



Efficient Attention Mechanism for Visual Dialog that Can Handle All the Interactions Between Multiple Inputs

Van-Quang Nguyen^{1(✉)}, Masanori Suganuma^{1,2}, and Takayuki Okatani^{1,2}

¹ Grad School of Information Sciences, Tohoku University, Sendai, Japan
{quang,suganuma,okatani}@vision.is.tohoku.ac.jp

² RIKEN Center for AIP, Tokyo, Japan

Abstract. It has been a primary concern in recent studies of vision and language tasks to design an effective attention mechanism dealing with interactions between the two modalities. The Transformer has recently been extended and applied to several bi-modal tasks, yielding promising results. For visual dialog, it becomes necessary to consider interactions between three or more inputs, i.e., an image, a question, and a dialog history, or even its individual dialog components. In this paper, we present a neural architecture named *Light-weight Transformer for Many Inputs* (LTMI) that can efficiently deal with all the interactions between multiple such inputs in visual dialog. It has a block structure similar to the Transformer and employs the same design of attention computation, whereas it has only a small number of parameters, yet has sufficient representational power for the purpose. Assuming a standard setting of visual dialog, a layer built upon the proposed attention block has less than one-tenth of parameters as compared with its counterpart, a natural Transformer extension. The experimental results on the VisDial datasets validate the effectiveness of the proposed approach, showing improvements of the best NDCG score on the VisDial v1.0 dataset from 57.59 to 60.92 with a single model, from 64.47 to 66.53 with ensemble models, and even to 74.88 with additional finetuning.

Keywords: Visual dialog · Attention · Multimodality

1 Introduction

Recently, an increasing amount of attention has been paid to problems lying at the intersection of the vision and language domains. Many pilot tasks in this intersecting region have been designed and introduced to the research community, together with datasets. Visual dialog has been developed aiming at a higher level of vision-language interactions [7], as compared with VQA (visual question answering)

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-58586-0_14) contains supplementary material, which is available to authorized users.

[2] and VCR (visual commonsense reasoning). It extends VQA to multiple rounds; given an image and a history of question-answer pairs about the image, an agent is required to answer a new question. For example, to answer the question ‘*What color are they?*’, the agent needs to understand the context from a dialog history to know what ‘*they*’ refers to and look at the relevant image region to find out a color.

In recent studies of vision-language tasks, a primary concern has been to design an attention mechanism that can effectively deal with interactions between the two modalities. In the case of visual dialog, it becomes further necessary to consider interactions between an image, a question, and a dialog history or additionally multiple question-answer pairs in the history. Thus, the key to success will be how to deal with such interactions between three and more entities. Following a recent study [36], we will use the term *utility* to represent each of these input entities for clarity, since the term *modality* is inconvenient to distinguish between the question and the dialog history.

Existing studies have considered attention from one utility to another based on different hypotheses, such as “question \rightarrow history \rightarrow image” path in [18, 28], and “question \rightarrow image \rightarrow history \rightarrow question” path in [12, 43], etc. These methods cannot take all the interactions between utilities into account, although the missing interactions could be crucial. Motivated by this, a recent study tries to capture all the possible interactions by using a factor graph [36]. However, building the factor graph is computationally inefficient, which seemingly hinders the method from unleashing the full potential of modeling all the interactions, especially when the dialog history grows long.

The Transformer [41] has become a standard neural architecture for various tasks in the field of natural language processing, especially since the huge success of its pretrained model, BERT [11]. Its basic mechanism has recently been extended to the bi-modal problems of vision and language, yielding promising results [6, 13, 26, 27, 47]. Then, it appears to be natural to extend it further to deal with many-to-many utility interactions. However, it is not easy due to several reasons. As its basic structure is designed to be deal with self-attention, even in the simplest case of bi-modality, letting X and Y be the two utilities, there are four patterns of attention, $X \rightarrow Y$, $Y \rightarrow X$, $X \rightarrow X$, and $Y \rightarrow Y$; we need an independent Transformer block for each of these four. When extending this to deal with many-to-many utility interactions, the number of the blocks and thus of their total parameters increases proportionally with the square of the number of utilities, making it computationally expensive. Moreover, it is not apparent how to aggregate the results from all the interactions.

To cope with this, we propose a neural architecture named *Light-weight Transformer for Many Inputs* (LTMI) that can deal with all the interactions between many utilities. While it has a block structure similar to the Transformer and shares the core design of attention computation, it differs in the following two aspects. One is the difference in the implementation of multi-head attention. Multi-head attention in the Transformer projects the input feature space linearly to multiple lower-dimensional spaces, enabling to handle

multiple attention maps, where the linear mappings are represented with learnable parameters. In the proposed model, we instead split the input feature space to subspaces mechanically according to its indexes, removing all the learnable parameters from the attention computation.

