

Visual question answering based on local-scene-aware referring expression generation

Jung-Jun Kim^{a,1}, Dong-Gyu Lee^{b,1}, Jialin Wu^c, Hong-Gyu Jung^a, Seong-Whan Lee^{d,*}

^a Department of Brain and Cognitive Engineering, Korea University, Anam-dong, Seongbuk-gu, Seoul 02841, Republic of Korea

^b Department of Artificial Intelligence, Kyungpook National University, Daehak-ro, Buk-gu, Daegu 41566, Republic of Korea

^c Department of Computer Science, University of Texas at Austin, Austin, TX 78712, United States

^d Department of Artificial Intelligence, Korea University, Anam-dong, Seongbuk-gu, Seoul 02841, Republic of Korea

ARTICLE INFO

Article history:

Received 23 August 2020

Received in revised form 14 January 2021

Accepted 4 February 2021

Available online 24 February 2021

Keywords:

Visual question answering

Joint-embedding multi-head attention

Referring expression generation

ABSTRACT

Visual question answering requires a deep understanding of both images and natural language. However, most methods mainly focus on visual concept; such as the relationships between various objects. The limited use of object categories combined with their relationships or simple question embedding is insufficient for representing complex scenes and explaining decisions. To address this limitation, we propose the use of text expressions generated for images, because such expressions have few structural constraints and can provide richer descriptions of images. The generated expressions can be incorporated with visual features and question embedding to obtain the question-relevant answer. A joint-embedding multi-head attention network is also proposed to model three different information modalities with co-attention. We quantitatively and qualitatively evaluated the proposed method on the VQA v2 dataset and compared it with state-of-the-art methods in terms of answer prediction. The quality of the generated expressions was also evaluated on the RefCOCO, RefCOCO+, and RefCOCOg datasets. Experimental results demonstrate the effectiveness of the proposed method and reveal that it outperformed all of the competing methods in terms of both quantitative and qualitative results.

© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

Over the past few years, visual question answering (VQA) has attracted substantial attention from both the computer vision and natural language processing communities (Anderson et al., 2018; Antol et al., 2015; Ben-Younes, Cadene, Cord, & Thome, 2017; Fukui et al., 2016; Goyal, Khot, Summers-Stay, Batra, & Parikh, 2017; Lu et al., 2018; Teney, Anderson, He, & Van Den Hengel, 2018; Yu, Yu, Cui, Tao, & Tian, 2019). Compared to the traditional tasks of computer vision or natural language processing, such as object detection (Ren, He, Girshick, & Sun, 2015), image captioning (Liu, Wang, & Yang, 2017; Lu, Xiong, Parikh, & Socher, 2017; Pedersoli, Lucas, Schmid, & Verbeek, 2017; Rennie, Marcheret, Mroueh, Ross, & Goel, 2017; Yu, Tan, Bansal, & Berg, 2017), tracking (Park, Choi, Jain, & Lee, 2013; Roh, Kim, Park, & Lee, 2007), face recognition (Kang, Han, Jain, & Lee, 2014; Maeng, Liao, Kang, Lee, & Jain, 2012), action recognition (Kim, Lee, & Lee, 2020; Lee & Lee, 2019; Roh, Shin, & Lee, 2010), or machine translation (Bahdanau, Cho, & Bengio, 2015; Cho et al., 2014), the VQA is a challenging task that requires a more fine-grained

semantic understanding of both questions and images jointly, as well as common sense knowledge to answer accurately. The recently collected VQA v2 dataset (Antol et al., 2015; Goyal et al., 2017) provides complementary pairs of questions and answers. Each pair includes the same question for two semantically similar images, but the corresponding answers are different. Because the two images are semantically similar, a VQA model must be able to perform fine-grained reasoning by understanding the details of the scene context to answer questions correctly. In this paper, we propose a framework that can answer visual questions about a given image and generate human-interpretable textual explanations for answers.

Most VQA research focuses on directly exploiting visual features, such as object attributes or the relationships between objects, to understand images (Anderson et al., 2018; Hong, Fu, Uh, Mei, & Byun, 2019; Lu et al., 2018, 2017; Pedersoli et al., 2017). Some studies have utilized image captions or descriptive paragraphs to represent the semantic meaning of the image. The semantic understanding with higher-level information of the scene such as contexts and relationships which are difficult to directly extract from low-level features, can be inferred through the descriptive paragraphs. Such textual representations can contain the richer and necessary information to improve model performances (Kim & Bansal, 2019; Wu, Hu, & Mooney,

* Corresponding author.

E-mail address: sw.lee@korea.ac.kr (S.-W. Lee).

¹ Equally contributed.

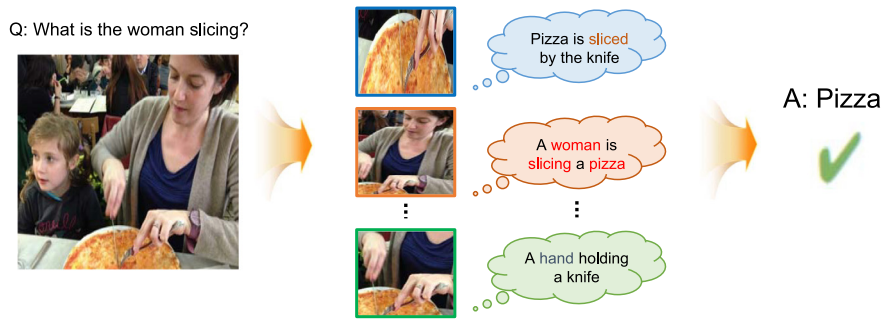


Fig. 1. Visualization of the concept of our proposed approach to VQA. Multiple object expressions for an image are generated using our local-scene-aware speaker module. Then, the VQA model answers questions based on question-relevant object expressions using one of three input modalities.

2019). Appropriate captions are highly helpful for VQA, which is demonstrated by the fact VQA models that employ human-annotated captions without images can outperform many VQA models that use image features (Li, Fu, Yu, Mei, & Luo, 2018).

However, representing an image using a single sentence has limitations in terms of expressive power, and it is easy to overlook important attributes or objects when images become complex or when a question asks a model for specific information. In terms of object captions, comparison with a few limited objects yields poor precision and can often generate non-existent objects in complex scenes (Liu et al., 2017; Mao et al., 2016; Tanaka et al., 2019; Yu, Poirson, Yang, Berg, & Berg, 2016; Yu et al., 2017).

To address this issue, we propose a novel speaker module that can produce referring expressions for each specific object. This module is tailored for answer prediction and is aware of local contextual information. The referring expression for a target object, which is likely to be a question-relevant object in an image, is generated by considering contextualized spatial neighbors. The speaker module can describe detailed local scenes, allowing a VQA model to learn target-object-relevant answers.

