

Relation-Aware Graph Attention Network for Visual Question Answering

Linjie Li, Zhe Gan, Yu Cheng, Jingjing Liu
 Microsoft Dynamics 365 AI Research

{lindsey.li, zhe.gan, yu.cheng, jingj1}@microsoft.com

Abstract

In order to answer semantically-complicated questions about an image, a Visual Question Answering (VQA) model needs to fully understand the visual scene in the image, especially the interactive dynamics between different objects. We propose a Relation-aware Graph Attention Network (ReGAT), which encodes each image into a graph and models multi-type inter-object relations via a graph attention mechanism, to learn question-adaptive relation representations. Two types of visual object relations are explored: (i) *Explicit Relations* that represent geometric positions and semantic interactions between objects; and (ii) *Implicit Relations* that capture the hidden dynamics between image regions. Experiments demonstrate that ReGAT outperforms prior state-of-the-art approaches on both VQA 2.0 and VQA-CP v2 datasets. We further show that ReGAT is compatible to existing VQA architectures, and can be used as a generic relation encoder to boost the model performance for VQA.

1. Introduction

Recent advances in deep learning have driven tremendous progress in both Computer Vision and Natural Language Processing (NLP). Interdisciplinary area between language and vision, such as image captioning, text-to-image synthesis and visual question answering (VQA), has attracted rapidly growing attention from both vision and NLP communities. Take VQA as an example - the goal (and the main challenge) is to train a model that can achieve comprehensive and semantically-aligned understanding of multimodal input. Specifically, given an image and a natural language question grounded on the image, the task is to associate visual features in the image with the semantic meaning in the question, in order to correctly answer the question.

Most state-of-the-art approaches to VQA [56, 11, 38, 33, 49] focus on learning a multimodal joint representation of images and questions. Specifically, a Convolutional Neural Network (CNN) or Region-based CNN (R-CNN) is com-

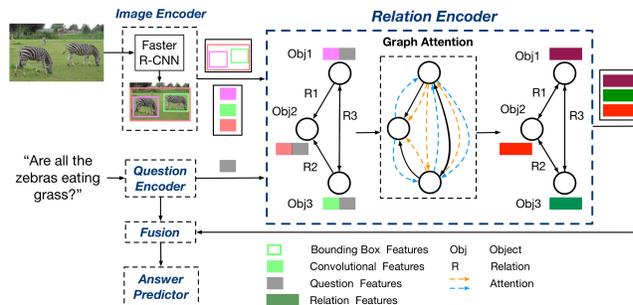


Figure 1. An overview of the ReGAT model. Both explicit relations (semantic and spatial) and implicit relations are considered. The proposed relation encoder captures question-adaptive object interactions via Graph Attention.

monly used as a visual feature extractor for image encoding, and a Recurrent Neural Network (RNN) is used for question encoding. After obtaining a sparse set of image regions from the visual feature extractor, multimodal fusion is applied to learn a joint representation that represents the alignment between each individual region and the question. This joint representation is then fed into an answer predictor to produce an answer.

This framework has proven to be useful for the VQA task, but there still exists a significant semantic gap between image and natural language. For example, given an image of a group of zebras (see Figure 1), the model may recognize the black and white pixels, but not which white and black pixels are from which zebra. Thus, it is difficult to answer questions such as “*Is the zebra at the far right a baby zebra?*” or “*Are all the zebras eating grass?*”. A VQA system needs to recognize not only the objects (“*zebras*”) and the surrounding environment (“*grass*”), but also the semantics about actions (“*eating*”) and locations (“*at the far right*”) in both images and questions.

In order to capture this type of action and location information, we need to go beyond mere object detection in image understanding, and learn a more holistic view of the visual scene in the image, by interpreting the dynamics and interactions between different objects in an image. One possible solution is to detect the relative geometrical positions of neighboring objects (e.g., <motorcycle-next

to-car>), to align with spacial descriptions in the question. Another direction is to learn semantic dependencies between objects (e.g., <girl-eating-cake>) to capture the interactive dynamics in the visual scene.

Motivated by this, we propose a Relation-aware Graph Attention Network (ReGAT) for VQA, introducing a novel relation encoder that captures these inter-object relations beyond static object/region detection. These visual relation features can reveal more fine-grained visual concepts in the image, which in turn provides a holistic scene interpretation that can be used for answering semantically-complicated questions. In order to cover the high variance in image scenes and question types, both explicit (e.g., spatial/positional, semantic/actionable) relations and implicit relations are learned by the relation encoder, where images are represented as graphs and interactions between objects are captured via a graph attention mechanism.

Furthermore, the graph attention is learned based on the context of the question, permitting the injection of semantic information from the question into the relation encoding stage. In this way, the features learned by the relation encoder not only capture object-interactive visual contents in the image, but also absorb the semantic clues in the question, to dynamically focus on particular relation types and instances for each question on the fly.

Figure 1 shows an overview of the proposed model. First, a Faster R-CNN is used to generate a set of object region proposals, and a question encoder is used for question embedding. The convolutional and bounding-box features of each region are then injected into the relation encoder to learn the relation-aware, question-adaptive, region-level representations from the image. These relation-aware visual features and the question embeddings are then fed into a multimodal fusion module to produce a joint representation, which is used in the answer prediction module to generate an answer.

In principle, our work is different from (and compatible to) existing VQA systems. It is pivoted on a new dimension: using question-adaptive inter-object relations to enrich image representations in order to enhance VQA performance. The contributions of our work are three-fold:

- We propose a novel graph-based relation encoder to learn both explicit and implicit relations between visual objects via graph attention networks.
- The learned relations are question-adaptive, meaning that they can dynamically capture visual object relations that are most relevant to each question.
- We show that our ReGAT model is a generic approach that can be used to improve state-of-the-art VQA models on the VQA 2.0 dataset. Our model also achieved state-of-the-art performance on the more challenging VQA-CP v2 dataset.