The other difference from the Transformer is that LTMI is designed to receive multiple utilities and compute all the interactions to one utility from all the others including itself. This yields the same number of attended features as the input utilities, which are then concatenated in the direction of the feature space dimensions and then linearly projected back to the original feature space. We treat the parameters of the last linear projection as only learnable parameters in LTMI. This design makes it possible to retain sufficient representational power with a much fewer number of parameters, as compared with a natural extension of the Transformer block to many utilities. By using the same number of blocks in parallel as the number of utilities, we can deal with all the interactions between the utilities; see Fig. 2 for example. Assuming three utilities and the feature space dimensionality of 512, a layer consisting of LTMI has 2.38M parameters, whereas its counterpart based on naive Transformer extension has 28.4M parameters.

2 Related Work

2.1 Attention Mechanisms for Vision-Language Tasks

Attention mechanisms are currently indispensable to build neural architectures for vision-language tasks, such as VQA [4, 16, 20, 29, 31, 45, 48, 49] and visual grounding [10, 46, 52], etc. Inspired by the recent success of the Transformer for language tasks [11, 41], several studies have proposed its extensions to bi-modal vision-language tasks [6, 13, 26, 27, 40, 47]. Specifically, for VQA, it is proposed to use intra-modal and inter-modal attention blocks and stack them alternately to fuse question and image features [13]; it is also proposed to use a cascade of modular co-attention layers that compute the self-attention and guided-attention of question and image features [47]. The method of pretraining a Transformer model used in BERT [11] is employed along with Transformer extension to bi-modal tasks for several vision-language tasks [6, 26, 27]. They first pretrain the models on external datasets, such as COCO Captions [5] or Conceptual Captions dataset [38], and then fine-tune them on several target tasks.

2.2 Visual Dialog

The task of visual dialog has recently been proposed by two groups of researchers concurrently [7, 9]. De Vries et al. introduced the GuessWhat?! dataset, which is built upon goal-oriented dialogs held by two agents to identify unknown objects in an image through a set of yes/no questions [9]. Das et al. released the VisDial dataset, which is built upon dialogs consisting of pairs of a question and an answer about an image that are provided in the form of natural language texts [7]. Kottur et al. recently introduced CLEVR-Dialog as the diagnostic dataset for visual dialog [23].

Most of the existing approaches employ an encoder-decoder architecture [39]. They can be categorized into the following three groups by the design of the encoder: i) fusion-based methods, e.g., LF [7] and HRE [7], which fuses the inputs by their concatenation followed by the application of a feed-forward or recurrent network, and Synergistic [14], which fuses the inputs at multiple stages; ii) attention-based methods that compute attended features of the input image, question, and history utilities, e.g., MN [7], CoAtt [43], HClAE [28], Synergistic [14], ReDAN [12], FGA [36], and CDF [19]; ReDAN compute the attention over several reasoning steps, FGA models all the interactions over many utilities via a factor graph; iii) methods that attempt to resolve visual co-reference, e.g., RvA [32] and CorefNMN [22], which use neural modules to form an attention mechanism, DAN [18], which employs a network having two attention modules, and AMEM [37], which utilizes a memory mechanism for attention. As for the decoder, there are two designs: i) discriminative decoders that rank the candidate answers using the cross-entropy loss [7] or the n-pair loss [28]; and ii) generative decoders that yield an answer by using a MLE loss [7], weighted likelihood estimation [50], or a combination with adversarial learning [28, 43], which trains a discriminator on both positive and negative answers, then transferring it to the generator with auxiliary adversarial learning.

Other approaches include GNN [51], which models relations in a dialog by an unknown graph structure; the employment of reinforcement learning [3, 8]; and HACAN [44] which adopts policy gradient to learn the impact of history by intentionally imposing the wrong answer into dialog history. In [30, 42], pre-trained vision-language models are adopted, which consist of many Transformer blocks with hundreds of millions parameters, leading to some performance gain. Qi et al. [34] present model-agnostic principles for visual dialog to maximize performance.

3 Efficient Attention Mechanism for Many Utilities

3.1 Attention Mechanism of Transformer

As mentioned earlier, the Transformer has been applied to several bi-modal vision-language tasks, yielding promising results. The Transformer computes and uses attention from three types of inputs, Q (query), K (key), and V (value). Its computation is given by

$$\mathcal{A}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d}} \right) V, \quad (1)$$

where Q , K , and V are all collection of features, each of which is represented by a d -dimensional vector. To be specific, $Q = [q_1, \dots, q_M]^T \in \mathbb{R}^{M \times d}$ is a collection of M features; similarly, K and V are each a collection of N features, i.e., $K, V \in \mathbb{R}^{N \times d}$. In Eq. (1), V is attended with the weights computed from the similarity between Q and K .

