensus graph reasoning to alleviate the hallucination issue. As shown in Figure 1, the consensus representation is inferred by aligning visual graph to language graph. Utilizing such consensus, the model can capture the accurate and fine-grained information among objects to predict the non-visual function words reasonably (*e.g.*, *relationship verbs*: “ride”, “play”, *attribute adjective*: “red”, “striped”), while these words are inherently challenging to predict in GIC manner. Besides, the appropriate correlations between image and the annotated caption can be mined in the consensus to against the semantic inconsistency across the visual-language domain. Exactly as illus-

*These authors contributed equally to this work.

†Corresponding Authors

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1. An example of how consensus facilitates the grounded image captioning. Pink, green and blue represent the objects, relationships and attributes, respectively.

trated in Figure 1, although visual and language graphs are quite varied, the CGRL employs the aligned consensus that can maintain semantic consistency, and then generates the accurate description which captures both the correct linguistic characteristics and visual relevance. In addition, the object words “woman”, “shirt”, “laptop” and “table” are also grounded appropriately on the spatial regions.

Specifically, in our setting, the training pipeline of CGRL consists of three parts: 1) *Consensus Reasoning Module*, we first build the image scene graph SG^V and the language scene graph SG^L from the image and its ground truth (GT) at the training stage. Then we infer the consensus representation by aligning the SG^V to SG^L . This is a challenging task: as shown in Figure 1, the category and number of visual concepts in SG are varied in visual and language domain; besides, for a graph, the structural information (e.g., the diversity among objects, attributes and relationships) also needs to be aligned. To infer the consensus, we propose a Generative Adversarial Structure Network (GASN) that aligns the SG^V to SG^L . We first encode the SG at three unified levels (object, relationship, attribute) by GCN, then we simultaneously align the nodes and edges in the encoded graph by GASN to exploit the semantic consistency across domains. The representation of the aligned result can be regarded as the consensus for better GIC, 2) *Sentence Decoder*, we first exploit the latent spatial relations between image proposals and link them as the augmented region information for the sentence generator. Then the sentence decoder learns to decide the utilization of the augmented regions and consensus representation to describe an image more reasonably and accurately. 3) *Grounding Module*, we build a grounding and localizing mechanism. It not only encourages the model to dynamically ground the regions based on the current semantic context to predict words, but also localizes regions by the generated object words. Such a setting can boost the accuracy of the object word generation.

To summarize, the major contributions of our paper are as follows: 1) We present a novel perspective that alleviates the issue of object hallucination (CHAIR_i: -9%) by utilizing the consensus to maintain the semantic consistency across vision-language modality; 2) We propose a novel Consensus Graph Representation Learning framework (CGRL) that organically incorporates the consensus representation into the GIC pipeline for better grounded caption generation; 3) We

propose an adversarial graph alignment method to reason the consensus representation, which mainly addresses the problem of data representation and structural information for graph alignment; 4) We demonstrate the superiority of CGRL by automatic metrics and human judgments in boosting the captioning quality (Cider: +2.9) and grounding accuracy (F1_{LOC}: +2.3) over the state-of-the-art baselines.

Related Work

Image Captioning. Image captioning has been actively studied in recent vision and language research, the prevailing image captioning techniques often incorporate the encoder-decoder pipeline inspired by the first successful model (Vinyals et al. 2015). Benefiting from the rapid development of deep learning, the image captioning models have achieved striking advances by attention mechanism (Wang et al. 2019), reinforcement learning (Rennie et al. 2017; Zhang et al. 2017b) and generative adversarial networks (Chen et al. 2017; Xu et al. 2019). Although these methods have reached impressive performance on automatic metrics, they usually neglect how well the generated caption words are grounded in the image, making models less explainable and trustworthy.

Visual Grounding. Visual grounding models encourage captioning generator link phrases to specific spatial regions of the images, which present a potential way to improve the explainability of models. The most common way adopted by grounding models (Rohrbach et al. 2016; Xiao, Sigal, and Jae Lee 2017; Zhou et al. 2019) is to predict the next word through an attention mechanism, which is deployed over the NPs (Noun Phrases) with supervised bounding boxes as input. However, in GIC, visual grounding as an auxiliary task that has not truly solve the problem of semantic inconsistency can still hallucinate objects in the generated caption.

Scene Graph. Recently, SG construction has become popular research topics with significant advancements (Zellers et al. 2018; Yang et al. 2018; Gu et al. 2019b; Zhang et al. 2017a) based on the Visual Genome (Krishna et al. 2017) dataset. The SG contains the structured semantic information, it can represent scenes as directed graphs. Using this priors knowledge is natural to improve the performance of vision-language task, e.g., image captioning (Yang et al. 2019; Gu et al. 2019a), VQA (Shi, Zhang, and Li 2019; Teney, Liu, and van den Hengel 2017). However, directly fed the SG to the captioning model may lead to the non-correspondence problem between vision and language. Thus, how to infer the consensus knowledge from SG is the key to promote the vision-language field further.