2. Related Work

2.1. Visual Question Answering

The current dominant framework for VQA systems consists of an image encoder, a question encoder, multimodal fusion, and an answer predictor. In lieu of directly using visual features from CNN-based feature extractors, [56, 11, 41, 33, 49, 38, 63, 36] explored various image attention mechanisms to locate regions that are relevant to the question. To learn a better representation of the question, [33, 38, 11] proposed to perform question-guided image attention and image-guided question attention collaboratively, to merge knowledge from both visual and textual modalities in the encoding stage. [15, 25, 60, 4, 24] explored higher-order fusion methods to better combine textual information with visual information (e.g., using bilinear pooling instead of simpler first-order methods such as summation, concatenation and multiplication).

To make the model more interpretable, some literature [30, 59, 29, 54, 55, 53] also exploited high-level semantic information in the image, such as attributes, captions and visual relation facts. Most of these methods applied VQA-independent models to extract semantic knowledge from the image, while [34] built a Relation-VQA dataset and directly mined VQA-specific relation facts to feed additional semantic information to the model. A few recent studies [48, 35, 29] investigated how to incorporate memory to aid the reasoning step, especially for difficult questions.

However, the semantic knowledge brought in by either memory or high-level semantic information is usually converted into textual representation, instead of directly used as visual representation, which contains richer and more indicative information about the image. Our work is complementary to these prior studies in that we encode object relations directly into image representation, and the relation encoding step is generic and can be naturally fit into any state-of-the-art VQA model.

2.2. Visual Relationship

Visual relationship has been explored before deep learning became popular. Early work [10, 14, 7, 37] presented methods to re-score the detected objects by considering object relations (e.g., co-occurrence [10], position and size [5]) as post-processing steps for object detection. Some previous work [16, 17] also probed the idea that spatial relationships (e.g., “above”, “around”, “below” and “inside”) between objects can help improve image segmentation.

Visual relationship has proven to be crucial to many computer vision tasks. For example, it aided the cognitive task of mapping images to captions [13, 12, 58] and improved image search [47, 23] and object localization [45, 21]. Recent work on visual relationship [45, 43, 9] focused more on non-spatial relation, or known as “semantic

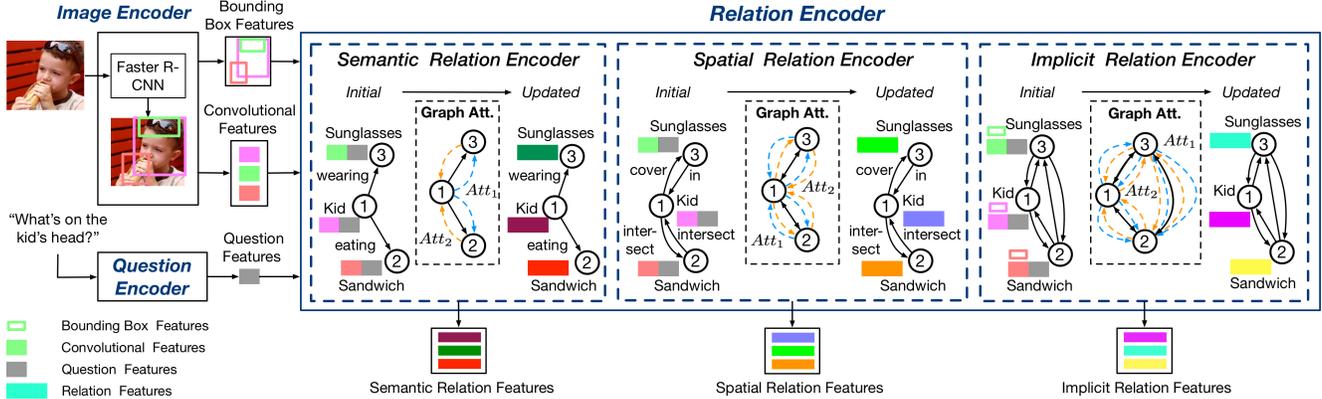


Figure 2. Model architecture of the proposed ReGAT for visual question answering. Faster R-CNN is employed to detect a set of object regions. These region-level features are then fed into different relation encoders to learn relation-aware question-adaptive visual features, which will be fused with question representation to predict an answer. Multimodal fusion and answer predictor are omitted for simplicity.

relation” (i.e., actions of, or interactions between objects). A few neural network architectures have been designed for the visual relationship prediction task [32, 8, 61].

2.3. Relational Reasoning

We name the visual relationship aforementioned as *explicit* relation, which has been shown to be effective for image captioning [58]. Specifically, [58] exploited pre-defined semantic relations learned from the Visual Genome dataset [28] and spatial relations between objects. A graph was then constructed based on these relations, and a Graph Convolutional Network (GCN) [26] was used to learn representations for each object.

Another line of research focuses on *implicit* relations, where no explicit semantic or spatial relations are used to construct the graph. Instead, all the relations are *implicitly* captured by an attention module or via higher-order methods over the fully-connected graph of an input image [46, 21, 6, 57], to model the interactions between detected objects. For example, [46] reasons over all the possible pairs of objects in an image via the use of simple MLPs. In [6], a bilinear fusion method, called MuRel cell, was introduced to perform pairwise relationship modeling.

Some other work [50, 39, 52] have been proposed for learning question-conditioned graph representations for images. Specifically, [39] introduced a graph learner module that is conditioned on question representations to compute the image representations using pairwise attention and spatial graph convolutions. [50] exploited structured question representations such as parse trees, and used GRU to model contextualized interactions between both objects and words. A more recent work [52] introduced a sparser graph defined by inter/intra-class edges, in which relationships are implicitly learned via a language-guided graph attention mechanism. However, all these work still focused on *implicit* relations, which are less interpretable than *explicit* relations.