The above computation is usually multi-plexed in the way called multi-head attention. It enables to use a number of attention distributions in parallel, aiming

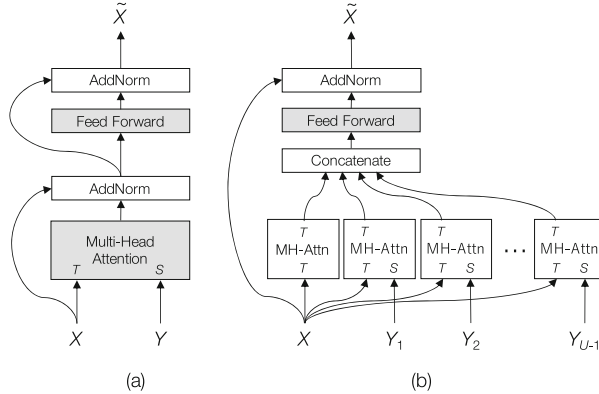


Fig. 1. (a) Source-to-target attention for bi-modal problems implemented by the standard Transformer block; the source Y is attended by weights computed from the similarity between the target X and Y . (b) The proposed block that can deal with many utilities; the source features $\{Y_1, \dots, Y_{U-1}\}$ are attended by weights computed between them and the target X . Shaded boxes have learnable weights

at an increase in representational power. The outputs of H ‘heads’ are concatenated, followed by linear transformation with learnable weights $W^O \in \mathbb{R}^{d \times d}$ as

$$\mathcal{A}^M(Q, K, V) = [\text{head}_1, \dots, \text{head}_H] W^O. \quad (2)$$

Each head is computed as follows:

$$\text{head}_h = \mathcal{A}(QW_h^Q, KW_h^K, VW_h^V), \quad h = 1, \dots, H, \quad (3)$$

where $W_h^Q, W_h^K, W_h^V \in \mathbb{R}^{d \times d_H}$ each are learnable weights inducing a linear projection from the feature space of d -dimensions to a lower space of $d_H (= d/H)$ -dimensions. Thus, one attentional block $\mathcal{A}^M(Q, K, V)$ has the following learnable weights:

$$(W_1^Q, W_1^K, W_1^V), \dots, (W_H^Q, W_H^K, W_H^V) \text{ and } W^O. \quad (4)$$

3.2 Application to Bi-modal Tasks

While Q , K , and V in NLP tasks are of the same modality (i.e., language), the above mechanism has been extended to bi-modality and applied to vision-language tasks in recent studies [6, 13, 26, 27, 40, 47]. They follow the original idea of the Transformer, considering attention from source features Y to target features X as

$$\mathcal{A}_Y(X) = \mathcal{A}^M(X, Y, Y). \quad (5)$$

In MCAN [47], language feature is treated as the source and visual feature is as the target. In [26] and others [6, 13, 27, 40], co-attention, i.e., attention in the

both directions, is considered. Self-attention, i.e., the attention from features to themselves, is given as a special case by

$$\mathcal{A}_X(X) = \mathcal{A}^M(X, X, X). \quad (6)$$

In the above studies, the Transformer block with the source-to-target attention and that with the self-attention are independently treated and are stacked, e.g., alternately or sequentially.

3.3 Light-Weight Transformer for Many Inputs

Now suppose we wish to extend the above attention mechanism to a greater number of utilities¹; we denote the number by U . If we consider every possible source-target pairs, there are $U(U - 1)$ cases in total, as there are U targets, for each of which $U - 1$ sources exist. Then we need to consider attention computation $\mathcal{A}_Y(X)$ over $U - 1$ sources Y 's for each target X . Thus, the straightforward extension of the above attention mechanism to U utilities needs $U(U - 1)$ times the number of parameters listed in Eq. (4). If we stack the blocks, the total number of parameters further increases proportionally.