Method

Task Description

Before presenting our method, we first introduce some basic notions and terminologies. Given an image \mathcal{I} , the goal of grounded image captioning (GIC) is to generate the natural language sentence \mathcal{S} , and localizing the object words with certain regions $\mathcal{R}^o = \{r_1^o, r_2^o, \dots, r_k^o\}$ in the corresponding image, where K is the number of grounding words appearing in sentence. Let \mathcal{M} be the GIC model with initialized

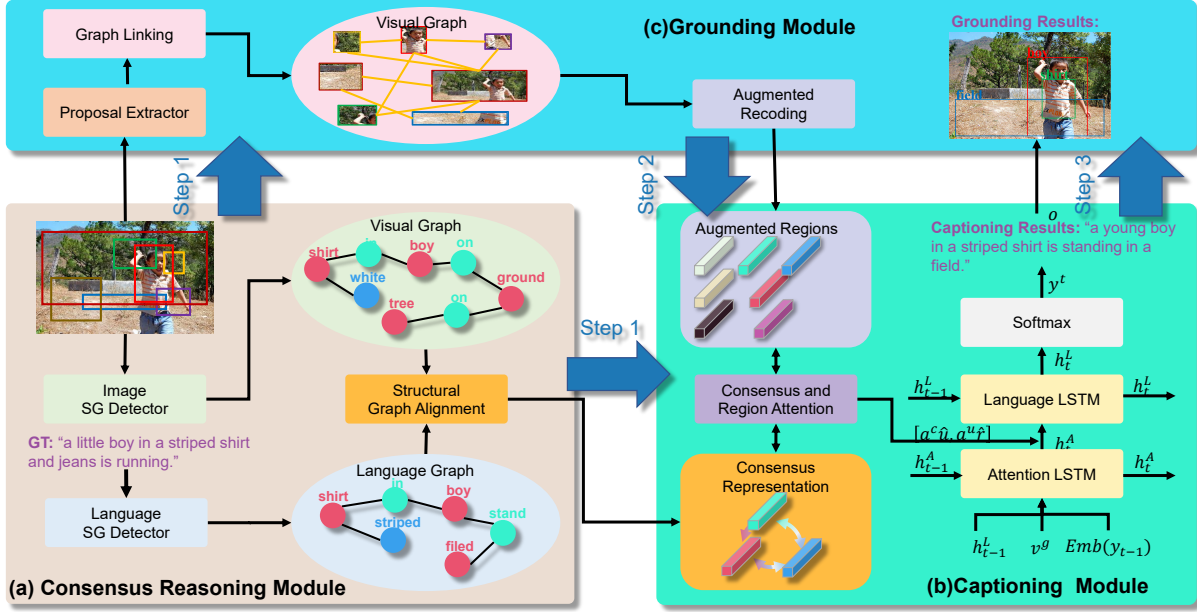


Figure 2. Overview of our method. It consists of three modules: (a) *Consensus Reasoning Module* first transforms an image and its GT into the \mathcal{SG} . Then we develop a Generative Adversarial Structure Network (GASN) that aligns the graphs from the visual domain to the language domain and learn the consensus representation. (Section 3.2); *Captioning Module* is encouraged to predict the textual word based on the consensus and augmented regions. (Section 3.3); *Grounding Module* compute the accuracy of localization of object words on the GT sentence, and the captioning model dynamically attends region proposals in the word prediction stage. (Section 3.4). The blue arrow represents the order of interaction between modules.

parameter Θ , trained to generate a sentence \hat{S} and localize the object regions $\hat{\mathcal{R}}^o$ for an image. *i.e.*, $(\hat{S}, \hat{\mathcal{R}}^o) = \mathcal{M}(\mathcal{I})$. We define the loss for a training pair as $\mathcal{L}((\hat{S}, \hat{\mathcal{R}}^o), \mathcal{M}(\mathcal{I}; \Theta))$. Fig.2 illustrates the overview of our Consensus Graph Representation Learning (CGRL) system, we illustrate our method for the GIC task as an example for clarity.

Consensus Reasoning Module

In our setting, we regard the pre-extracted \mathcal{SG} from an image and its corresponding GT sentence as the visual graph \mathcal{SG}^v and the language graph \mathcal{SG}^l , respectively. We aim to translate the \mathcal{SG}^v to \mathcal{SG}^l , the result of translation can be regarded as the consensus knowledge. Generally, the $\mathcal{SG} = (\mathcal{N}, \mathcal{E})$ defined in our task contains a set of nodes \mathcal{N} and edges \mathcal{E} . The node set \mathcal{N} contains three types of nodes: object node o , attribute node a , and relationship node r .

Scene Graph Encoder We first follow (Yang et al. 2019) and (Schuster et al. 2015) to extract the \mathcal{SG}^v and \mathcal{SG}^l , respectively. To represent the nodes of \mathcal{SG} at a unified level, we first denote the nodes $\{o, a, r\}$ by the label embeddings $\{\mathbf{e}_o, \mathbf{e}_a, \mathbf{e}_r\} \in \mathbb{R}^e$, corresponding respectively to objects, attributes and relationships both in \mathcal{SG}^v and \mathcal{SG}^l . Then we employ node encoder of the Graph Convolutional Network (GCN) (Marcheggiani and Titov 2017; Yang et al. 2019) to encode the three nodes in \mathcal{SG} at a unified representation $\mathcal{U} = \{\mathbf{u}_o, \mathbf{u}_a, \mathbf{u}_r\} \in \mathbb{R}^u$. Note that all the GCNs are defined with the same structure but independent parameters.