Our contributions Our work is inspired by [21, 58]. However, different from them, ReGAT considers both explicit and implicit relations to enrich image representations. For explicit relations, our model uses Graph Attention Network (GAT) rather than a simple GCN as used in [58]. As opposed to GCNs, the use of GAT allows for assigning different importances to nodes of the same neighborhood. For implicit relations, our model learns a graph that is adaptive to each question by filtering out question-irrelevant relations, instead of treating all the relations equally as in [21]. In experiments, we conduct detailed ablation studies to demonstrate the effectiveness of each individual design.

3. Relation-aware Graph Attention Network

Here is the problem definition of the VQA task: given a question q grounded in an image I , the goal is to predict an answer $\hat{a} \in \mathcal{A}$ that best matches the ground-truth answer a^* . As common practice in the VQA literature, this can be defined as a classification problem:

$$\hat{a} = \arg \max_{a \in \mathcal{A}} p_{\theta}(a|I, q), \quad (1)$$

where p_{θ} is the trained model.

Figure 2 gives a detailed illustration of our proposed model, consisting of an Image Encoder, a Question Encoder, and a Relation Encoder. For the Image Encoder, Faster R-CNN [2] is used to identify a set of objects $\mathcal{V} = \{v_i\}_{i=1}^K$, where each object v_i is associated with a visual feature vector $v_i \in \mathbb{R}^{d_v}$ and a bounding-box feature vector $b_i \in \mathbb{R}^{d_b}$ ($K = 36$, $d_v = 2048$, and $d_b = 4$ in our experiments). Each $b_i = [x, y, w, h]$ 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. For the Question Encoder, we use a bidirectional RNN with Gated Recurrent Unit (GRU) and perform self attention on the sequence of RNN hidden states to generate question embed-

ding $q \in \mathbb{R}^{d_q}$ ($d_q = 1024$ in our experiments). The following sub-sections will explain the details of the Relation Encoder.

3.1. Graph Construction

Fully-connected Relation Graph By treating each object v_i in the image as one vertex, we can construct a fully-connected undirected graph $\mathcal{G}_{imp} = (\mathcal{V}, \mathcal{E})$, where \mathcal{E} is the set of $K \times (K - 1)$ edges. Each edge represents an implicit relation between two objects, which can be reflected by the learned weight assigned to each edge through graph attention. All the weights are learned implicitly without any prior knowledge. We name the relation encoder built on this graph \mathcal{G}_{imp} the *implicit* relation encoder.

Pruned Graph with Prior Knowledge On the other hand, if explicit relations between vertices are available, one can readily transform the fully-connected graph \mathcal{G}_{imp} into an explicit relation graph, by pruning the edges where the corresponding explicit relation does not exist. For each pair of objects i, j , if $\langle i-p-j \rangle$ is a valid relation, an edge is created from i to j , with an edge label p . In addition, we assign each object node i with a self-loop edge and label this edge as *identical*. In this way, the graph becomes sparse, and each edge encodes prior knowledge about one inter-object relation in the image. We name the relation encoder built upon this graph the *explicit* relation encoder.

The explicit nature of these features requires pre-trained classifiers to extract the relations in the form of discrete class labels, which represent the dynamics and interactions between objects explicit to the human eye. Different types of explicit relations can be learned based on this pruned graph. In this paper, we explore two instances: spatial and semantic graphs, to capture positional and actionable relations between objects, which is imperative for visual question answering.

Spatial Graph Let $spa_{i,j} = \langle object_i - predicate - object_j \rangle$ denote the spatial relation that represents the relative geometric position of $object_i$ against $object_j$. In order to construct a spatial graph \mathcal{G}_{spa} , given two object region proposals $object_i$ and $object_j$, we classify $spa_{i,j}$ into 11 different categories [58] (e.g., $object_i$ is inside $object_j$ (class 1), $object_j$ is inside $object_i$ (class 2), as illustrated in Figure 3(a)), including a *no-relation* class retained for objects that are too far away from each other. Note that edges formed by spatial relations are symmetrical: if $\langle object_i - p_{i,j} - object_j \rangle$ is a valid spatial relation, there must be a valid spatial relation $spa_{j,i} = \langle object_j - p_{j,i} - object_i \rangle$. However, the two predicates $p_{i,j}$ and $p_{j,i}$ are different.

Semantic Graph In order to construct semantic graph \mathcal{G}_{sem} , semantic relations between objects need to be extracted (e.g., $\langle subject - predicate - object \rangle$).

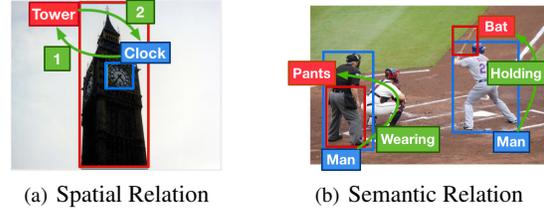


Figure 3. Illustration of spatial and semantic relations. The green arrows denote the direction of relations (subject \rightarrow object). Labels in green boxes are class labels of relations. Red and Blue boxes contain class labels of objects.

This can be formulated as a classification task [58] by training a semantic relation classifier on a visual relationship dataset (e.g., Visual Genome [27]). Given two object regions i and j , the goal is to determine which predicate p represents a semantic relation $\langle i-p-j \rangle$ between these two regions. Here, the relations between the subject j and the object i are not interchangeable, meaning that the edges formed by semantic relations are not symmetric. For a valid $\langle i-p_{i,j}-j \rangle$, there may not exist a relation $\langle j-p_{j,i}-i \rangle$ within our definition. For example, $\langle man - holding - bat \rangle$ is a valid relation, while there is no semantic relation from *bat* to *man*.

The classification model takes in three inputs: feature vector of the subject region v_i , feature vector of the object region v_j , and region-level feature vector $v_{i,j}$ of the union bounding box containing both i and j . These three types of feature are obtained from pre-trained object detection model, and then transformed via an embedding layer. The embedded features are then concatenated and fed into a classification layer to produce softmax probability over 14 semantic relations, with an additional *no-relation* class. The trained classifier is then used to predict relations between any pair of object regions in a given image. Examples of semantic relations are shown in Figure 3(b).