To cope with this, we remove all the weights from Eq. (5). To be specific, for each head $h(= 1, \dots, H)$, we choose and freeze (W_h^Q, W_h^K, W_h^V) as

$$W_h^Q = W_h^K = W_h^V = \underbrace{[O_{d_H}, \dots, O_{d_H}]_{(h-1)d_H}}_{(h-1)d_H}, \underbrace{[I_{d_H}, O_{d_H}, \dots, O_{d_H}]_{(H-h)d_H}}_{(H-h)d_H}^\top, \quad (7)$$

where O_{d_H} is a $d_H \times d_H$ zero matrix and I_{d_H} is a $d_H \times d_H$ identity matrix. In short, the subspace for each head is determined to be one of H subspaces obtained by splitting the d -dimensional feature space with its axis indexes. Besides, we set $W^O = I$, which is the linear mapping applied to the concatenation of the heads' outputs. Let $\bar{\mathcal{A}}_Y(X)$ denote this simplified attention mechanism.

Now let the utilities be denoted by $\{X, Y_1, \dots, Y_{U-1}\}$, where $X \in \mathbb{R}^{M \times d}$ is the chosen target and others $Y_i \in \mathbb{R}^{N_i \times d}$ are the sources. Then, we compute all the source-to-target attention as $\bar{\mathcal{A}}_{Y_1}(X), \dots, \bar{\mathcal{A}}_{Y_{U-1}}(X)$. In the standard Transformer block (or rigorously its natural extensions to bi-modal problems), the attended features are simply added to the target as $X + \mathcal{A}_Y(X)$, followed by normalization and subsequent computations. To recover some of the loss in representational power due to the simplification yielding $\bar{\mathcal{A}}_Y(X)$, we propose a different approach to aggregate $\bar{\mathcal{A}}_{Y_1}(X), \dots, \bar{\mathcal{A}}_{Y_{U-1}}(X)$ and X . Specifically, we concatenate all the source-to-target attention plus the self-attention $\bar{\mathcal{A}}_X(X)$ from X to X as

$$X_{\text{concat}} = [\bar{\mathcal{A}}_X(X), \bar{\mathcal{A}}_{Y_1}(X), \dots, \bar{\mathcal{A}}_{Y_{U-1}}(X)], \quad (8)$$

¹ As we stated in Introduction, we use the term *utility* here to mean a collection of features.

where $X_{\text{concat}} \in \mathbb{R}^{M \times U d}$. We then apply linear transformation to it given by $W \in \mathbb{R}^{U d \times d}$ and $b \in \mathbb{R}^d$ with a single fully-connected layer, followed by the addition of the original X and layer normalization as

$$\tilde{X} = \text{LayerNorm}(\text{ReLU}(X_{\text{concat}}W + \mathbf{1}_M \cdot b^\top) + X), \quad (9)$$

where $\mathbf{1}_M$ is M -vector with all ones. With this method, we aim at recovery of representational power as well as the effective aggregation of information from all the utilities.

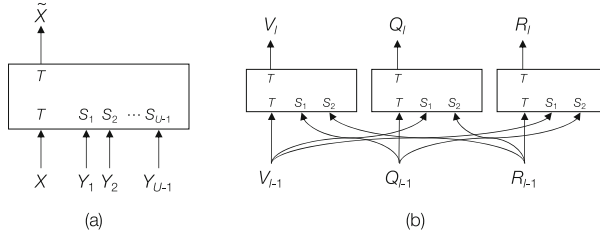


Fig. 2. (a) Simplified symbol of the proposed block shown in Fig. 1(b). (b) Its application to Visual Dialog

3.4 Interactions Between All Utilities

We have designed a basic block (Fig. 1(b)) that deals with attention from many sources to a single target. We wish to consider all possible interactions between all the utilities, not a single utility being the only target. To do this, we use U basic blocks to consider all the source-to-target attention. Using the basic block as a building block, we show how an architecture is designed for visual dialog having three utilities, visual features V , question features Q , and dialog history features R , in Fig. 2(b).

4 Implementation Details for Visual Dialog

4.1 Problem Definition

The problem of Visual Dialog is stated as follows. An agent is given the image of a scene and a dialog history containing T entities, which consists of a caption and question-answer pairs at $T - 1$ rounds. Then, the agent is further given a new question at round T along with 100 candidate answers for it and requested to answer the question by choosing one or scoring each of the candidate answers.