Objects Encoding. For an object o_i in \mathcal{SG} , it can act different roles (“*subject*” or “*object*”) due to different edge di-

rections, *i.e.*, two triples with different relationships, $\langle o_i - r_{i,j} - o_j \rangle$ and $\langle o_k - r_{k,i} - o_i \rangle$. Such associated objects by the cascaded encoding scheme can represent the global information of objects for an image or a sentence. Therefore we compute \mathbf{u}_o by explicit role modeling:

$$\mathbf{u}_o = \sum_{i=1}^{N_o} (a_i^o \cdot \frac{1}{N_{o_i}} (\sum_{\mathbf{e}_{o_i} \in sub} E_{subject}(\mathbf{e}_{o_i}, \mathbf{e}_{r_{i,j}}, \mathbf{e}_{o_j}) + \sum_{\mathbf{e}_{o_i} \in obj} E_{object}(\mathbf{e}_{o_k}, \mathbf{e}_{r_{k,i}}, \mathbf{e}_{o_i}))) \quad (1)$$

where $N_{o_i} = |sub(o_i)| + |obj(o_i)|$ is the total number of the relationship triplets for object o_i . N_o is the number of objects. a_i^o is the soft attention computed by all the objects. $E_{subject}(\cdot)$ and $E_{object}(\cdot)$ are the graph convolutional operation for objects as a “*subject*” or an “*object*”.

Attributes Encoding. Given o_i in \mathcal{SG} , it usually includes several attributes $\{a_1^i, \dots, a_{N_{a_i}}^i\}$, where N_{a_i} is the total number of its attributes. Therefore, \mathbf{u}_a can be computed as:

$$\mathbf{u}_a = \sum_{i=1}^{N_a} (a_i^a \cdot \frac{1}{N_{a_i}} \sum_{i \in N_{a_i}} E_{attribute}(\mathbf{e}_{o_i}, \mathbf{e}_{a_i})) \quad (2)$$

where $E_{attribute}(\cdot)$ is the graph convolutional operation for object o_i and its attributes. N_a is the number of attributes. a_i^a is the soft attention computed by all the attributes.

Relationships Encoding. Between two salient objects o_i and o_j , their relation is given by the triplet $\langle o_i - r_{i,j} - o_j \rangle$.

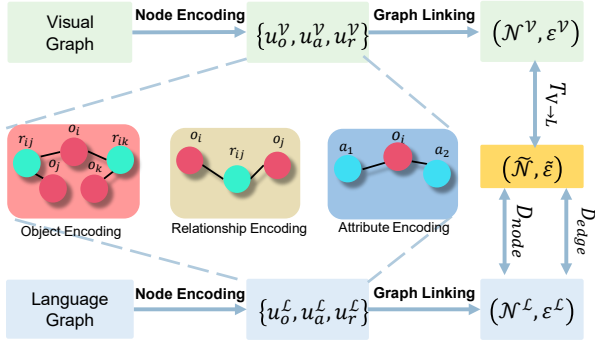


Figure 3. An overview of the consensus representation generation. The translator $T_{P \rightarrow K}(\cdot)$ aims to fool the node discriminator D_{node} and edge discriminator D_{edge} through adversarial training.

Similarly, the relationship encoding \mathbf{u}_r is produced by:

$$\mathbf{u}_r = \sum_{i=1}^{N_{r_i}} \sum_{j=1}^{N_{r_j}} (a_i^r \cdot E_{relationship}(\mathbf{e}_{o_i}, \mathbf{e}_{r_{i,j}}, \mathbf{e}_{o_j})) \quad (3)$$

where $E_{relationship}(\cdot)$ is the graph convolutional operation for relational object o_i and o_j . N_{r_i} and N_{r_j} is the number of relationships for object o_i and o_j . a_i^r is the soft attention computed by all the relationships.

After graph encoding, for each \mathcal{SG}^V and \mathcal{SG}^L , we have three node embeddings:

$$\mathcal{U}^M = \{\mathbf{u}_o^M, \mathbf{u}_a^M, \mathbf{u}_r^M\}, M \in \{V, L\} \quad (4)$$

Adversarial Consensus Representation Generation To summarize and infer the consensus representation to assist the model for better GIC, we need to translate the \mathcal{SG}^V to \mathcal{SG}^L . Instead of directly modeling the distribution alignment across domains, we take the discrepancy in the modality of \mathcal{SG} directly into account by aligning the data representation and structural information of \mathcal{SG}^V with the \mathcal{SG}^L .

In our setting, we link the \mathcal{U} to build a complete graph $\mathcal{G}_c = (\mathcal{N}_c, \mathcal{E}_c)$, where nodes are $\mathcal{N}_c = \{\mathbf{u}_o, \mathbf{u}_a, \mathbf{u}_r\}$ and edge is the distance between each pair of nodes. The edge values \mathcal{E}_c are calculated based on cosine similarity. For example, $\mathbf{e}_{c_{oa}} = \exp(-\frac{\mathbf{u}_o \cdot \mathbf{u}_a}{\|\mathbf{u}_o\| \|\mathbf{u}_a\|})$ is the value of edge between nodes \mathbf{u}_o and \mathbf{u}_a . To reason the consensus, we develop a Generative Adversarial Structure Network (GASN) to align the nodes and edges from \mathcal{SG}^V to \mathcal{SG}^L . As shown in Figure 3, the translator $T_{V \rightarrow L}(\cdot)$ is trained to generate the $(\tilde{\mathcal{N}}_c, \tilde{\mathcal{E}}_c)$, the overall loss \mathcal{L}_G of GASN consists of two parts:

$$\mathcal{L}_G = \lambda_N \cdot \mathcal{L}(N) + \lambda_E \cdot \mathcal{L}(E) \quad (5)$$

where $\mathcal{L}(N)$ and $\mathcal{L}(E)$ are the cross-entropy losses for the node alignment and edge alignment, respectively.