As shown in Fig. 1, an image can be represented by multiple text expressions for each detected object from object detector. The relationships between multiple objects are efficiently expressed in the format of text. The proposed method predicts the answer “pizza” by simultaneously considering text expressions “woman slicing pizza” as well as the attributes “woman”, “slicing”, and “pizza”. The generated expression “a woman is slicing a pizza” serves as the foundation for the final answer. Attention mechanisms have been applied in numerous uni-modal deep neural networks (Bahdanau et al., 2015; Mnih, Heess, Graves, et al., 2014; Shih, Singh, & Hoiem, 2016). In recent studies, the simultaneous learning of co-attention for visual and textual modalities has yielded fine-grained representations of images and questions for VQA tasks (Anderson et al., 2018; Fukui et al., 2016; Lu et al., 2017). Inspired by the Transformer model (Vaswani et al., 2017) and Co-attention networks (Kim, Jun, & Zhang, 2018; Nguyen & Okatani, 2018; Yu et al., 2019), we designed a joint-embedding multi-head attention (JE-MHA) network, which consists of a combination of three multi-head attention (MHA) modules for learning multi-modal embeddings. The JE-MHA network takes three types of inputs: visual feature, question and text expressions. It is because generated expressions are likely to be question-relevant representations, owing to the representation power of the speaker module, semantic understanding of a target image associated with the provided questions allows a VQA model to learn the correct answers to questions.

Overall, our primary contributions can be summarized as follows: (1) we propose a novel speaker module that can contextually understand target objects by considering neighboring objects. This allows us to generate local-scene-aware object text

expressions. (2) we designed a VQA framework with the proposed JE-MHA network; this framework can model the question-relevant text expressions with images and questions. (3) the answer prediction performance of the proposed VQA model was validated with complementary pairs. The quality of the generated explanations was also evaluated and compared to that of other methods. The results indicate that the proposed model focuses on important factors for answer prediction.

The remainder of this paper is organized as follows. In Section 2, we discuss studies related to our work. In Section 3, we provide the details of the proposed methods for VQA and expression generation. In Section 4, we present an analysis of the results of our experiments. Finally, in Section 5, we conclude this paper.

2. Related work

2.1. Visual question answering

Recent studies on deep learning based VQA systems have used bottom-up visual features (Anderson et al., 2018; Kim & Bansal, 2019; Li et al., 2018; Wu et al., 2019; Yu et al., 2019). Such systems detect object attributes using pre-trained detectors, trained on the noisy labeled Visual Genome dataset (Krishna et al., 2017). Object-level features help VQA systems attain a semantic understanding of images. One successful approach to VQA is the Transformer model (Vaswani et al., 2017) with Co-attention (Kim et al., 2018; Nguyen & Okatani, 2018; Yu et al., 2019). These methods employ a pre-trained text corpus that was collected from a larger corpus in the natural language processing field. A large number of VQA systems have adopted similar ideas of combining attention over visual features and language features to increase model capacity for question-relevant image understanding (Fukui et al., 2016; Lu, Yang, Batra, & Parikh, 2016; Xu & Saenko, 2016).

Meanwhile, more recently, some researchers have adopted text information as an input modality for VQA, such text is generated from images to supplement the visual relationships between object representations (Kim & Bansal, 2019; Li et al., 2018; Wu et al., 2019). Li et al. (2018) used a pre-trained captioner to generate general captions based on a fixed annotator. The generated captions for image were fed into an answer predictor in a VQA system. Therefore, these captions were not necessarily relevant to the target questions. Wu et al. (2019) proposed a question-relevant image captioner to generate captions that are more likely to help the VQA system answer questions. A novel caption-embedding module recognized important words in captions, and produced caption embeddings tailored for answer prediction. Kim and Bansal (2019) proposed a visual and textual question answering model to generate caption paragraphs that describe image details to use text as an additional input feature. Their model combines information from text and images via

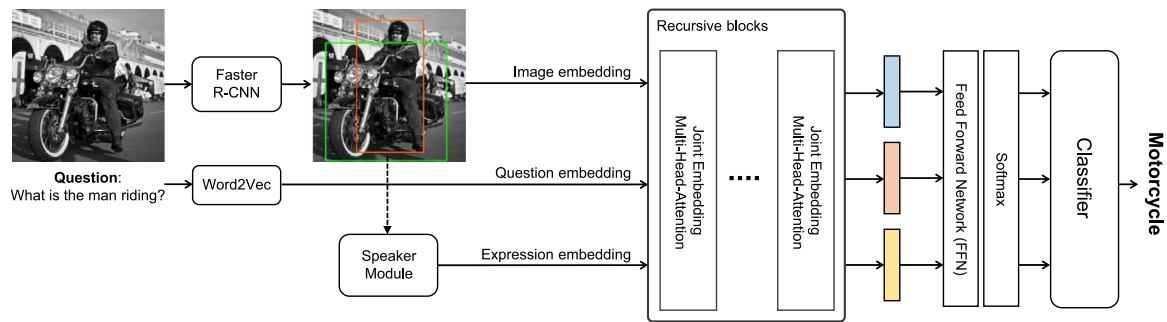


Fig. 2. Overall framework of the proposed VQA system. The JE-MHA network takes RoI pooled image embeddings, question embeddings, and generated expressions as input features. The model simultaneously learns multi-modal attention through a recursive structure and a classifier predicts final answers.

early fusion, late fusion, and later fusion. However, when a scene becomes complex, image captions can easily disregard important objects. Therefore, it is difficult to represent such scenes accurately. To address this issue, we propose the JE-MHA network, which can consider image, question and expression embeddings, simultaneously. The proposed speaker module generates referring expressions considering local-scene-aware neighboring objects. Common contexts of targets can help identify important objects and their relationships in a complex scene.

2.2. Referring expression generation

Referring expression generation is the task of generating an informative sentence for a particular object in a complex scene. Therefore, referred sentences are more unique and specific than general image captions. Most studies on referring expressions generation have directly compared the difference of visual attributes, positions and sizes between the target object and positive pairs in the scene (Liu et al., 2017; Mao et al., 2016; Tanaka et al., 2019; Yu et al., 2016, 2017). The positive pair means that the target object and its pair object belong to the same category, usually the most relevant object within an image, while the negative pair means that the pair object has different categories from the target object (Yu et al., 2016). Such methods generate sentences regarding particular objects while reducing the ambiguity of expressions between multiple objects in a scene. Therefore, large-scale referring expression generation datasets (RefCOCO, RefCOCO+, RefCOCOg) have been developed, for this task (Kazemzadeh, Ordonez, Matten, & Berg, 2014). Mao et al. (2016) introduced a max-margin maximum mutual information (MMI) training method that solves the problem of sentence ambiguity and improves the performance of both generation and comprehension tasks. Their method improved expressive power for both generation and comprehension tasks. Yu et al. (2016) proposed an explicit encoding method that compares visual feature between objects of the same category in an image. They focused on the most salient objects, rather than general positive pair objects. Liu et al. (2017) explored the role of attributes of objects and their paired descriptions. Expressions were generated by incorporating visual and learned attributes into the same semantic space. Yu et al. (2017) proposed a unified framework for both referring expression generation and comprehension tasks. A speaker module and listener module were jointly trained by minimizing the distance between paired objects and sentences mapped in an embedding space. Tanaka et al. (2019) proposed a method that can be interpreted accurately and quickly by humans. When a target is not salient, their model generates expressions by extracting information from the target and environment. A delicate understanding of scene context is essential for referring expression generation. In this paper, the representation power of detailed object relationships was maximized based on the

referring expressions of each target object. This method helps a VQA model understand the context of complex scenes. Additionally, the explanation quality of answers was also increased via referring expression generation.