4.2 Representation of Utilities

We first extract features from an input image, a dialog history, and a new question at round T to obtain their representations. For this, we follow the standard method employed in many recent studies. For the image utility, we use the bottom-up mechanism [1], which extracts region-level image features using the Faster-RCNN [35] pre-trained on the Visual Genome dataset [24]. For each region (i.e., a bounding box = an object), we combine its CNN feature and geometry to get a d -dimensional vector v_i ($i = 1, \dots, K$), where K is the predefined number of regions. We then define $V = [v_1, v_2, \dots, v_K]^T \in \mathbb{R}^{K \times d}$. For the question utility, after embedding each word using an embedding layer initialized by pretrained GloVe vectors, we use two-layer Bi-LSTM to transform them to q_i ($i = 1, \dots, N$), where N is the number of words in the question. We optionally use the positional embedding widely used in NLP studies. We examine its effects in an ablation test. We then define $Q = [q_1, \dots, q_N]^T \in \mathbb{R}^{N \times d}$. For the dialog history utility, we choose to represent it as a single utility here. Thus, each of its entities represents the initial caption or the question-answer pair at one round. As with the question utility, we use the same embedding layer and a two-layer Bi-LSTM together with the positional embeddings for the order of dialog rounds to encode them with a slight difference in formation of an entity vector r_i ($i = 1, \dots, T$), where T is the number of Q&A plus the caption. We then define $R = [r_1, \dots, r_T]^T \in \mathbb{R}^{T \times d}$. More details are provided in the supplementary material.

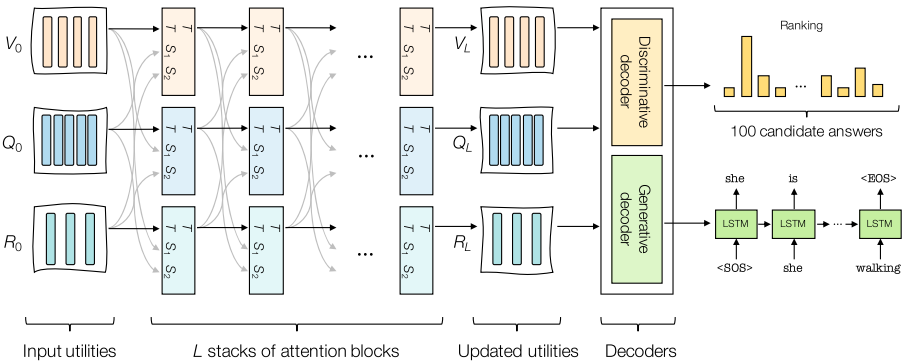


Fig. 3. The entire network built upon the proposed LTMI for Visual Dialog

4.3 Overall Network Design

Figure 3 shows the entire network. It consists of an encoder and a decoder. The encoder consists of L stacks of the proposed attention blocks; a single stack has U blocks in parallel, as shown in Fig. 2(b). We set $V_0 = V$, $Q_0 = Q$, and

$R_0 = R$ as the inputs of the first stack. After the l -th stack, the representations of the image, question, and dialog history utilities are updated as V_l , Q_l , and R_l , respectively. In the experiments, we apply dropout with the rate of 0.1 to the linear layer inside every block. There is a decoder(s) on top of the encoder. We consider a discriminative decoder and a generative decoder, as in previous studies. Their design is explained below.

4.4 Design of Decoders

Decoders receive the updated utility representations, V_L , Q_L , and R_L at their inputs. We convert them independently into d -dimensional vectors c_V , c_Q , and c_R , respectively. This conversion is performed by a simple self-attention computation. We take c_V as an example here. First, attention weights over the entities of V_L are computed by a two-layer network as

$$a_V = \text{softmax}(\text{ReLU}(V_L W_1 + \mathbf{1}_K b_1^\top) W_2 + \mathbf{1}_K b_2), \quad (10)$$

where $W_1 \in \mathbb{R}^{d \times d}$, $W_2 \in \mathbb{R}^{d \times 1}$, $b_1 \in \mathbb{R}^d$, $b_2 \in \mathbb{R}^1$, and $\mathbf{1}_K$ is K -vector with all ones. Then, c_V is given by

$$c_V = \sum_{i=1}^K v_{L,i}^\top a_{V,i}, \quad (11)$$

where $v_{L,i}$ is the i -th row vector of V_L and $a_{V,i}$ is the i -th attention weight (a scalar). The others, i.e., c_Q and c_R , can be obtained similarly.

These vectors are integrated and used by the decoders. In our implementation for visual dialog, we found that c_R does not contribute to better results; thus we use only c_V and c_Q . Note that this does not mean the dialog utility R is not necessary; it is interacted with other utilities inside the attention computation, contributing to the final prediction. The two d -vectors c_V and c_Q are concatenated as $[c_V^\top, c_Q^\top]^\top$, and this is projected to d -dimensional space, yielding a context vector $c \in \mathbb{R}^d$.

We design the discriminative and generative decoders following the previous studies. Receiving c and the candidate answers, the two decoders compute the score of each candidate answer in different ways. See details in the supplementary material.