3.2. Relation Encoder

Question-adaptive Graph Attention The proposed relation encoder is designed to encode relational dynamics between objects in an image. For the VQA task, there might be different types of relations that are useful for different question types. Thus, in designing the relation encoder, we use a question-adaptive attention mechanism to inject semantic information from the question into relation graphs, to dynamically assign higher weights to those relations that are mostly relevant to each question. This is achieved by first concatenating the question embedding q with each of the K visual features v_i , denoted as

$$v'_i = [v_i || q] \quad \text{for } i = 1, \dots, K. \quad (2)$$

Self-attention is then performed on the vertices, which generates hidden relation features $\{v_i^*\}_{i=1}^K$ that characterize the

relations between a target object and its neighboring objects. Based on this, each relation graph goes through the following attention mechanism:

$$\mathbf{v}_i^* = \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \cdot \mathbf{W} \mathbf{v}'_j \right). \quad (3)$$

For different types of relation graph, the definition of the attention coefficients α_{ij} varies, so does the projection matrix $\mathbf{W} \in \mathbb{R}^{d_h \times (d_q + d_v)}$ and the neighborhood \mathcal{N}_i of object i . $\sigma(\cdot)$ is a nonlinear function such as ReLU. To stabilize the learning process of self-attention, we also extend the above graph attention mechanism by employing multi-head attention, where M independent attention mechanisms are executed, and their output features are concatenated, resulting in the following output feature representation:

$$\mathbf{v}_i^* = \parallel_{m=1}^M \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^m \cdot \mathbf{W}^m \mathbf{v}'_j \right). \quad (4)$$

In the end, \mathbf{v}_i^* is added to the original visual feature \mathbf{v}_i to serve as the final relation-aware feature.

Implicit Relation Since the graph for learning implicit relation is fully-connected, \mathcal{N}_i contains all the objects in the image, including object i itself. Inspired by [21], we design the attention weight α_{ij} to not only depend on visual-feature weight α_{ij}^v , but also bounding-box weight α_{ij}^b . Specifically,

$$\alpha_{ij} = \frac{\alpha_{ij}^b \cdot \exp(\alpha_{ij}^v)}{\sum_{j=1}^K \alpha_{ij}^b \cdot \exp(\alpha_{ij}^v)}, \quad (5)$$

where α_{ij}^v represents the similarity between the visual features, computed by scaled dot-product [51]:

$$\alpha_{ij}^v = (\mathbf{U} \mathbf{v}'_i)^\top \cdot \mathbf{V} \mathbf{v}'_j, \quad (6)$$

where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{d_h \times (d_q + d_v)}$ are projection matrices. α_{ij}^b measures the relative geometric position between any pair of regions:

$$\alpha_{ij}^b = \max\{0, \mathbf{w} \cdot f_b(\mathbf{b}_i, \mathbf{b}_j)\}, \quad (7)$$

where $f_b(\cdot, \cdot)$ first computes a 4-dimensional relative geometry feature $(\log(\frac{|x_i - x_j|}{w_i}), \log(\frac{|y_i - y_j|}{h_i}), \log(\frac{w_j}{w_i}), \log(\frac{h_j}{h_i}))$, then embeds it into a d_h -dimensional feature by computing cosine and sine functions of different wavelengths. $\mathbf{w} \in \mathbb{R}^{d_h}$ transforms the d_h -dimensional feature into a scalar weight, which is further trimmed at 0. Unlike how we assume no-relation for objects that are too far away from each other in the explicit relation setting, the restrictions for implicit relations are learned through \mathbf{w} and the zero trimming operation.

Explicit Relation We consider semantic relation encoder first. Since edges in the semantic graph \mathcal{E}_{sem} now contain label information and are directional, we design the attention mechanism in (3) to be sensitive to both directionality (v_i -to- v_j , v_j -to- v_i and v_i -to- v_i) and labels. Specifically,

$$\mathbf{v}_i^* = \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \cdot (\mathbf{W}_{dir(i,j)} \mathbf{v}'_j + \mathbf{b}_{lab(i,j)}) \right), \quad (8)$$

$$\alpha_{ij} = \frac{\exp((\mathbf{U} \mathbf{v}'_i)^\top \cdot \mathbf{V}_{dir(i,j)} \mathbf{v}'_j + \mathbf{c}_{lab(i,j)})}{\sum_{j \in \mathcal{N}_i} \exp((\mathbf{U} \mathbf{v}'_i)^\top \cdot \mathbf{V}_{dir(i,j)} \mathbf{v}'_j + \mathbf{c}_{lab(i,j)})},$$

where $\mathbf{W}_{\{\cdot\}}, \mathbf{V}_{\{\cdot\}}$ are matrices, and $\mathbf{b}_{\{\cdot\}}, \mathbf{c}_{\{\cdot\}}$ are bias terms. $dir(i, j)$ selects the transformation matrix wrt the directionality of each edge, and $lab(i, j)$ represents the label of each edge. Consequently, after encoding all the regions $\{\mathbf{v}'_i\}_{i=1}^K$ via the above graph attention mechanism, the refined region-level features $\{\mathbf{v}_i^*\}_{i=1}^K$ are endowed with the prior semantic relations between objects.

As opposed to graph convolutional networks, this graph attention mechanism effectively assigns different weights of importance to nodes of the same neighborhood. Combining with the question-adaptive mechanism, the learned attention weights can reflect which relations are relevant to a specific question. The relation encoder can work in the same manner on the spatial graph \mathcal{E}_{spa} , with a different set of parameters to be learned, thus details are omitted for simplicity.

3.3. Multimodal Fusion and Answer Prediction

After obtaining relation-aware visual features, we want to fuse question information \mathbf{q} with each visual representation \mathbf{v}_i^* through a multi-modal fusion strategy. Since our relation encoder preserves the dimensionality of visual features, it can be incorporated with any existing multi-modal fusion method to learn a joint representation \mathbf{J} :

$$\mathbf{J} = f(\mathbf{v}^*, \mathbf{q}; \Theta), \quad (9)$$

where f is a multi-modal fusion method and Θ are trainable parameters of the fusion module.