Node Alignment. For node representation alignment, we utilize a node discriminator D_N to discriminate which source the latent vector comes from, and force the representation of \mathcal{SG}^V obey the \mathcal{SG}^L distribution. Given a pair

$(\mathcal{N}_c^V, \mathcal{N}_c^L)$, the discriminator D_N is trained with the following objective:

$$\mathcal{L}(N) = \mathbb{E}_L[\log D_N(\mathcal{N}_c^L)] + \phi_N \cdot \mathbb{E}_L[\log D_N(\mathcal{N}_c^{\hat{L}})] + \psi_N \cdot \mathbb{E}_V[\log(1 - D_N(T_{V \rightarrow L}(\mathcal{N}_c^V))] \quad (6)$$

where $\mathcal{U}^{\hat{L}}$ is the other \mathcal{SG}^S representation from the textual corpus, ϕ_N and ψ_N are the hyper-parameters.

Edge Alignment. Similarly, the edge alignment is based on an edge discriminator D_E . As the structural information is recorded in the edges of structural graph \mathcal{G}_s , the alignment of structures can be modeled as the alignment between edges of \mathcal{E}_c^V and \mathcal{E}_c^L . Thus, a consistency loss is designed to align the edges.

$$\mathcal{L}(E) = \mathbb{E}_L[\log D_E(\mathcal{H}(\mathcal{E}_s^L, \mu^L))] + \phi_E \cdot \mathbb{E}_L[\log D_E(\mathcal{H}(\mathcal{E}_s^{\hat{L}}, \mu^{\hat{L}}))] + \psi_E \cdot \mathbb{E}_V[\log(1 - D_E(T_{V \rightarrow L}(\mathcal{H}(\mathcal{E}_s^V, \mu^V)))] \quad (7)$$

where $\mathcal{H}(\mathcal{E}, \mu) = 1/(1 + e^{-(\mathcal{E}-\mu)})$ is a logistic function with sigmoid's midpoint μ . μ is the mean of \mathcal{E}_s . Constrained by the above Equation, similar edges in \mathcal{E}_s^V remain closely \mathcal{E}_s^L located in the feature space. ϕ_E and ψ_E are the hyper-parameters.

On the basis of the aforementioned graph alignment, each node in the aligned graph $(\tilde{\mathcal{N}}_c, \tilde{\mathcal{E}}_c)$ is concatenated by itself and the product of edge and its adjacent nodes. Thus, the consensus representation $\tilde{\mathcal{U}} = \{\tilde{\mathbf{u}}_o, \tilde{\mathbf{u}}_a, \tilde{\mathbf{u}}_r\}$ is obtained, and then fed into the captioning module, providing a consensus for image captioning and object grounding.

Captioning Module

In this module, we exploit the latent information of spatial region proposals and aforementioned consensus representation for caption generation. a

Augmented Region Proposals. The region proposals $\mathcal{R} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N\}$ may contains multiple visual concepts, each $\mathbf{r}_i = [x, y, w, h]$ corresponds to a 4-dimensional spatial coordinate, where (x, y) denotes the coordinate of the top-left point of the bounding box, and (h, w) corresponds to the height, width of the box. To exploit their latent spatial relations, following (Yao et al. 2018), we link the proposals by their relative distance to a visual graph $\mathcal{G}_s = (\mathcal{N}_s, \mathcal{E}_s)$. Except for the two regions are too far away from each other, i.e., their Intersection over Union < 0.5 , the spatial relation between them is tend to be weak and no edge is established in this case. Finally, the proposals are re-encoded by a GCN, which allows information passing across all proposals.

$$\tilde{\mathbf{r}}_i = \sum_{j=1}^{N_{r_i}} E_{region}(\mathbf{r}_i, \mathbf{s}_{ij}, \mathbf{r}_j) \quad (8)$$

where N_{r_j} is the total number of the region proposals that associated with \mathbf{r}_i . $E_{region}(\cdot)$ is the graph convolutional operation for region proposals. Thus, the augmented region proposals are given by $\tilde{\mathcal{R}} = \{\tilde{\mathbf{r}}_1, \tilde{\mathbf{r}}_2, \dots, \tilde{\mathbf{r}}_N\}$.

Sentence Generator. We extend the state-of-the-art of Top-Down visual attention language model (?) as our captioning model. This model consists of two-layer LSTMs : Attention LSTM and Language LSTM (see Fig. 2b). The first

one for encoding the image features \mathbf{v}^g and embedding of the previously generated word \mathbf{y}_{t-1} into the hidden state \mathbf{h}_t^A . The second one for caption generation. We allow the Language LSTM to select dynamically how much from the augmented regions and consensus to generate words. Specifically, a soft attention is developed on $\hat{\mathcal{U}}$ and $\hat{\mathcal{R}}$ by \mathbf{h}_t^A from Attention LSTM, and then encoding the attended augmented regions $\hat{\mathcal{R}}$ and consensus $\hat{\mathcal{U}}$ into the hidden state \mathbf{h}_t^L . In this setting, the captioning model can select the relevant information it needs on the basis of the current semantic context to describe an image reasonably.

Using the notation $\mathbf{y}_{1:T}$ is refer to a sequence of words $\{\mathbf{y}_1, \dots, \mathbf{y}_T\}$. Each step t , the conditional distribution over possible words is generated by the Softmax, the cross-entropy loss $\mathcal{L}(S)$ for captioning can be computed by:

$$\mathcal{L}(S) = -\lambda_L \cdot \sum_{i=0}^T \log(p(\mathbf{y}_t | \mathbf{y}_{0:T-1})) \quad (9)$$

Grounding Module

In this section, we aim to evaluate how well the captioning model grounds the visual objects.