3. Proposed method

We present an explainable VQA framework with local-scene-aware referring expression generation. The overall framework and speaker module are illustrated in Fig. 2 and Fig. 3, respectively. Given an image, the speaker module generates referring expressions for each object considering the context of the target object. Three different embeddings corresponding to each modality, namely region of interest (RoI) pooled image embeddings, question embeddings, and expression embeddings are used as inputs for the VQA model. A combination of MHA modules comprises the JE-MHA network, which operates recursively to obtain three final feature vectors. Finally, the answer predictor computes a confidence-level for each answer candidate.

3.1. Speaker module

In this section, we present our approach to generating object expressions that are more detailed and precise for images with crowded backgrounds. The proposed speaker module considers neighboring features to understand the target object's context. The referring expressions are generated for every object proposal of the object detector to not only bring as much additional knowledge as possible for the accurate answer but also provide a human-interpretable explanation. We also encode context features by calculating visual differences between a target and its neighboring objects. We discuss local-scene-aware expression generation for each object. The speaker module can generate detailed and precise referring expressions for complex scenes.

3.1.1. Local scene context with neighbors

Pairs of objects in the same category, which are called positive pairs, have been frequently considered to generate unambiguous referring expressions in recent studies (Liu et al., 2017; Mao et al., 2016; Tanaka et al., 2019; Yu et al., 2016, 2017). The same category means that the pair of objects belongs to the same 'supercategory'. In this paper, we use RefCOCO supercategories (Kazemzadeh et al., 2014), which are semantic concepts encompassing several objects under a common theme. For example, the supercategory 'vehicle' contains 'bicycle', 'car', 'bus', 'truck', etc. However, in real-world images, there are often many surrounding objects that are not in the same category. The relative locations or sizes of objects are frequently considered to identify positive pairs in an image, but these are insufficient for understanding the overall context of a scene. The context of the

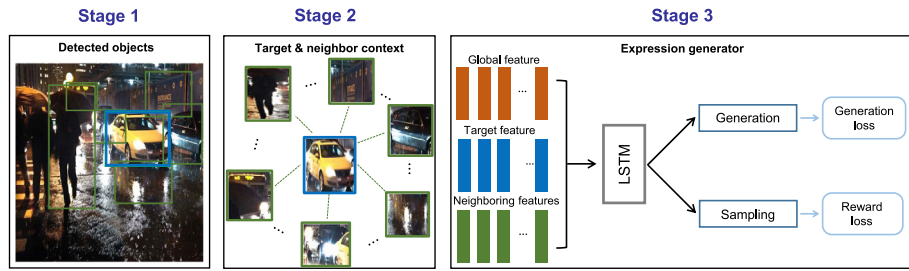


Fig. 3. High-level illustration of our speaker module. Each detected object constitutes a target object and its neighbors. The expression generator uses the global context, target object, and neighboring features to generate local-scene-aware text expressions. The LSTM is trained to minimize generation loss with reward loss.

surrounding objects can be helpful for generating high-quality referring expressions.

In this regard, we propose a local-scene-aware referring expression generation method that considers neighboring objects that are associated with target object features. Instead of utilizing only the locations and sizes of positive pairs of the target, differences in the visual features, locations, and sizes between the target and neighboring objects are extracted as visual comparison features. The explicit encoding of the visual difference between objects in the same category within an image is effective for the target context representation (Yu et al., 2016). Thus, we additionally adopt the visual comparison technique to minimize ambiguity by learning the difference between the target and its positive pairs. This allows the proposed method to understand the discriminative aspects of the target object and leads to more precise expression generation, particularly in the scene with complex background characteristics. For the object detector, we utilize the Faster-RCNN (Ren et al., 2015), to detect objects and extract local-scene-aware features, including targets and neighbors. We extract five features: the target object o_i , global context g_i from the entire image, target location/size l_i , target context (differences in visual features) δo_i , and target location/size context (differences in location/size) δl_i . The global context g_i is modeled as the averaged features of objects detected by Faster-RCNN within a scene. The target object o_i is extracted as an RoI-pooled feature using a region proposal network (Anderson et al., 2018). The location and size of a target object in an image, which are defined by the x, y coordinates of the top-left and bottom-right corners of the object, as well as its absolute size, are represented as $l_i = \begin{bmatrix} x_{tl} & y_{tl} & x_{br} & y_{br} & w_i \cdot h_i \\ W & H & W & H & W \cdot H \end{bmatrix}$ (see Fig. 3).

To identify meaningful neighbor objects around a target object, we first eliminate duplicate object detection results using non-maximum suppression (NMS). After that, we select the k -nearest neighbor objects based on Euclidean distance to extract the local scene context. Using the selected neighbors, we calculate the target context $\delta o_i = \frac{1}{k} \sum_{j \neq i} \frac{o_i - o_j}{\|o_i - o_j\|}$, where k is the number of objects selected as neighbors and the target location/size context $\delta l_i = \begin{bmatrix} \frac{[\Delta x_{tl}]_{ij}}{w_i} & \frac{[\Delta y_{tl}]_{ij}}{h_i} & \frac{[\Delta x_{br}]_{ij}}{w_i} & \frac{[\Delta y_{br}]_{ij}}{h_i} & \frac{w_j \cdot h_j}{w_i \cdot h_i} \end{bmatrix}$, which represents the relative location and size of the target object among the neighbor objects.

3.1.2. Referring expression generation

The local-scene-aware visual representation v_i that is used to generate discriminative sentences is a combination of the five aforementioned image features and is defined as $v_i = W_m[o_i, g_i, l_i, \delta o_i, \delta l_i]$. To generate object expressions for each referred object, $x_t = [v_i; w_t]$ that is a concatenation of v_i and a word embedding vector w_t is fed into a long short-term memory (LSTM) to minimize the negative log-likelihood with the parameters θ as follows:

$$L_s^1(\theta) = - \sum_i \log P(r_i | v_i; \theta). \quad (1)$$

We adopt the maximum mutual information (Mao et al., 2016) constraint that can reduce the ambiguity by minimizing the log-likelihood for positive pairs. Thus, the generated referring expressions for the target object, o_i , which is distinguished from the expressions for positive pairs. There are two constraints; (1) the ground-truth expression r_i should be more likely to be generated using the target object o_i than using any other randomly sampled object o_k . (2) the target object should be more likely to generate the ground-truth expression r_i , rather than other expressions r_j for positive pairs. The marginal loss $L_s^2(\theta)$ is calculated as follows:

$$L_s^2(\theta) = \sum_i \lambda_1^s \max(0, M_1 + \log P(r_i | v_k) - \log P(r_i | v_i)) + \lambda_2^s \max(0, M_2 + \log P(r_j | v_i) - \log P(r_i | v_i)), \quad (2)$$

where v_k denotes the visual representations of positive pairs of target objects, and λ_1^s, λ_2^s , and M_1, M_2 are marginal hyper parameters. The optimized marginal loss indicates that the generated expression is more focused on the target objects than the positive pair. It is because the MMI constraint is designed to maximize and to minimize the log-likelihood for the target object and the positive pair, respectively.