4.5 Multi-task Learning

We observe in our experiments that accuracy is improved by training the entire network using the two decoders simultaneously. This is simply done by minimizing the sum of the losses, \mathcal{L}_D for the discriminative one and \mathcal{L}_G for the generative one (we do not use weights on the losses):

$$\mathcal{L} = \mathcal{L}_D + \mathcal{L}_G. \quad (12)$$

The increase in performance may be attributable to the synergy of learning two tasks while sharing the same encoder. Details will be given in Sect. 5.3.

5 Experimental Results

5.1 Experimental Setup

Dataset. We use the VisDial v1.0 dataset in our experiments which consists of the train 1.0 split (123,287 images), the val 1.0 split (2,064 images), and test v1.0 split (8,000 images). Each image has a dialog composed of 10 question-answer pairs along with a caption. For each question-answer pair, 100 candidate answers are given. The val v1.0 split and 2,000 images of the train v1.0 split are provided with dense annotations (i.e., relevance scores) for all candidate answers. Although the test v1.0 split was also densely annotated, the information about the ground truth answers and the dense annotations are not publicly available. Additionally, we evaluate the method on the Audio Visual Scene-aware Dialog Dataset [15]; the results are shown in the supplementary.

Evaluation Metrics. From the visual dialog challenge 2018, normalized discounted cumulative gain (NDCG) has been used as the principal metric to evaluate methods on the VisDial v1.0 dataset. Unlike other classical retrieval metrics such as R@1, R@5, R@10, mean reciprocal rank (MRR), and mean rank, which are only based on a single ground truth answer, NDCG is computed based on the relevance scores of all candidate answers for each question, which can properly handle the case where each question has more than one correct answer, such as ‘*yes it is*’ and ‘*yes*’; such cases do occur frequently.

Other Configurations. We employ the standard method used by many recent studies for the determination of hyperparameters etc. For the visual features, we detect $K = 100$ objects from each image. For the question and history features, we first build the vocabulary composed of 11,322 words that appear at least five times in the training split. The captions, questions, and answers are truncated or padded to 40, 20, and 20 words, respectively. Thus, $N = 20$ for the question utility Q . T for the history utilities varies depending on the number of dialogs. We use pre-trained 300-dimensional GloVe vectors [33] to initialize the embedding layer, which is shared for all the captions, questions, and answers.

For the attention blocks, we set the dimension of the feature space to $d = 512$ and the number of heads H in each attention block to 4. We mainly use models having two stacks of the proposed attention block. We train our models on the VisDial v0.9 and VisDial v1.0 dataset using the Adam optimizer [21] with 5 epochs and 15 epochs respectively. The learning rate is warmed up from 1×10^{-5} to 1×10^{-3} in the first epoch, then halved every 2 epochs. The batch size is set to 32 for the both datasets.

5.2 Comparison with State-of-the-Art Methods

Compared Methods. We compare our method with previously published methods on the VisDial v0.9 and VisDial v1.0 datasets, including LF, HRE, MN [7], LF-Att, MN-Att (with attention) [7], SAN [45], AMEM [37], SF [17],

HCIAE [28] and Sequential CoAttention model (CoAtt) [43], Synergistic [14], FGA [36], GNN [51], RvA [32], CorefNMN [22], DAN [18], and ReDAN [12], all of which were trained without using external datasets or data imposition. Unless noted otherwise, the results of our models are obtained from the output of discriminative decoders.

Table 1. Comparison of the performances of different methods on the validation set of VisDial v1.0 with discriminative and generative decoders.

Model	Discriminative						Generative					
	NDCG↑	MRR↑	R@1↑	R@5↑	R@10↑	Mean↓	NDCG↑	MRR↑	R@1↑	R@5↑	R@10↑	Mean↓
MN [7]	55.13	60.42	46.09	78.14	88.05	4.63	56.99	47.83	38.01	57.49	64.08	18.76
CoAtt [43]	57.72	62.91	48.86	80.41	89.83	4.21	59.24	49.64	40.09	59.37	65.92	17.86
HCIAE [28]	57.75	62.96	48.94	80.5	89.66	4.24	59.70	49.07	39.72	58.23	64.73	18.43
ReDAN [12]	59.32	64.21	50.6	81.39	90.26	4.05	60.47	50.02	40.27	59.93	66.78	17.4
LTMI	62.72	62.32	48.94	78.65	87.88	4.86	63.58	50.74	40.44	61.61	69.71	14.93

Results on the val v1.0 Split. We first compare single-model performance on the val v1.0 split. We select here MN, CoAtt, HCIAE, and ReDAN for comparison, as their performances from the both decoders in all metrics are available in the literature. To be specific, we use the accuracy values reported in [12] for a fair comparison, in which these methods are reimplemented using the bottom-up-attention features. Similar to ours, all these methods employ the standard design of discriminative and generative decoders as in [7]. Table 1 shows the results. It is seen that our method outperforms all the compared methods on the NDCG metric with large margins regardless of the decoder type. Specifically, as compared with ReDAN, the current state-of-the-art on the VisDial v1.0 dataset, our model has improved NDCG from 59.32 to 62.72 and from 60.47 to 63.58 with discriminative and generative decoders, respectively.