Region Grounding. To further assist the language model to attend to the correct regions, following (Zhou et al. 2019), we develop the region attention loss $\mathcal{L}(R)$: we denote the indicators of positive/negative regions as $\rho = \{\rho_1, \dots, \rho_N\}$ in each time step, where $\rho_i = 1$ if the region ρ_i has over 0.5 IoU with GT box and otherwise 0. Combining the treating attention a^r (in section of Captioning Module), the proposal attention loss function is defined as:

$$\mathcal{L}(R) = -\lambda_R \cdot \sum_1^N \gamma_i \log a_i^r \quad (10)$$

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^p as below:

$$p^s(\mathbf{r}, a^r) = \text{Softmax}(W_s^T \mathbf{r} + a^r) \quad (11)$$

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} , the grounding loss function $\mathcal{L}(G)$ for word \mathbf{w} is defined as follows:

$$\mathcal{L}(L) = -\lambda_L \cdot \sum_{i=1}^N \gamma_i \log p^s(\mathbf{r}_i) \quad (12)$$

Training Algorithm

Algorithm 1 (see in appendix) details the pseudocode of our CGRL algorithm for GIC. First, we pretrain the language GCNs by reconstructing the sentences from the latent vector that is encoded from the language graph SG^L with the GCN (sentence $\rightarrow SG^L \rightarrow$ latent vector \rightarrow sentence $\bar{S} \rightarrow$ object grounding \bar{R}^o). Then, we keep the language GCN fixed, and learn the visual GCN by aligning the visual graph SG^V to

the language graph SG^L in the latent space. Basically, the encoded latent vector from the knowledge SG^L is used as supervised signals to learn the visual GCN.

Experiments

Dataset and Setting

Dataset. We benchmark our approach for GIC on the Flickr30k Entities dataset and compare our CRGL method to the state-of-the-art models. Moreover, the Flickr30k Entities collected 31k images with 275k bounding boxes with associated with natural language phrases. Each image is annotated with 5 crowdsourced captions. There are 290k images for training, 1k images for validation, and another 1k images for testing¹.

Settings. Given an image as the input, the region encoding r is extracted from pre-trained Faster R-CNN (Ren et al. 2015). The global visual feature \bar{v}^g contains its all-region features which are given by ResNeXt-101 (Xie et al. 2017). For image captioning, 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 *Unknown token* to replace the words out of the vocabulary list. The vocabulary has 7,000 words, and each word is represented by a 512-dimensional vector, the RNN encoding size $m = 1024$.

Comparing Methods

We compare the proposed Consensus Graph Representation Learning (CGRL) algorithm with existing SoTA method NBK (Lu et al. 2018), Cyclical (Ma et al. 2019) and GVD (Zhou et al. 2019) on Flickr30k Entities dataset.

We also conduct an ablation study to investigate the contributions of individual components in CGRL. In our experiment, we train the following variants of CGRL: **CGRL (w/o CR)**, which is a general GIC model that only use image features to generate grounded image caption. **CGRL (w/o ARP)**, generating captions without augmented region proposals to study the importance of region supervision. **CGRL (w/o OG)**, generating captions sequentially without object word grounding, which is similar to standard image captioning algorithm. **CGRL (w/o NA)**, generating captions without node alignment in Generative Adversarial Structure Network. **CGRL (w/o EA)**, generating captions without edge alignment in Generative Adversarial Structure Network.

Metrics

Automatic Metrics. We use the captioning evaluation tool² provided by the 2018 ActivityNet Captions Challenge, which includes *BLEU@1*, *BLEU@4*, *METEOR*, *CIDEr*, and *SPICE* to evaluate the captioning results. Following³, we compute *F1_{ALL}*, *F1_{LOC}*, *GRD*, and *ATT* to evaluate the grounding results. For hallucination evaluation, we compute the *CHAIR_i* and *CHAIR_s* provide by (Rohrbach et al. 2018). Besides, we also record the recall of object words

¹<https://hockenmaier.cs.illinois.edu/DenotationGraph/>

²https://github.com/ranjaykrishna/densevid_eval

³<https://github.com/facebookresearch/ActivityNet-Entities>

Method	CR.	Captioning Evaluation					Grounding Evaluation			
		BLEU@1	BLEU@4	METEOR	CIDEr	SPICE	GRD.	ATT.	F1 _{ALL}	F1 _{LOC}
NBK *		69.0	27.1	21.7	57.5	15.6	-	-	-	-
Cyclical *		68.9	26.6	22.3	60.9	16.3	-	-	4.85	13.4
GVD [†]		69.9	27.3	22.5	62.3	16.5	41.4	50.9	7.55	22.2
CGRL (w/o CR) [‡]	✓	70.0	27.4	22.4	62.9	16.5	41.6	51.0	7.48	22.1
CGRL (w/o ARP) [‡]	✓	69.9	27.0	22.2	61.2	16.3	27.4	29.5	4.49	12.8
CGRL (w/o OG) [‡]	✓	72.9	28.3	22.4	65.4	16.8	45.5	55.9	-	-
CGRL (w/o NA) [‡]	✓	70.9	26.8	21.3	62.1	16.3	41.9	50.7	7.59	22.6
CGRL (w/o EA) [‡]	✓	72.7	27.4	22.1	64.1	16.8	43.9	54.1	8.00	23.4
CGRL [‡]	✓	72.5	27.8	22.4	65.2	16.8	44.3	54.2	8.01	23.7

Table 1. Captioning results and grounding results on Flickr30k Entities test set. CR. indicates the consensus representation knowledge is used or not for GVD generation. * indicates the results are obtained from the original papers. [‡] captions are implemented by author from the code of original paper. [†] sentences obtained directly from the author. Larger value indicates better performance. Top one score on each metric are in bold. Acronym notations of each method see in Sec 4.3.