For the Answer Predictor, we adopt a two-layer multi-layer perceptron (MLP) as the classifier, with the joint representation \mathbf{J} as the input. Binary cross entropy is used as the loss function, similar to [2].

In the training stage, different relations encoders are trained independently. In the inference stage, we combine the three graph attention networks with a weighted sum of the predicted answer distributions. Specifically, the final answer distribution is calculated by:

$$Pr(a = a_i) = \alpha Pr_{sem}(a = a_i) + \beta Pr_{spa}(a = a_i) + (1 - \alpha - \beta) Pr_{imp}(a = a_i), \quad (10)$$

where α and β are trade-off hyper-parameters ($0 \leq \alpha + \beta \leq 1, 0 \leq \alpha, \beta \leq 1$). $Pr_{sem}(a = a_i), Pr_{spa}(a = a_i)$

and $Pr_{imp}(a = a_i)$ denote the predicted probability for answer a_i , from the model trained with semantic, spatial and implicit relations, respectively.

4. Experiments

We evaluate our proposed model on VQA 2.0 and VQA-CP v2 datasets [3, 19, 1]. In addition, Visual Genome [27] is used to pre-train the semantic relation classifier. It is also used to augment the VQA dataset when testing on the test-dev and test-std splits. We use accuracy as the evaluation metric:

$$\text{Acc}(\text{ans}) = \min\left(1, \frac{\#\text{humans provided ans}}{3}\right). \quad (11)$$

4.1. Datasets

VQA 2.0 dataset is composed of real images from MSCOCO [31] with the same train/validation/test splits. For each image, an average of 3 questions are generated. These questions are divided into 3 categories: `Y/N`, `Number` and `Other`. 10 answers are collected for each image-question pair from human annotators, and the most frequent answer is selected as the correct answer. Both open-ended and multiple-choice question types are included in this dataset. In this work, we focus on the open-ended task, and take the answers that appeared more than 9 times in the training set as candidate answers, which produces 3, 129 answer candidates. The model is trained on the training set, but when testing on the test set, both training and validation set are used for training, and the max-probable answer is selected as the predicted answer.

VQA-CP v2 dataset is a derivation of the VQA 2.0 dataset, which was introduced to evaluate and reduce the question-oriented bias in VQA models. In particular, the distribution of answers with respect to question types differs between training and test splits.

Visual Genome contains 108K images with densely annotated objects, attributes and relationships, which we used to pre-train the semantic relation classifier in our model. We filtered out those images that also appeared in the VQA validation set, and split the relation data into 88K for training, 8K for validation, and 8K for testing. Furthermore, we selected the top-14 most frequent predicates in the training data, after normalizing the predicates with relationship-aliases provided in Visual Genome. The final semantic relation classifier is trained over 14 relation classes plus a `no-relation` class.

4.2. Implementation Details

Each question is tokenized and each word is embedded using 600-dimensional word embeddings (including 300-dimensional GloVe word embeddings [42]). The sequence of embedded words is then fed into GRU for each time step

up to the 14th token (similar to [24]). Questions shorter than 14 words are padded at the end with zero vectors. The dimension of the hidden layer in GRU is set as 1024. We employ multi-head attention with 16 heads for all three graph attention networks. The dimension of relation features is set to 1024. For implicit relation, we set the embedded relative geometry feature dimension d_h to be 64.

For the semantic relation classifier, we extract pre-trained object detection features with known bounding boxes from Faster R-CNN [44] model in conjunction with ResNet-101 [20]. More specifically, the features are the output of the Pool5 layer after RoI pooling from Res4b22 feature map [58]. The Faster R-CNN model is trained over 1,600 selected object classes and 400 attribute classes, similar to the bottom-up attention [2].

Our model is implemented based on PyTorch [40]. In experiments, we use Adamax optimizer for training, with the mini-batch size as 256. For choice of learning rate, we employ the warm-up strategy [18]. Specifically, we begin with a learning rate of 0.0005, linearly increasing it at each epoch till it reaches 0.002 at epoch 4. After 15 epochs, the learning rate is decreased by 1/2 for every 2 epochs up to 20 epochs. Every linear mapping is regularized by weight normalization and dropout ($p = 0.2$ except for the classifier with 0.5).

4.3. Experimental Results

This sub-section provides experimental results on the VQA 2.0 and VQA-CP v2 datasets. By way of design, the relation encoder can be composed into different VQA architectures as a plug-and-play component. In our experiments, we consider three popular VQA models with different multimodal fusion methods: Bottom-up Top-down [2] (BUTD), Multimodal Tucker Fusion [4] (MUTAN), and Bilinear Attention Network [24] (BAN). Table 1 reports results on the VQA 2.0 validation set in the following setting:

- *Imp / Sem / Spa*: only one single type of relation (implicit, semantic or spatial) is used to incorporate bottom-up attention features.
- *Imp+Sem / Imp+Spa / Sem+Spa*: two different types of relations are used via weighted sum.
- *All*: all three types of relations are utilized, through weighted sum (e.g.: $\alpha = 0.4, \beta = 0.3$). See Eqn. (10) for details.

Compared to the baseline models, we can observe consistent performance gain for all three architectures after adding the proposed relation encoder. These results demonstrate that our ReGAT model is a generic approach that can be used to improve state-of-the-art VQA models. Furthermore, the results indicate that each single relation helps improve the performance, and pairwise combination of relations can achieve consistent performance gain. When all three types are combined, our model can achieve the best

Fusion Method	Model								
	Baseline	BiLSTM	Imp	Sem	Spa	Imp+Sem	Imp+Spa	Sem+Spa	All
BUTD [2]	63.15 (63.38 [†])	61.95	64.10	64.11	64.02	64.93	64.92	64.84	65.30
MUTAN [4]	58.16 (61.36 [†])	61.22	62.45	62.60	62.01	63.99	63.70	63.89	64.37
BAN [24]	65.36±0.14 (65.51 [†])	64.55	65.93±0.06	65.97±0.05	66.02±0.12	66.81	66.76	66.85	67.18

Table 1. Performance on VQA 2.0 validation set with different fusion methods. Consistent improvements are observed across 3 popular fusion methods, which demonstrates that our model is compatible to generic VQA frameworks. (†) Results based on our re-implementations.