Results on the Test-Standard v1.0 Split. We next consider performance on the test-standard v1.0 split. In our experiments, we encountered a phenomenon that accuracy values measured by NDCG and other metrics show a trade-off relation (see the supplementary material for details), depending much on the choice of metrics (i.e., NDCG or others) for judging convergence at the training time. This is observed in the results reported in [12] and is attributable to the inconsistency between the two types of metrics. Thus, we show two results here, the one obtained using NDCG for judging convergence and the one using MRR for it; the latter is equivalent to performing early stopping.

Table 2(a) shows single-model performances on the blind test-standard v1.0 split. With the outputs from the discriminative decoder, our model gains improvement of 3.33pp in NDCG from the best model. When employing the aforementioned early stopping, our model achieves at least comparable or better performance in other metrics as well.

Many previous studies report the performance of an ensemble of multiple models. To make a comparison, we create an ensemble of 16 models with some

Table 2. Comparison in terms of (a) single- and (b) ensemble-model performance on the blind test-standard v1.0 split of the VisDial v1.0 dataset and in terms of (c) the number of parameters of the attention mechanism. The result obtained by early stopping on MRR metric is denoted by \star and those with fine-tuning on dense annotations are denoted by \dagger .

a) Performance of single models							
Model	NDCG \uparrow	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	Mean \downarrow	
LF [7]	45.31	55.42	40.95	72.45	82.83	5.95	
HRE [7]	45.46	54.16	39.93	70.45	81.50	6.41	
MN [7]	47.50	55.49	40.98	72.30	83.30	5.92	
MN-Att [7]	49.58	56.90	42.42	74.00	84.35	5.59	
LF-Att [7]	49.76	57.07	42.08	74.82	85.05	5.41	
FGA [36]	52.10	63.70	49.58	80.97	88.55	4.51	
GNN [51]	52.82	61.37	47.33	77.98	87.83	4.57	
CorefNMN [22]	54.70	61.50	47.55	78.10	88.80	4.40	
RvA [32]	55.59	63.03	49.03	80.40	89.83	4.18	
Synergistic [14]	57.32	62.20	47.90	80.43	89.95	4.17	
DAN [18]	57.59	63.20	49.63	79.75	89.35	4.30	
LTMI\star	59.03	64.08	50.20	80.68	90.35	4.05	
LTMI	60.92	60.65	47.00	77.03	87.75	4.90	

b) Performance of ensemble models							
Model	NDCG \uparrow	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	Mean \downarrow	
FGA [36]	52.10	67.30	53.40	85.28	92.70	3.54	
Synergistic [14]	57.88	63.42	49.30	80.77	90.68	3.97	
DAN [18]	59.36	64.92	51.28	81.60	90.88	3.92	
ReDAN [12]	64.47	53.73	42.45	64.68	75.68	6.63	
LTMI	66.53	63.19	49.18	80.45	89.75	4.14	
P1_P2[34] \dagger	74.91	49.13	36.68	62.98	78.55	7.03	
VD-BERT[42] \dagger	75.13	50.00	38.28	60.93	77.28	6.90	
LTMI\dagger	74.88	52.14	38.93	66.60	80.65	6.53	

c) Num. of attention parameters and the metrics scores			
Model	# params	MRR \uparrow	NDCG \uparrow
DAN [18]	12.6M	63.20	57.59
RvA [32]	11.9M	63.03	55.59
Naive Transformer	56.8M	62.09	55.10
LTMI\star (MRR-based)	4.8M	64.08	59.92
LTMI (Q, V)	4.8M	60.65	60.92
LTMI (Q, V, R)	4.8M	60.76	61.12

differences, from initialization with different random seeds to whether to use sharing weights across attention blocks or not, the number of attention blocks (i.e. $L = 2, 3$), and the number of objects in the image (i.e. $K = 50, 100$). Aiming at achieving the best performance, we also enrich the image features by incorporating the class label and attributes of each object in an image, which are also obtained from the pretrained Faster-RCNN model. Details are given in the supplementary material. We take the average of the outputs (probability distributions) from the discriminative decoders of these models to rank the candidate answers. Furthermore, we also test fine-tuning each model with its discriminative decoder on the available dense annotations from the train v1.0 and val v1.0, where the cross-entropy loss with soft labels (i.e. relevance scores) is minimized for two epochs. Table 2(b) shows the results. It is observed that our ensemble model (w/o the fine-tuning) achieves the best NDCG = 66.53 in all the ensemble models.