Method	Hallucination Evaluation		
	CHAIR _i	CHAIR _s	RECALL _o
CGRL (w/o CR) [‡]	0.437	0.732	0.541
CGRL (w/o ARP) [‡]	0.385	0.718	0.573
CGRL (w/o NA) [‡]	0.392	0.718	0.567
CGRL (w/o EA) [‡]	0.361	0.712	0.598
CGRL (w/o OG) [‡]	0.379	0.712	0.629
CGRL [‡]	0.347	0.707	0.644

Table 2. Hallucination results on Flickr30k Entities test set. Lower value of CHAIR_i, CHAIR_s indicates better performance, while the large value of RECALL_o is better.

Metric	Human Evaluation		
	CGRL is Better	CGRL is Worse	Equal
R _{object}	0.26 (+5%)	0.21	0.53
R _{relationship}	0.35 (+12%)	0.23	0.42
R _{attribute}	0.17 (+3%)	0.14	0.69
DES	0.22 (+6%)	0.16	0.62

Table 3. Human evaluation of our model and its variant CGRL (w/o CR) on Flickr30k Entities test set. For each pairwise comparison are compared by 100 random images. Larger value indicates better performance.

RECALL_o that correctly predicted in each sentence. See the appendix for more details.

Human Evaluation. To better understand how satisfactory are the sentences generated from our CGRL, we also conducted a human study with 5 experienced workers to compare the descriptions generated by CGRL and CGRL (w/o CR), and asked them which one is more descriptive and captures more key visual concepts. Besides, we validate relevance of object/relationship/attribute words in generated captions by R_{object}, R_{relationship} and R_{attribute}. On the other hand, we also allow humans to evaluate which caption is more descriptive subjectively by DES. For each pairwise comparison, 100 images are randomly sampled from the Karpathy split for the workers to evaluate.

Experimental Results

Captioning Results. Table 1 shows the overall qualitative results of our model and SoTA on the test set of Flickr30k Entities dataset. In general, CGRL achieves the best performance on almost all the metrics to SoTA, METEOR also gets a comparable result with a small gap (within 0.1). Notice that we obtain 0.6 improvements on SPICE, since our method learned the consensus representation of SG, which gives a positional semantic prior to improve this score.

Grounding Results. Table 1 also presents our method effectively improves the accuracy of GRD., ATT., F1_{ALL} and F1_{LOC}. According to our observation, CGRL almost obtains the best performance for the majority of grounding metrics (except for ATT.). Since the captioning benefit from the consensus representation and grounding regions to generate correct words. Notice that CGRL (w/o OG) outperform other variants on attention correctness ATT, our intuition is that the CGRL (w/o OG) without object grounding may pay more attention to the word localization on GT sentence.

Hallucination Results. Table 2 presents object hallucination on the test set. We note an interesting phenomenon that the consensus based methods tend to perform better on the CHAIR_i, CHAIR_s and RECALL_o metrics than methods without consensus by a large margin. This proves the consensus contains key knowledge that can help the captioning model capture more correct objects in image. What’s more, based on this basis, the hallucination of objects is decreased, we believe that the consensus 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 images more accurately, i.e., some captioning models can describe an image in more details but with the lower captioning scores. Thus, we conducted a human evaluation. The results of the comparisons are shown in Table 3. It is seen from the table that CGRL outperforms CGRL (w/o CR) in terms of R_{object}, R_{relationship} and R_{attribute} by a large margin. The results indicate that our model with consensus can capture more key visual concepts in generated captions. This proves the effectiveness of CGRL for grounded image captioning.

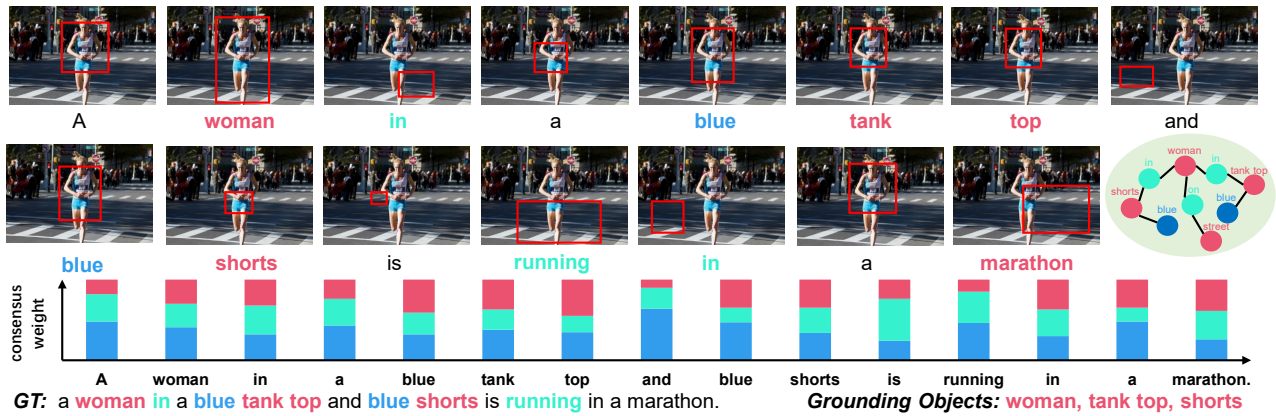


Figure 4. Visualization of image caption generation. Red box corresponds to the region with the highest attention a^r . Histogram presents the consensus weight for $\tilde{\mathbf{u}}_o$, $\tilde{\mathbf{u}}_a$, $\tilde{\mathbf{u}}_r$, respectively.