We also adopt a reinforcer module (Yu et al., 2017) to generate precise referring expressions for each target object by calculating reward loss. Specifically, a two-layered multi layer perceptron (MLP) network is used to evaluate the consistency between generated expressions and visual features. Expression representations and visual representations are combined by the MLP to generate reward loss. We use the local-scene-aware target object feature v_i as a visual feature and employ an LSTM to encode the generated expressions as sentence features. The evaluation scores are then used as rewards. The policy-gradient technique is used to optimize the reward function as follows:

$$\nabla_{\theta} J = -E_{P(w_{1:T} | v_i)} [F(w_{1:T}, v_i) \nabla_{\theta} \log P(w_{1:T} | v_i; \theta)], \quad (3)$$

where w and T denote a word vector and the number of words in a sentence, respectively. The LSTM loss is used to minimize the expression of positive pairs while maximizing the expression of the target object by combining generation loss with reward loss. Consequently, the speaker module can generate precise text expressions regarding the target object, which makes answers more accurate and increases the quality of the text expressions fed into the VQA.

3.2. Image, question, and expression embedding

In the proposed framework, the JE-MHA network takes three feature representations as inputs, namely; image embeddings, question embeddings, and expression embeddings. With regard to the image embedding, we use 36 detected object regions per image as a form of Up-down attention (Anderson et al., 2018). A Faster R-CNN head (Ren et al., 2015) with a ResNet-101 network (He, Zhang, Ren, & Sun, 2016) is employed for object

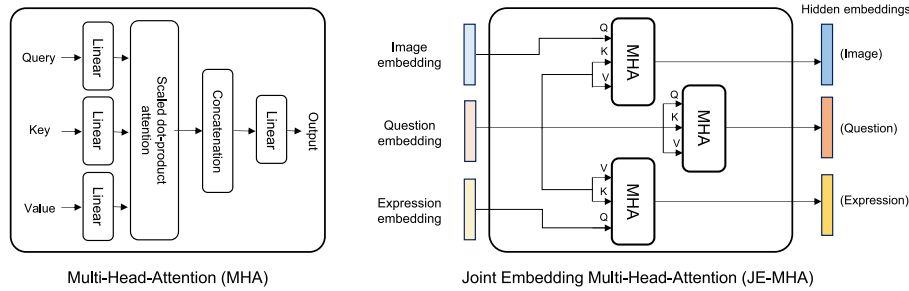


Fig. 4. (left) MHA module takes queries, keys and values as inputs. Outputs are computed as weighted sums of inputs. (right) JE-MHA network takes image, question, and expression embeddings as inputs. Each JE-MHA module outputs three different embeddings, which are used as inputs for the next module.

detection. The detection head was first pre-trained on the Visual Genome dataset (Krishna et al., 2017), which was learned to detect 1,600 object categories and 400 attributes. To generate an output set of image features \mathbf{X} , we extract RoI pooled image features from the detected regions and perform NMS for each object category. The Word2Vec method is applied for question embedding. We used a publicly available pre-trained model for the word embedding matrix of the 300-dimensional glove vectors and fine-tune weights during training (Pennington, Socher, & Manning, 2014). Expression embedding is performed using a standard gated recurrent unit (Cho et al., 2014) with 512 hidden units. The hidden unit in the final stage is extracted as the embedding feature x_c for the text expression of each object.

3.3. Joint embedding multi-head-attention network

The visual, question, and expression embeddings are fed into the JE-MHA network module to generate fused features. In this section, we elucidate the working of the JE-MHA module, which consist of a combination of three MHA modules. Multiple JE-MHAs are utilized recursively, and final results are combined through a feed-forward network (FFN).

3.3.1. Multi-head-attention

An attention function maps a query and set of key–value pairs to an output, as discussed in Vaswani et al. (2017). We compute the dot-product of the queries for all keys and divide each query by $\sqrt{d_k}$ to scale each query dimension. Then, the softmax function is applied to obtain weights for the individual values as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V. \quad (4)$$

The MHA is constructed by using linear projection to attend each query jointly based on key and value features in the same feature space. Then, a scaled dot-product is calculated h times in parallel, as follows:

$$\text{MHA}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O, \quad (5)$$

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V). \quad (6)$$

The dot-product-based attention functions (Lu, Batra, Parikh, & Lee, 2019; Lu, Goswami, Rohrbach, Parikh, & Lee, 2019; Yu et al., 2019), which consist of keys, queries, and values as input features are more effective at embedding two different input modalities compared to traditional attention techniques (e.g., additive attention) (Anderson et al., 2018; Lu et al., 2016; Teney et al., 2018). Therefore, we utilize dot-product attention for our MHA to obtain more discriminative embeddings compared to additive attention techniques.

3.3.2. JE-MHA network with image, question, and expression features

The JE-MHA network utilizes three MHA modules to construct connections between image, question, and expression embedding features based on co-attention. Let $\mathbf{Q} \in \mathbb{R}^{N \times d}$ denote a set of question features containing N tokens. Each token has d -dimensional features. Let $\mathbf{X} \in \mathbb{R}^{M \times d}$ denote the image embeddings of the d -dimensional features for M detected objects. The attention module first utilizes the image embedding features \mathbf{X} to produce query vectors q_x . It then uses the question features \mathbf{Q} to generate a key k_q and value v_q . The inner-product of the query q_x and key k_q are computed to learn the softmax result for the attention weights. The attention is then used to weight the value vectors v_q to generate weighted features, as defined in Eq. (4).

In our multi-head setting, h instances of the attended feature are computed using h different projections of the original queries, keys, and values to improve the quality of attended features (Vaswani et al., 2017). The projected vectors have d/h dimensions to ensure that the numbers of parameters are on the same scale. The output of the JE-MHA module following six recursive processes based on the outputs of each MHA is a concatenation of the h heads of attended features with dimensions of $M \times (d/h)$, resulting in a final output size of $F \in \mathbb{R}^{M \times d}$. Fig. 4 presents the overall composition of the JE-MHA network and MHA module.

3.3.3. Recursive block of the JE-MHA network

In the proposed VQA system, we combine JE-MHA modules recursively. The input embeddings \mathbf{Q}^i , \mathbf{X}_o^i , and \mathbf{X}_c^i are used as the question, image, and expression features for the i th recursion, respectively. Additionally, we use MHA and FFN_d to denote an MHA module and the FFN layer with d output units, respectively. Based on the inputs from the i th recursion, we compute the $(i + 1)$ -th recursion as follows:

$$\mathbf{Q}^{i+1} = \text{MHA}(\mathbf{Q}^i, \mathbf{Q}^i, \mathbf{Q}^i), \quad (7)$$

$$\mathbf{H}_o^{i+1} = \text{LayerNorm}(\text{MHA}(\mathbf{X}_o^i, \mathbf{Q}^i, \mathbf{Q}^i) + \mathbf{X}_o^i), \quad (8)$$

$$\mathbf{X}_o^{i+1} = \text{LayerNorm}(\text{FFN}_d(\mathbf{H}_o^{i+1}) + \mathbf{H}_o^{i+1}), \quad (9)$$

$$\mathbf{H}_c^{i+1} = \text{LayerNorm}(\text{MHA}(\mathbf{X}_c^i, \mathbf{Q}^i, \mathbf{Q}^i) + \mathbf{X}_c^i), \quad (10)$$

$$\mathbf{X}_c^{i+1} = \text{LayerNorm}(\text{FFN}_d(\mathbf{H}_c^{i+1}) + \mathbf{H}_c^{i+1}). \quad (11)$$

From the three final outputs of the recursive JE-MHA module, namely \mathbf{Q}^t , \mathbf{X}_o^t , and \mathbf{X}_c^t , we construct a classifier to predict the confidence for each answer candidate. To compute the vector representations for \mathbf{Q}^t , \mathbf{X}_o^t , and \mathbf{X}_c^t , we adopt the attention fusion approach as shown below:

$$\mathbf{q} = (\text{softmax}(\text{FFN}_1(\mathbf{Q}^t)))^T \mathbf{Q}^t, \quad (12)$$

$$\mathbf{x}_c = (\text{softmax}(\text{FFN}_1(\mathbf{X}_c^t)))^T \mathbf{X}_c^t, \quad (13)$$

$$\mathbf{x}_o = (\text{softmax}(\text{FFN}_1(\mathbf{X}_o^t)))^T \mathbf{X}_o^t, \quad (14)$$

Finally, the answer confidence is computed as a projection of the joint question, image, and expression features as follows;

$$A = \text{FFN}_{|dict|}(\mathbf{q} \odot (\mathbf{x}_c + \mathbf{x}_o)), \quad (15)$$

where $|dict|$ denotes the word corpus of training data.