Model	SOTA [6]	Baseline	Sem	Spa	Imp	All
Acc.	39.54	39.24	39.54	40.30	39.58	40.42

Table 2. Model accuracy on the VQA-CP v2 benchmark (open-ended setting on the test split).

Model	Test-dev				Test-std
	Overall	Y/N	Num	Other	
BUTD [49]	65.32	81.82	44.21	56.05	65.67
MFH [60]	68.76	84.27	50.66	60.50	-
Counter [62]	68.09	83.14	51.62	58.97	68.41
Pythia [22]	70.01	-	-	-	70.24
BAN [24]	70.04	85.42	54.04	60.52	70.35
v-AGCN [57]	65.94	82.39	56.46	45.93	66.17
Graph learner [39]	-	-	-	-	66.18
MuRel [6]	68.03	84.77	49.84	57.85	68.41
ReGAT (ours)	70.27	86.08	54.42	60.33	70.58

Table 3. Model accuracy on the VQA 2.0 benchmark (open-ended setting on the test-dev and test-std split).

performance. Our best results are achieved by combining the best single relation models through weighted sum. To verify the performance gain is significant, we performed t-test on the results of our BAN baseline and our proposed model with each single relation. We report the standard deviation in Table 1, and the p-value is 0.001459. The improvement from our method is significant at $p < 0.05$. We also compare with an additional baseline model that uses BiLSTM as the contextualized relation encoder, the results show that using BiLSTM hurts the performance.

To demonstrate the generalizability of our ReGAT model, we also conduct experiments on the VQA-CP v2 dataset, where the distributions of the training and test splits are very different from each other. Table 2 shows results on VQA-CP v2 test split. Here we use BAN with four glimpses as the baseline model. Consistent with what we have observed on VQA 2.0, our ReGAT model surpasses the baseline by a large margin. With only single relation, our model has already achieved state-of-the-art performance on VQA-CP v2 (40.30 vs. 39.54). When adding all the relations, the performance gain was further lifted to +0.88.

Table 3 shows single-model results on VQA 2.0 test-dev and test-std splits. The top five rows show results from models without relational reasoning, and the bottom four rows are results from models with relational reasoning. Our model surpasses all previous work with or without relational reasoning. Our final model uses bilinear attention with four glimpses as the multimodal fusion method. Compared to

Att.	Q-adaptive	Semantic	Spatial	Implicit	All
No	No	63.20	63.04	n/a	n/a
Yes	No	63.90	63.85	63.36	64.98
No	Yes	63.31	63.13	n/a	n/a
Yes	Yes	64.11	64.02	64.10	65.30

Table 4. Performance on VQA 2.0 validation set for ablation study (Q-adaptive: question-adaptive; Att: Attention).

BAN [24], which uses eight bilinear attention maps, our model outperforms BAN with fewer glimpses. Pythia [22] achieved 70.01 by adding additional grid-level features and using 100 object proposals from a fine-tuned Faster R-CNN on the VQA dataset for all images. Our model, without any feature augmentation used in their work, surpasses Pythia’s performance by a large margin.

4.4. Ablation Study

In Table 4, we compare three ablated instances of ReGAT with its complete form. Specifically, we validate the importance of concatenating question features to each object representation and attention mechanism. All the results reported in Table 4 are based on BUTD model architecture. To remove attention mechanism from our relation encoder, we simply replace graph attention network with graph convolution network, which can also learn node representation from graphs but with simple linear transformation.

Firstly, we validate the effectiveness of using attention mechanism to learn relation-aware visual features. Adding attention mechanism leads to a higher accuracy for all three types of relation. Comparison between line 1 and line 2 shows a gain of +0.70 for semantic relation and +0.81 for spatial relation. Secondly, we validate the effectiveness of question-adaptive relation features. Between line 1 and line 3, we see a gain of approximately +0.1 for both semantic and spatial relations. Finally, attention mechanism and question-adaptive features are added to give the complete ReGAT model. This instance gives the highest accuracy (line 4). Surprisingly, by comparing line 1 and line 4, we can observe that combining graph attention with question-adaptive gives better gain than simply adding the individual gains from the two methods. It is worth mentioning that for implicit relation, adding question-adaptive improves the model performance by +0.74, which is higher than the gain from question-adaptive for the two explicit relations. When all the relations are considered, we observe consistent performance gain by adding the question-adaptive mechanism.

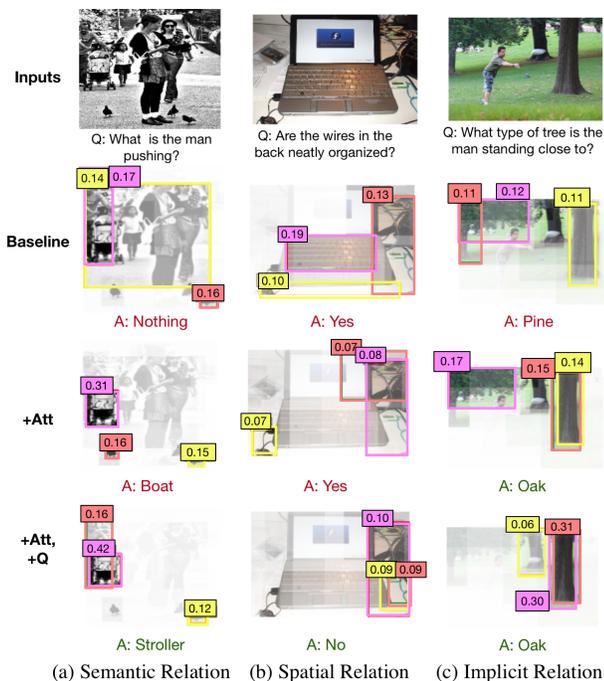


Figure 4. Visualization of attention maps learned from ablated instances: The three bounding boxes shown in each image are the top-3 attended regions. The numbers are attention weights.