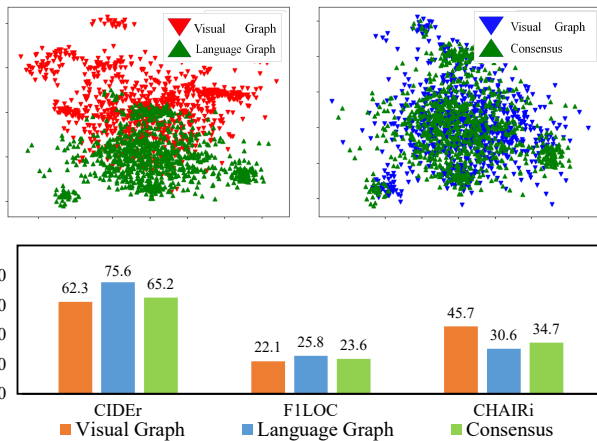


Figure 5. Visualization of features (1000 cases) in 2D space by t-SNE and the GIC performance of different knowledge. To facilitate the comparison, the value of CHAIR_i is multiplied by 100.

On another side, the sentence produced by CGRL achieves a higher score of the *DES* evaluation. We believe this result benefits from the consensus in CGRL.

Quantitative Analysis

We further validate the several key issues of the proposed method by answering three questions as follows.

Q1: How much improvement of semantic inconsistency has the consensus brought? First, the visualization of features alignment indicates that our GASN is able to capture the semantic consistency across vision-language domain in Figure 4. To justify the contribution of consensus, we investigated the importance of different knowledge by three variants of CGRL. In Figure 4, the best performance is to use the $\mathcal{SG}^{\mathcal{L}}$ directly. This is reasonable since the $\mathcal{SG}^{\mathcal{L}}$ contains the key concepts to reconstruct its description naturally. In addition, consensus-based outperforms $\mathcal{SG}^{\mathcal{V}}$ by a large margin on all the important metrics, especially the CHAIRi (-11.0). The experimental results indicate that the consensus can improve the quality of GIC greatly.

Q2: How does the CGRL generate the captions based on consensus and region attention?

We visualize the process of caption generation in Figure 8(a)(In appendix). On the one hand, from the the histogram, it clearly demonstrates the CGRL select *how much* from the consensus $\tilde{\mathbf{u}}_o$, $\tilde{\mathbf{u}}_a$, $\tilde{\mathbf{u}}_r$, respectively. For example, consensus knowledge is more important while generating the relational words “run”. On the other hand, CGRL attends to the appropriate proposal to predict a word, such the word “woman”. In summary, on the basis of the consensus attention and region attention, our model can select the concerned information to generate the correct word.

Conclusion

In this paper, we propose a novel Consensus Graph Representation Learning (CGRL) framework to train a GIC model. We design a consensus based method that aligns the visual graph to language graph through a novel generative adversarial structure network, which aims to maintain the semantic consistency between multi-modals. Based on the consensus, our model not only can generate more accurate caption, but also ground appropriate regions and greatly alleviate the hallucination problem. We hope our CGRL can provide a complement for existing literature of visual captioning and benefit further study of vision and language.

Acknowledgments

This work has been supported in Apart by National Key Research and Development Program of China (2018AAA010010), NSFC (U19B2043, 61976185), Zhejiang NSF (LR21F020004, LR19F020002), Key Research & Development Project of Zhejiang Province (2018C03055), Funds from City Cloud Technology (China) Co. Ltd., Zhejiang University iFLYTEK Joint Research Center, Zhejiang University-Tongdun Technology Joint Laboratory of Artificial Intelligence, Chinese Knowledge Center of Engineering Science and Technology (CKCEST), Hikvision-Zhejiang University Joint Research Center, the Fundamental Research Funds for the Central Universities, and Engineering Research Center of Digital Library, Ministry of Education.