4. Experiments

We validated the effectiveness of the proposed method by comparing it to state-of-the-art methods on the VQA v2.0 dataset (Antol et al., 2015) to evaluate the accuracy of answer prediction and on the RefCOCO, RefCOCO+, and RefCOCOg datasets (Kazemzadeh et al., 2014) to highlight the explainability of the proposed method.

4.1. Datasets

Visual question answering. We evaluated our proposed VQA model on the recent VQA v2.0 dataset (Antol et al., 2015). This dataset contains 1.1 million question and answer pairs based on MSCOCO images. The training, validation and test splits of the VQA 2.0 dataset were collected from the training, validation, and test splits of MSCOCO dataset, respectively. This dataset includes complementary VQA pairs that ask the same question for two semantically similar images with different answers. Because these pairs of images are semantically similar, fine-grained reasoning is required to answer the question correctly. Therefore, our explanations focus on the important factors for predicting answers that can be compared for complementary pairs. Each question and expression have fixed lengths of 14 and 20 words, respectively.

Referring expression generation. We performed experiments on three expression datasets (RefCOCO, RefCOCO+, and RefCOCOg) (Kazemzadeh et al., 2014) collected from MSCOCO images. The RefCOCO and RefCOCO+ datasets are constructed based on the Image CLEF IAPR image retrieval dataset. For the RefCOCOg dataset, all the training, validation and test splits were constructed on images sampled from the training split of the MSCOCO dataset. Since text expressions, which are generated from images, are utilized as an important component of the proposed method, we validated the quality of our generated text expressions. RefCOCO and RefCOCO+ were collected with time constraints for easily identifying target objects, whereas RefCOCOg was collected with no interactive settings and contains longer expressions than those in RefCOCO and RefCOCO+. Each dataset has 3.9 total objects and 1.63 objects of the same type per image. Overall, the three datasets contain 142,210, 141,565, and 104,560 expressions, respectively.

4.2. Implementation details

We optimized our speaker module using the Adam optimizer (Kingma & Ba, 2014) with a batch size of 128 and initialized the learning rate to $4e-4$. The learning rate was set to decay by 0.5 every 500 iterations. The hidden state size of the LSTM generator for word embeddings was set to 512. Additionally, we empirically found that considering 20 neighbors with a 0.7 IoU in the NMS stage yields optimal results. For the reinforcer, our model generates three sample sentences and estimates the corresponding rewards. During the test, we used beam search with a beam size of 10. For question and expression embedding, text preprocessing and tokenizing for UpDn features was performed using Stanford's CoreNLP (Manning et al., 2014) following (Anderson et al., 2018). The questions and expressions were first converted to lower case, trimmed to a maximum of 14 words, then tokenized based on white spaces. During the training phase, we trained our system

Table 1

VQA performance comparisons between the proposed method and competing methods on the VQA v2 test-dev and test-std split.

Method	Test-dev				Test-std
	All	Yes/No	Num	Other	All
Bottom-up (Anderson et al., 2018)	65.32	81.82	44.21	56.05	65.67
MuRel (Cadene et al., 2019)	68.03	84.77	49.84	57.85	68.41
Caption VQA (Wu et al., 2019)	66.10	82.80	44.10	56.80	68.40
MCAN VQA (Yu et al., 2019)	70.63	86.82	53.26	60.72	70.90
MFH (Yu et al., 2018)	68.02	–	–	–	–
MFH+Bottom-Up (Yu et al., 2018)	68.76	84.27	49.56	59.89	–
BANGICnt (Kim et al., 2018)	70.04	85.42	54.04	60.52	70.35
DFAF (Gao et al., 2019)	70.22	86.09	55.32	60.49	70.34
DCN (Nguyen & Okatani, 2018)	66.87	83.51	46.61	57.26	66.97
Counter (Zhang et al., 2018)	68.09	83.14	51.62	58.97	68.41
JE-MHA (ours)	69.70	86.30	50.80	59.90	70.79

on the VQA v2 training set for 13 epochs using the binary cross-entropy loss with soft scores, and BertAdam optimizer (Devlin, Chang, Lee, & Toutanova, 2019) with a batch size of 64. After 10 epochs, the learning rate decayed by 0.2 every two epochs according to the method described in Yu et al. (2019). For the evaluation of VQA tasks, during the training phase of the speaker module, we used the RefCOCOg dataset only to exploit their abundant expressions for the target object.

4.3. Results for visual question answering with explanations

In this experiment, we first evaluated the performance of our proposed VQA method on the VQA v2 dataset. We compared its performance with the state-of-the-art methods, i.e., Bottom-up (Anderson et al., 2018), MuRel (Cadene, Ben-Younes, Cord, & Thome, 2019), MCAN VQA (Yu et al., 2019), Caption VQA (Wu et al., 2019), MFH (Yu, Yu, Xiang, Fan, & Tao, 2018), MFH+Bottom-Up (Yu et al., 2018), BANGICnt (Kim et al., 2018), DFAF (Gao et al., 2019), DCN (Nguyen & Okatani, 2018), and Counter (Zhang, Hare, & Prügel-Bennett, 2018). After that, we also performed ablation studies to evaluate the contributions of input modalities and the effectiveness of using generated referring expressions.

Table 1 lists the answer prediction accuracies of each category and the corresponding average values on the VQA v2 test-dev split and test-standard split. We can see that the proposed JE-MHA VQA achieved 69.70% and 70.79% accuracy on average, while the recent state-of-the-art method, MCAN VQA (Yu et al., 2019) achieves 70.63% and 70.90% on the test-dev split and test-std split of the VQA v2 dataset, respectively. This is because the proposed method showed a little short ability to count numbers. Even though the proposed method showed slightly lower accuracy on the test-dev split, it has the advantage of a capability of human-understandable explanation generation. We believe that the explainability can be the most promising characteristic of the proposed framework. Furthermore, the proposed method showed it showed one of highest accuracy in test-std split, and the state-of-the-art performance on the validation split, of the VQA v2 dataset. We conducted a comparison with additional ablation studies on validation split of VQA v2 dataset to show the effect of text feature: We compared the JE-MHA VQA model in two different settings: (1) visual features only and (2) visual features + text features.