To better understand how these two components help answer questions, we further visualize and compare the attention maps learned by the ablated instances in Section 4.5.

4.5. Visualization

To better illustrate the effectiveness of adding graph attention and question-adaptive mechanism, we compare the attention maps learned by the complete ReGAT model in a single-relation setting with those learned by two ablated models. As shown in Figure 4, the second, third and last rows correspond to line 1, 3 and 4 in Table 4, respectively. Comparing row 2 with row 3 leads to the observation that graph attention helps to capture the interactions between objects, which contributes to a better alignment between image regions and questions. Row 3 and row 4 show that adding the question-adaptive attention mechanism produces sharper attention maps and focuses on more relevant regions. These visualization results are consistent with the quantitative results reported in Table 4.

Figure 5 provides visualization examples on how different types of relations help improve the performance. In each example, we show the top-3 attended regions and the learned relations between these regions. As shown in these examples, each relation type contributes to a better alignment between image regions and questions. For example, in Figure 5(a), semantic relations “*Holding*” and “*Riding*” resonate with the same words that appeared in the corresponding questions. Figure 5(b) shows how spatial relations

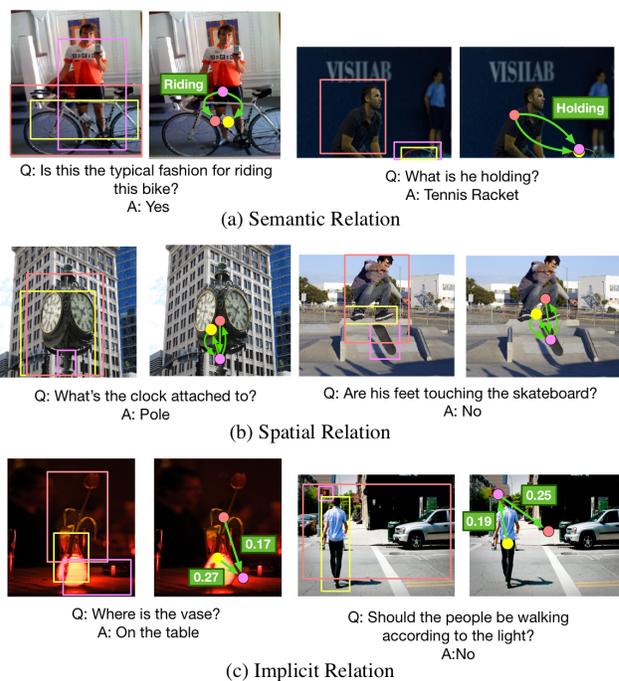


Figure 5. Visualization of different types of visual object relation in VQA task. The 3 bounding boxes shown in each image are the top-3 attended regions. Green arrows indicate relations from subject to object. Labels and numbers in green boxes are class labels for explicit relation and attention weights for implicit relation.

capture the relative geometric positions between regions.

To visualize implicit relations, Figure 5(c) shows the attention weights to the top-1 region from every other region. Surprisingly, the learned implicit relations are able to capture both spatial and semantic interactions. For example, the top image in Figure 5(c) illustrates spatial interaction “*on*” between the table and the vase, and the bottom image illustrates the semantic interaction “*walking*” between the traffic light and the person.

5. Conclusion

We have presented Relation-aware Graph Attention Network (ReGAT), a novel framework for visual question answering, to model multi-type object relations with question-adaptive attention mechanism. ReGAT exploits two types of visual object relations: Explicit Relations and Implicit Relations, to learn a relation-aware region representation through graph attention. Our method achieves state-of-the-art results on both VQA 2.0 and VQA-CP v2 datasets. The proposed ReGAT model is compatible with generic VQA models. Comprehensive experiments on two VQA datasets show that our model can be infused into state-of-the-art VQA architectures in a plug-and-play fashion. For future work, we will investigate how to fuse the three relations more effectively and how to utilize each relation to solve specific question types.

References

- [1] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Anirudha Kembhavi. Don't just assume; look and answer: Overcoming priors for visual question answering. In *CVPR*, 2018. 6
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 3, 5, 6, 7
- [3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. VQA: Visual Question Answering. In *ICCV*, 2015. 6
- [4] Hedi Ben-Younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. Mutan: Multimodal tucker fusion for visual question answering. In *ICCV*, 2017. 2, 6, 7
- [5] Irving Biederman, Robert J Mezzanotte, and Jan C Rabinowitz. Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive psychology*, 1982. 2
- [6] Remi Cadene, Hedi Ben-younes, Matthieu Cord, and Nicolas Thome. Murel: Multimodal relational reasoning for visual question answering. In *CVPR*, 2019. 3, 7
- [7] Myung Jin Choi, Antonio Torralba, and Alan S Willsky. A tree-based context model for object recognition. *PAMI*, 2012. 2
- [8] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *CVPR*, 2017. 3
- [9] Santosh K Divvala, Ali Farhadi, and Carlos Guestrin. Learning everything about anything: Webly-supervised visual concept learning. In *CVPR*, 2014. 2
- [10] Santosh K Divvala, Derek Hoiem, James H Hays, Alexei A Efros, and Martial Hebert. An empirical study of context in object detection. In *CVPR*, 2009. 2
- [11] Haoqi Fan and Jiatong Zhou. Stacked latent attention for multimodal reasoning. In *CVPR*, 2018. 1, 2
- [12] Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, et al. From captions to visual concepts and back. In *CVPR*, 2015. 2
- [13] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. Every picture tells a story: Generating sentences from images. In *ECCV*, 2010. 2
- [14] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *PAMI*, 2010. 2
- [15] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016. 2
- [16] Carolina Galleguillos, Andrew Rabinovich, and Serge Belongie. Object categorization using co-occurrence, location and appearance. In *CVPR*, 2008. 2
- [17] Stephen Gould, Jim Rodgers, David Cohen, Gal Elidan, and Daphne Koller. Multi-class segmentation with relative location prior. *IJCV*, 2008. 2
- [18] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 6
- [19] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, 2017. 6
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [21] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *CVPR*, 2018. 2, 3, 5
- [22] Yu Jiang, Vivek Natarajan, Xinlei Chen, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia v0. 1: the winning entry to the vqa challenge 2018. *arXiv preprint arXiv:1807.09956*, 2018. 7
- [23] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *CVPR*, 2015. 2
- [24] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *NeurIPS*, 2018. 2, 6, 7
- [25] Jin-Hwa Kim, Kyoung-Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Hadamard product for low-rank bilinear pooling. In *ICLR*, 2017. 2
- [26] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017. 3
- [27] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332*, 2016. 4, 6
- [28] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 3
- [29] Guohao Li, Hang Su, and Wenwu Zhu. Incorporating external knowledge to answer open-domain visual questions with dynamic memory networks. In *CVPR*, 2018. 2
- [30] Qing Li, Jianlong Fu, Dongfei Yu, Tao Mei, and Jiebo Luo. Tell-and-answer: Towards explainable visual question answering using attributes and captions. In *EMNLP*, 2018. 2
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 6
- [32] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*, 2016. 3