References

- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, 382–398. Springer.
- Berg, R. v. d.; Kipf, T. N.; and Welling, M. 2017. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*.
- Chen, H.; Zhang, H.; Chen, P.-Y.; Yi, J.; and Hsieh, C.-J. 2017. Show-and-fool: Crafting adversarial examples for neural image captioning. *arXiv preprint arXiv:1712.02051*.
- Chunseong Park, C.; Kim, B.; and Kim, G. 2017. Attend to you: Personalized image captioning with context sequence memory networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 895–903.
- Deshpande, A.; Aneja, J.; Wang, L.; Schwing, A. G.; and Forsyth, D. 2019. Fast, diverse and accurate image captioning guided by part-of-speech. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10695–10704.
- Fukui, A.; Park, D. H.; Yang, D.; Rohrbach, A.; Darrell, T.; and Rohrbach, M. 2016. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*.
- Gu, J.; Joty, S.; Cai, J.; Zhao, H.; Yang, X.; and Wang, G. 2019a. Unpaired Image Captioning via Scene Graph Alignments. *arXiv preprint arXiv:1903.10658*.
- Gu, J.; Zhao, H.; Lin, Z.; Li, S.; Cai, J.; and Ling, M. 2019b. Scene graph generation with external knowledge and image reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1969–1978.
- He, X.; Yang, Y.; Shi, B.; and Bai, X. 2019. VD-SAN: Visual-densely semantic attention network for image caption generation. *Neurocomputing* 328: 48–55.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision* 123(1): 32–73.
- Li, L.; Tang, S.; Deng, L.; Zhang, Y.; and Tian, Q. 2017. Image caption with global-local attention. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Liang, X.; Shen, X.; Xiang, D.; Feng, J.; Lin, L.; and Yan, S. 2016. Semantic object parsing with local-global long short-term memory. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3185–3193.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2018. Neural baby talk. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7219–7228.
- Ma, C.-Y.; Kalantidis, Y.; AlRegib, G.; Vajda, P.; Rohrbach, M.; and Kira, Z. 2019. Learning to Generate Grounded Image Captions without Localization Supervision. *arXiv preprint arXiv:1906.00283*.
- Marcheggiani, D.; Bastings, J.; and Titov, I. 2018. Exploiting semantics in neural machine translation with graph convolutional networks. *arXiv preprint arXiv:1804.08313*.
- Marcheggiani, D.; and Titov, I. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. *arXiv preprint arXiv:1703.04826*.
- Pei, W.; Zhang, J.; Wang, X.; Ke, L.; Shen, X.; and Tai, Y.-W. 2019. Memory-attended recurrent network for video captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8347–8356.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.
- Ren, Z.; Wang, X.; Zhang, N.; Lv, X.; and Li, L.-J. 2017. Deep reinforcement learning-based image captioning with embedding reward. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 290–298.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7008–7024.
- Rohrbach, A.; Hendricks, L. A.; Burns, K.; Darrell, T.; and Saenko, K. 2018. Object Hallucination in Image Captioning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4035–4045.
- Rohrbach, A.; Rohrbach, M.; Hu, R.; Darrell, T.; and Schiele, B. 2016. Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision*, 817–834. Springer.
- Schuster, S.; Krishna, R.; Chang, A.; Fei-Fei, L.; and Manning, C. D. 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.
- Shi, J.; Zhang, H.; and Li, J. 2019. Explainable and explicit visual reasoning over scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8376–8384.
- Shuster, K.; Humeau, S.; Hu, H.; Bordes, A.; and Weston, J. 2019. Engaging image captioning via personality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 12516–12526.
- Teney, D.; Liu, L.; and van den Hengel, A. 2017. Graph-structured representations for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1–9.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3156–3164.
- Wang, E. K.; Zhang, X.; Wang, F.; Wu, T.-Y.; and Chen, C.-M. 2019. Multilayer dense attention model for image caption. *IEEE Access* 7: 66358–66368.

- Wang, L.; Cao, Z.; De Melo, G.; and Liu, Z. 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.
- Wang, X.; Shrivastava, A.; and Gupta, A. 2017. A-fast-rnn: Hard positive generation via adversary for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2606–2615.
- Xiao, F.; Sigal, L.; and Jae Lee, Y. 2017. Weakly-supervised visual grounding of phrases with linguistic structures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5945–5954.
- Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1492–1500.
- Xu, D.; Zhu, Y.; Choy, C. B.; and Fei-Fei, L. 2017. Scene graph generation by iterative message passing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5410–5419.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, 2048–2057.
- Xu, Y.; Wu, B.; Shen, F.; Fan, Y.; Zhang, Y.; Shen, H. T.; and Liu, W. 2019. Exact Adversarial Attack to Image Captioning via Structured Output Learning with Latent Variables. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4135–4144.
- Yang, J.; Lu, J.; Lee, S.; Batra, D.; and Parikh, D. 2018. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 670–685.
- Yang, X.; Tang, K.; Zhang, H.; and Cai, J. 2019. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10685–10694.
- Yao, L.; Mao, C.; and Luo, Y. 2019a. Clinical text classification with rule-based features and knowledge-guided convolutional neural networks. *BMC medical informatics and decision making* 19(3): 71.
- Yao, L.; Mao, C.; and Luo, Y. 2019b. Graph convolutional networks for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7370–7377.
- Yao, T.; Pan, Y.; Li, Y.; and Mei, T. 2018. Exploring visual relationship for image captioning. In *Proceedings of the European conference on computer vision (ECCV)*, 684–699.
- Zanfir, M.; Marinoiu, E.; and Sminchisescu, C. 2016. Spatio-temporal attention models for grounded video captioning. In *asian conference on computer vision*, 104–119. Springer.
- Zellers, R.; Yatskar, M.; Thomson, S.; and Choi, Y. 2018. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5831–5840.
- Zhang, H.; Kyaw, Z.; Chang, S.-F.; and Chua, T.-S. 2017a. Visual translation embedding network for visual relation detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5532–5540.
- Zhang, L.; Sung, F.; Liu, F.; Xiang, T.; Gong, S.; Yang, Y.; and Hospedales, T. M. 2017b. Actor-critic sequence training for image captioning. *arXiv preprint arXiv:1706.09601*.
- Zhou, L.; Kalantidis, Y.; Chen, X.; Corso, J. J.; and Rohrbach, M. 2019. Grounded video description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6578–6587.
- Zhou, P.; Shi, W.; Tian, J.; Qi, Z.; Li, B.; Hao, H.; and Xu, B. 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.
- Zhuang, B.; Liu, L.; Shen, C.; and Reid, I. 2017. Towards context-aware interaction recognition for visual relationship detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 589–598.