The experimental result demonstrates the effectiveness of each component of the proposed framework as shown in Table 2. The textual feature without visual information showed 56.8% on average, while the proposed model achieved 66.7% with both visual and textual information. From this experiment, we can see that a baseline model using the generated expressions underperforms the proposed method. The results of the Tell-Answer

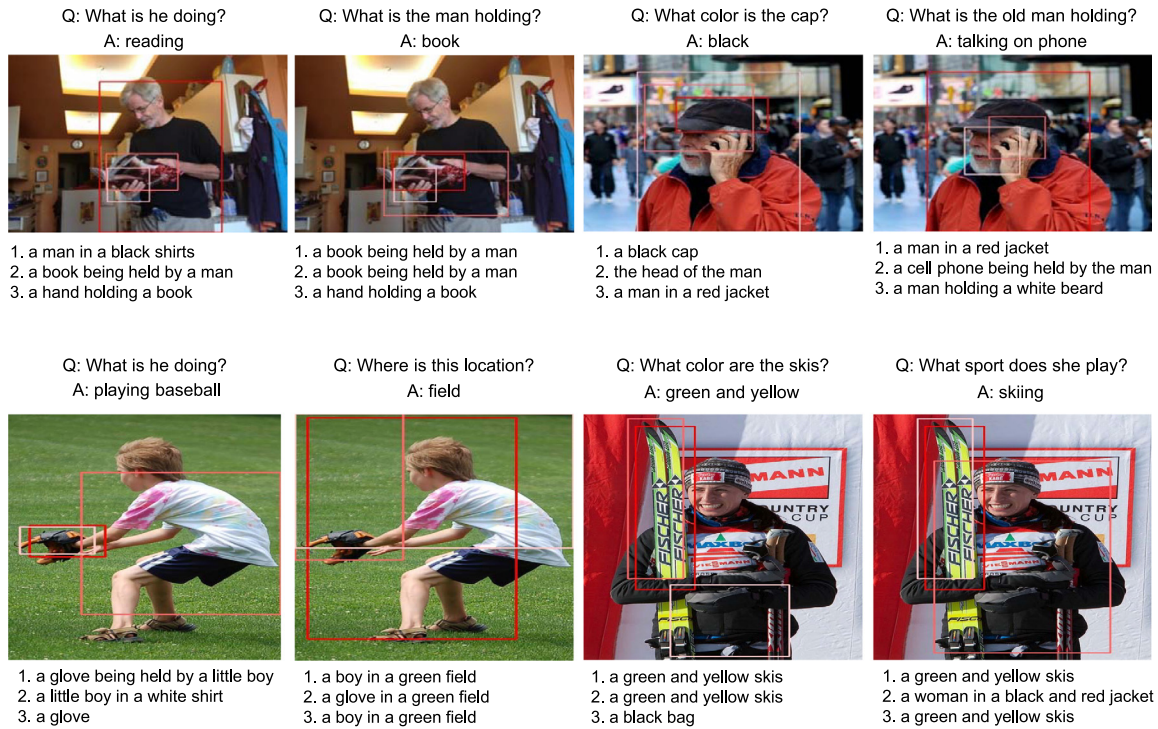


Fig. 5. Example VQA results with generated text explanations. The influential objects with the top three attention weights are indicated by bounding boxes. Each box is colored according to the importance of the corresponding visual feature, where darker red colors indicate higher weight. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

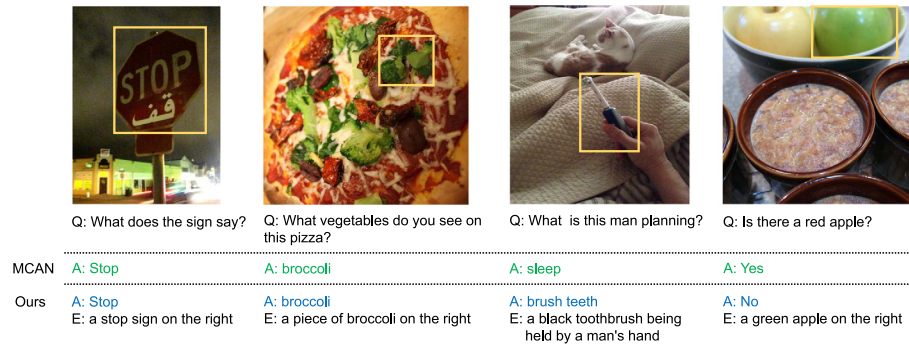


Fig. 6. Qualitative comparison to the MCAN method in terms of answer prediction. The yellow boxes indicate the most-attended objects in our model and the corresponding expressions are provided below (E:). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2

VQA performance comparisons with various input modalities on the VQA v2 validation split. “Visual” and “Textual” denote that the VQA models use visual or textual inputs, respectively.

Method	Visual	Textual	All	Yes/No	Num	Other
Tell-Answer (Li et al., 2018)		✓	54.9	76.3	36.6	42.2
JE-MHA VQA		✓	56.8	76.5	37.2	46.9
MCAN VQA (Yu et al., 2019)	✓		66.4	84.3	48.2	57.7
Caption VQA (Wu et al., 2019)	✓	✓	65.8	82.6	43.9	56.4
JE-MHA VQA (visdif+text)	✓	✓	66.7	84.6	48.3	58.0

method (Li et al., 2018) indicate that learning appropriate text expressions without visual features for VQA tasks can also yield good answer prediction performances. The proposed method also exhibits the best accuracy when only textual inputs are used. It is noteworthy that the proposed method not only provides an excellent answer prediction ability but also can provide explanations for the answers.

Fig. 6 presents comparisons between the proposed method and the MCAN (Yu et al., 2019) model. The green text in first lines represents the original MCAN predictions, the blue text in second lines represents the answers of our model, and the “E” phrases in third lines represent the explanations generated from our model. Our approach not only yields superior performance, but also makes decisions understandable for humans by presenting natural language expressions.

Additional examples of answer predictions with generated explanations are presented in Fig. 5. Two different questions are asked for each of the four images, and the information used for predicting answers is presented in order of importance. The bounding boxes represent the most influential objects and their textual expressions with the top three attention weights. The generated human-readable sentences can also provide insights for interpreting the VQA model. One can see that the bounding boxes of detected objects and generated explanations vary depending on the questions. The representation power of the



Fig. 7. Qualitative results for our referring expression generator on the RefCOCOg validation set and comparisons to the SLR model. Because the RefCOCOg dataset consist of longer sentences than RefCOCO, RefCOCO+, we trained our expression generator using the RefCOCOg dataset. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3

Comparison of our results to the state-of-the-art methods for referring expression generation for the RefCOCO, RefCOCO+, and RefCOCOg datasets. “+rerank” indicates re-ranking the generated expressions according to the listener module. “SLR” denotes the original SLR model and “re-SLR” is a re-implemented version that uses ResNet image features as inputs.