- [33] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *NIPS*, 2016. 1, 2
- [34] Pan Lu, Lei Ji, Wei Zhang, Nan Duan, Ming Zhou, and Jianyong Wang. R-vqa: Learning visual relation facts with semantic attention for visual question answering. In *SIGKDD*, 2018. 2
- [35] Chao Ma, Chunhua Shen, Anthony Dick, Qi Wu, Peng Wang, Anton van den Hengel, and Ian Reid. Visual question answering with memory-augmented networks. In *CVPR*, 2018. 2
- [36] Mateusz Malinowski, Carl Doersch, Adam Santoro, and Peter Battaglia. Learning visual question answering by bootstrapping hard attention. In *ECCV*, 2018. 2
- [37] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *CVPR*, 2014. 2
- [38] Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. Dual attention networks for multimodal reasoning and matching. In *CVPR*, 2017. 1, 2
- [39] Will Norcliffe-Brown, Stathis Vafeias, and Sarah Parisot. Learning conditioned graph structures for interpretable visual question answering. In *NeurIPS*, 2018. 3, 7
- [40] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [41] Badri Patro and Vinay P Nambodiri. Differential attention for visual question answering. In *CVPR*, 2018. 2
- [42] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 6
- [43] Vignesh Ramanathan, Congcong Li, Jia Deng, Wei Han, Zhen Li, Kunlong Gu, Yang Song, Samy Bengio, Charles Rosenberg, and Li Fei-Fei. Learning semantic relationships for better action retrieval in images. In *CVPR*, 2015. 2
- [44] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 6
- [45] Mohammad Amin Sadeghi and Ali Farhadi. Recognition using visual phrases. In *CVPR*, 2011. 2
- [46] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In *NIPS*, 2017. 3
- [47] Sebastian Schuster, Ranjay Krishna, Angel Chang, Li Fei-Fei, and Christopher D Manning. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *Proceedings of the fourth workshop on vision and language*, 2015. 2
- [48] Zhou Su, Chen Zhu, Yinpeng Dong, Dongqi Cai, Yurong Chen, and Jianguo Li. Learning visual knowledge memory networks for visual question answering. In *CVPR*, 2018. 2
- [49] Damien Teney, Peter Anderson, Xiaodong He, and Anton van den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In *CVPR*, 2018. 1, 2, 7
- [50] Damien Teney, Lingqiao Liu, and Anton van den Hengel. Graph-structured representations for visual question answering. In *CVPR*, 2017. 3
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 5
- [52] Peng Wang, Qi Wu, Jiewei Cao, Chunhua Shen, Lianli Gao, and Anton van den Hengel. Neighbourhood watch: Referring expression comprehension via language-guided graph attention networks. In *CVPR*, 2019. 3
- [53] Peng Wang, Qi Wu, Chunhua Shen, and Anton van den Hengel. The vqa-machine: Learning how to use existing vision algorithms to answer new questions. In *CVPR*, 2017. 2
- [54] Qi Wu, Chunhua Shen, Peng Wang, Anthony Dick, and Anton van den Hengel. Image captioning and visual question answering based on attributes and external knowledge. *PAMI*, 2017. 2
- [55] Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Ask me anything: Free-form visual question answering based on knowledge from external sources. In *CVPR*, 2016. 2
- [56] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *CVPR*, 2016. 1, 2
- [57] Zhuoqian Yang, Jing Yu, Chenghao Yang, Zengchang Qin, and Yue Hu. Multi-modal learning with prior visual relation reasoning. *arXiv preprint arXiv:1812.09681*, 2018. 3, 7
- [58] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *ECCV*, 2018. 2, 3, 4, 6
- [59] Dongfei Yu, Jianlong Fu, Tao Mei, and Yong Rui. Multi-level attention networks for visual question answering. In *CVPR*, 2017. 2
- [60] Zhou Yu, Jun Yu, Chenchao Xiang, Jianping Fan, and Dacheng Tao. Beyond bilinear: Generalized multimodal factorized high-order pooling for visual question answering. *IEEE Transactions on Neural Networks and Learning Systems*, 2018. 2, 7
- [61] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *CVPR*, 2017. 3
- [62] Yan Zhang, Jonathon S. Hare, and Adam Prügel-Bennett. Learning to count objects in natural images for visual question answering. In *ICLR*, 2018. 7
- [63] Chen Zhu, Yanpeng Zhao, Shuaiyi Huang, Kewei Tu, and Yi Ma. Structured attentions for visual question answering. In *ICCV*, 2017. 2