Method	Features from.	RefCOCO				RefCOCO+				RefCOCOg	
		Test A		Test B		Test A		Test B		Val	
		Meteor	CIDEr	Meteor	CIDEr	Meteor	CIDEr	Meteor	CIDEr	Meteor	CIDEr
SLR (Yu et al., 2017)	VGGNet	0.268	0.697	0.329	1.323	0.204	0.494	0.202	0.709	0.154	0.592
SLR+rerank	VGGNet	0.296	0.717	0.340	1.320	0.213	0.520	0.215	0.735	0.159	0.662
re-SLR (Tanaka et al., 2019)	VGGNet	0.279	0.729	0.334	1.315	0.201	0.491	0.211	0.757	0.146	0.679
re-SLR+rerank	VGGNet	0.278	0.717	0.332	1.262	0.198	0.476	0.206	0.721	0.150	0.676
re-SLR (Tanaka et al., 2019)	ResNet	0.296	0.804	0.341	1.358	0.220	0.579	0.221	0.798	0.153	0.742
RefGTA SR (Tanaka et al., 2019)	ResNet	0.307	0.865	0.343	1.381	0.242	0.671	0.220	0.812	0.164	0.738
RefGTA SR+rerank	ResNet	0.310	0.842	0.348	1.356	0.241	0.656	0.219	0.782	0.167	0.773
RefGTA SLR	ResNet	0.310	0.859	0.342	1.375	0.241	0.663	0.225	0.812	0.164	0.763
RefGTA SLR+rerank	ResNet	0.313	0.837	0.341	1.329	0.242	0.664	0.228	0.787	0.170	0.777
Our SR	ResNet	0.119	0.230	0.100	0.162	0.152	0.316	0.114	0.398	0.086	0.223
Our SR+rerank	ResNet	0.163	0.303	0.136	0.191	0.176	0.405	0.124	0.442	0.092	0.258
Our SLR	ResNet	0.322	0.905	0.342	1.393	0.260	0.722	0.235	0.853	0.177	0.896
Our SLR+rerank	ResNet	0.324	0.905	0.346	1.362	0.276	0.769	0.235	0.832	0.177	0.854

speaker module is highlighted by the diverse expression generation in these examples.

4.4. Results for referring expression generation

Since the proposed method aims to generate human-understandable text explanations for the answer predictions, we evaluated the quality of generated explanations by comparing them to those provided by state-of-the-art referring expression generation methods, including the speaker–listener–reinforcer (SLR) model (Yu et al., 2017), re-implemented SLR (re-SLR) (Tanaka et al., 2019) and referring expression based on grand theft auto V (RefGTA) (Tanaka et al., 2019).

The experimental results in Table 3 demonstrate that a reranking mechanism combined with a listener module improves expression quality in general, but such a combination is not very effective for the RefGTA model. In particular, the speaker–reinforcer (SR) model yields worse performance than the SLR model. However, in our expression generation module, because our listener module can locate the referred objects accurately, better expressions are generated and our method outperforms all competing methods, including the RefGTA model.

For the qualitative comparison, we also compared our generated expressions to those generated by SLR (Yu et al., 2017) on RefCOCOg dataset, which was used to generate textual expressions for VQA. Fig. 7 presents the generated expressions and target objects. One can observe that if the target object is occluded by other objects, our model understands the surrounding environment more accurately and generates less ambiguous sentences compared to SLR. This is important because our speaker module can generate referring expressions in complex situations, wherein ambiguous objects must be considered.

4.4.1. Comprehensive evaluation

In this section, we evaluate the performance of the speaker, listener, and reinforcer modules comprehensively to show the effects of each module for expression generation. Our model is compared to the RefGTA (Tanaka et al., 2019) and SLR (Yu et al., 2017) models, which are based on SLR strategies, for evaluation. We calculated scores for the reinforcer/speaker modules using the ground-truth bounding boxes for all objects given r , $o^* = \text{argmax}_i F(r, o_i)$, and $o^* = \text{argmax}_i P(r|o_i)$.

The results of our expression comprehension experiments are listed in Table 4. The listener module plays a more important role in our model than in other models. This is because our system considers the neighboring features of each target object. As a result, our method can improve the performance of the listener module by selecting a local-scene-aware context feature to generate unambiguous expressions.

4.5. Effectiveness of adjusting top-down attention


The proposed method also has an advantage in its ability to adjust top-down attention weights based on local-scene-aware context analysis. In addition to providing explanations, our model can focus and understand detailed images and questions.

Fig. 8 demonstrates that our model can attend regions in complex scenes accurately. We compared our model to the Tell-Answer model (Li et al., 2018) to verify its effectiveness. Each box in the figure above represents the most attended region for predicting a final answer. In other words, the boxes represent the main focal areas of our model. In the second example, our model also handles ambiguous references such as “yellow one” well. One can see that using generated expressions helps our model

Table 4
Comprehensive evaluations on the RefCOCO, RefCOCO+, and RefCOCOg datasets. “Ensemble” indicates using both the speaker and listener or reinforcer. Our three modules (speaker, listener, reinforcer) exhibit improved performance in most cases compared to the state-of-the-art methods.

Method	RefCOCO		RefCOCO+		RefCOCOg
	Test A	Test B	Test A	Test B	Val
SLR (ensemble) (Yu et al., 2017)	80.08%	81.73%	65.40%	60.73%	74.19%
re-SLR(ensemble) (Tanaka et al., 2019)	78.43%	81.33%	64.57%	60.48%	70.95%
re-SLR (listener)	81.14%	80.80%	68.16%	59.69%	72.36%
RefGTA SLR (listener) (Tanaka et al., 2019)	79.05%	80.31%	65.75%	62.18%	73.39%
Our SLR (listener)	82.67%	78.83%	75.80%	63.69%	78.35%
RefGTA SR (reinforcer)	80.44%	81.04%	67.81%	58.97%	74.94%
Our SR (reinforcer)	82.17%	79.04%	74.88%	62.81%	78.41%
re-SLR (speaker)	80.70%	81.71%	68.91%	60.77%	72.55%
RefGTA SR (speaker)	82.45%	82.00%	72.07%	61.06%	70.35%
RefGTA SLR (speaker)	83.05%	81.84%	72.37%	59.13%	74.79%
Our SR (speaker)	67.98%	64.83%	61.04%	48.40%	63.85%
Our SLR (speaker)	83.21%	78.55%	76.40%	63.75%	78.31%

Q: What animal is next to the man?
A: giraffe



C: Man posing in front of a pair of giraffes in background

E: A giraffe next to the man

Q: What is the yellow one?
A: bananas



C: There are bananas, pineapples, oranges, sandwiches, and drinks at the stand

E: A bunch of bananas

Fig. 8. Two examples of object-expression-based attention adjustment. The object expressions help the VQA model adjust visual attention more accurately. “C” and “E” denote image captions and explanations, respectively.

accurately attend regions and understand questions in relatively complex VQA scenes.

5. Conclusion

We proposed a novel JE-MHA framework for visual question answering with explanations. The proposed speaker module generates local-scene-aware text expressions considering the spatial context of neighboring features. The resulting high-quality object expressions are combined with visual features and question embeddings through the proposed JE-MHA network. Owing to representation power of text expressions, and co-attention of multimodal embedding, our VQA model can successfully predict answers for questions and provide rich explanations regarding scenes and questions. Experimental results showed state-of-the-art performance on the VQA v2 dataset for answer prediction. We

also showed expression quality using the RefCOCO, RefCOCO+, and RefCOCOg datasets by comparing different expression generation methods. Qualitative results for the generated expressions indicated that the proposed method can accurately answer questions when scenes are complex. In the future, we will extend our work on VQA with explanations to video data.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by Institute for Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2017-0-01779, A machine learning and statistical inference framework for explainable artificial intelligence, and No. 2019-0-00079, Artificial Intelligence Graduate School Program, Korea University).

References

Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., et al. (2018). Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6077–6086).

Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., et al. (2015). VQA: Visual question answering. In *Proceedings of the IEEE international conference on computer vision* (pp. 2425–2433).

Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In *International conference on learning representations* (pp. 1–15).

Ben-Younes, H., Cadene, R., Cord, M., & Thome, N. (2017). Mutan: Multimodal tucker fusion for visual question answering. In *Proceedings of the IEEE international conference on computer vision* (pp. 2612–2620).

Cadene, R., Ben-Younes, H., Cord, M., & Thome, N. (2019). Murel: Multimodal relational reasoning for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1989–1998).

Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., et al. (2014). Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the conference on empirical methods in natural language processing* (pp. 1724–1734).

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the annual meeting of the association for computational linguistics* (pp. 4171–4186).

Fukui, A., Park, D. H., Yang, D., Rohrbach, A., Darrell, T., & Rohrbach, M. (2016). Multimodal compact bilinear pooling for visual question answering and visual grounding. In *Proceedings of the conference on empirical methods in natural language processing* (pp. 457–468).

Gao, P., Jiang, Z., You, H., Lu, P., Hoi, S. C., Wang, X., et al. (2019). Dynamic fusion with intra-and inter-modality attention flow for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6639–6648).

Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., & Parikh, D. (2017). Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6904–6913).

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).

Hong, J., Fu, J., Uh, Y., Mei, T., & Byun, H. (2019). Exploiting hierarchical visual features for visual question answering. *Neurocomputing*, 351, 187–195.

Kang, D., Han, H., Jain, A. K., & Lee, S.-W. (2014). Nighttime face recognition at large standoff: Cross-distance and cross-spectral matching. *Pattern Recognition*, 47(12), 3750–3766.

Kazemzadeh, S., Ordonez, V., Matten, M., & Berg, T. (2014). Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing* (pp. 787–798).

Kim, H., & Bansal, M. (2019). Improving visual question answering by referring to generated paragraph captions. In *Proceedings of the annual meeting of the association for computational linguistics* (pp. 3606–3612).

- Kim, J.-H., Jun, J., & Zhang, B.-T. (2018). Bilinear attention networks. In *Advances in neural information processing systems* (pp. 1564–1574).
- Kim, Y.-J., Lee, D.-G., & Lee, S.-W. (2020). Three-stream fusion network for first-person interaction recognition. *Pattern Recognition*, 103, Article 107279.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., et al. (2017). Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1), 32–73.
- Lee, D.-G., & Lee, S.-W. (2019). Prediction of partially observed human activity based on pre-trained deep representation. *Pattern Recognition*, 85, 198–206.
- Li, Q., Fu, J., Yu, D., Mei, T., & Luo, J. (2018). Tell-and-answer: Towards explainable visual question answering using attributes and captions. In *Proceedings of the conference on empirical methods in natural language processing* (pp. 1338–1346).
- Liu, J., Wang, L., & Yang, M.-H. (2017). Referring expression generation and comprehension via attributes. In *Proceedings of the IEEE international conference on computer vision* (pp. 4856–4864).
- Lu, J., Batra, D., Parikh, D., & Lee, S. (2019). Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in neural information processing systems* (pp. 13–23).
- Lu, J., Goswami, V., Rohrbach, M., Parikh, D., & Lee, S. (2019). 12-in-1: Multi-task vision and language representation learning. arXiv preprint arXiv:1912.02315.
- Lu, P., Ji, L., Zhang, W., Duan, N., Zhou, M., & Wang, J. (2018). R-VQA: Learning visual relation facts with semantic attention for visual question answering. In *Proceedings of the international conference on knowledge discovery & data mining* (pp. 1880–1889).
- Lu, J., Xiong, C., Parikh, D., & Socher, R. (2017). Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 375–383).
- Lu, J., Yang, J., Batra, D., & Parikh, D. (2016). Hierarchical question-image co-attention for visual question answering. In *Advances in neural information processing systems* (pp. 289–297).
- Maeng, H., Liao, S., Kang, D., Lee, S.-W., & Jain, A. K. (2012). Nighttime face recognition at long distance: Cross-distance and cross-spectral matching. In *Asian conference on computer vision* (pp. 708–721).
- Manning, C. D., Surdeanu, M., Bauer, J., Finkel, J. R., Bethard, S., & McClosky, D. (2014). The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: System demonstrations* (pp. 55–60).
- Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A. L., & Murphy, K. (2016). Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 11–20).
- Mnih, V., Heess, N., Graves, A., et al. (2014). Recurrent models of visual attention. In *Advances in neural information processing systems* (pp. 2204–2212).
- Nguyen, D.-K., & Okatani, T. (2018). Improved fusion of visual and language representations by dense symmetric co-attention for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6087–6096).
- Park, U., Choi, H.-C., Jain, A. K., & Lee, S.-W. (2013). Face tracking and recognition at a distance: A coaxial and concentric PTZ camera system. *IEEE Transactions on Information Forensics and Security*, 8(10), 1665–1677.
- Pedersoli, M., Lucas, T., Schmid, C., & Verbeek, J. (2017). Areas of attention for image captioning. In *Proceedings of the IEEE international conference on computer vision* (pp. 1242–1250).
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing* (pp. 1532–1543).
- Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems* (pp. 91–99).
- Rennie, S. J., Marcheret, E., Mroueh, Y., Ross, J., & Goel, V. (2017). Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7008–7024).
- Roh, M.-C., Kim, T.-Y., Park, J., & Lee, S.-W. (2007). Accurate object contour tracking based on boundary edge selection. *Pattern Recognition*, 40(3), 931–943.
- Roh, M.-C., Shin, H.-K., & Lee, S.-W. (2010). View-independent human action recognition with volume motion template on single stereo camera. *Pattern Recognition Letters*, 31(7), 639–647.
- Shih, K. J., Singh, S., & Hoiem, D. (2016). Where to look: Focus regions for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4613–4621).
- Tanaka, M., Itamochi, T., Narioka, K., Sato, I., Ushiku, Y., & Harada, T. (2019). Generating easy-to-understand referring expressions for target identifications. In *Proceedings of the IEEE international conference on computer vision* (pp. 5794–5803).
- Teney, D., Anderson, P., He, X., & Van Den Hengel, A. (2018). Tips and tricks for visual question answering: Learnings from The 2017 challenge. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4223–4232).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998–6008).
- Wu, J., Hu, Z., & Mooney, R. (2019). Generating question relevant captions to aid visual question answering. In *Proceedings of the annual meeting of the association for computational linguistics* (pp. 3585–3594).
- Xu, H., & Saenko, K. (2016). Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *Proceedings of the European conference on computer vision* (pp. 451–466). Springer.
- Yu, L., Poirson, P., Yang, S., Berg, A. C., & Berg, T. L. (2016). Modeling context in referring expressions. In *Proceedings of the European conference on computer vision* (pp. 69–85).
- Yu, L., Tan, H., Bansal, M., & Berg, T. L. (2017). A joint speaker-listener-reinforcer model for referring expressions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7282–7290).
- Yu, Z., Yu, J., Cui, Y., Tao, D., & Tian, Q. (2019). Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6281–6290).
- Yu, Z., Yu, J., Xiang, C., Fan, J., & Tao, D. (2018). Beyond bilinear: Generalized multimodal factorized high-order pooling for visual question answering. *IEEE Transactions on Neural Networks and Learning Systems*, 29(12), 5947–5959.
- Zhang, Y., Hare, J., & Prügel-Bennett, A. (2018). Learning to count objects in natural images for visual question answering. arXiv preprint arXiv:1802.05766.