With optional fine-tuning, our ensemble model further gains a large improvement in NDCG, resulting in the third place in the leaderboard. The gap in NDCG to the first place (VD-BERT) is only 0.25pp, while our model yields performance that is better in all the other metrics, i.e., by 2.14pp, 5.67pp, and 3.37pp in MRR, R@5, and R@10, respectively, and 5.36% reduction in Mean.

Table 2(c) shows the number of parameters of the multi-modal attention mechanism employed in the recent methods along with their NDCG scores on the VisDial v1.0 test-standard set. We exclude the parameters of the networks computing the input utilities and the decoders, as they are basically shared among these methods. ‘Naive Transformer’ consists of two stacks of transformer blocks with simple extension to three utilities as mentioned in Sect. 1. The efficiency

of our models can be observed. Note also that the gap between (Q, V) and (Q, V, R) is small, contrary to the argument in [34].

Table 3. Ablation study on the components of our method on the val v1.0 split of VisDial dataset. \uparrow indicates the higher the better.

(a)					(b)				
Component	Details	A-NDCG \uparrow	D-NDCG \uparrow	G-NDCG \uparrow	Component	Details	A-NDCG \uparrow	D-NDCG \uparrow	G-NDCG \uparrow
Number of attention blocks	1	65.37	62.06	62.95	Context feature aggregation	[Q]	65.12	61.50	63.19
	2	65.75	62.72	63.58		[Q, V]	65.75	62.72	63.58
	3	65.42	62.48	63.22		[Q, V, R]	65.53	62.37	63.38
Self-Attention	No	65.38	61.76	63.31	Decoder Type	Gen	-	-	62.35
	Yes	65.75	62.72	63.58		Disc	-	61.80	-
						Both	65.75	62.72	63.58
Attended features aggregation	Add	64.12	60.28	61.49	The number of objects in an image	36	65.25	62.40	63.08
	Concat	65.75	62.72	63.58		50	65.24	62.29	63.12
						100	65.75	62.72	63.58
Shared Attention weights	No	65.75	62.72	63.58	Positional and spatial embeddings	No	65.18	61.84	62.96
	Yes	65.57	62.50	63.24		Yes	65.75	62.72	63.58

5.3 Ablation Study

To evaluate the effect of each of the components of our method, we perform the ablation study on the val v1.0 split of VisDial dataset. We evaluate here the accuracy of the discriminative decoder and the generative decoder separately. We denote the former by D-NDCG and the latter by G-NDCG, and the accuracy of their averaged model by A-NDCG (i.e., averaging the probability distributions over the candidate answers obtained by the discriminative and generative decoders). The results are shown in Table 3(a–b).

The first block of Table 3(a) shows the effect of the number of stacks of the proposed attention blocks. We observe that the use of two to three stacks achieves good performance on all three measures. More stacks did not bring further improvement, and thus are omitted in the table.

The second block of Table 3(a) shows the effect of self-attention, which computes the interaction within a utility, i.e., $\bar{\mathcal{A}}_X(X)$. We examine this because it can be removed from the attention computation. It is seen that self-attention does contribute to good performance. The third block shows the effects of how to aggregate the attended features. It is seen that their concatenation yields better performance than their simple addition. The fourth block shows the impact of sharing the weights across the stacks of the attention blocks. If the weights can be shared as in [25], it contributes a further decrease in the number of parameters. We observe that the performance does drop if weight sharing is employed, but the drop is not very large.

The first block of Table 3(b) shows the effect of how to aggregate the context features c_V , c_Q , and c_R in the decoder(s), which are obtained from the outputs of our encoder. As mentioned above, the context vector c_R of the dialog history does not contribute to the performance. However, the context vector c_v of the image is important for achieving the best performance. The second block of

Table 3(b) shows the effects of simultaneously training the both decoders (with the entire model). It is seen that this contributes greatly to the performance; this indicates the synergy of learning two tasks while sharing the encoder, resulting better generalization as compared with those trained with a single decoder.

We have also confirmed that the use of fewer objects leads to worse results. Besides, the positional embedding for representing the question and history utilities as well as the spatial embedding (i.e., the bounding box geometry of objects) for image utility representation have a certain amount of contribution.

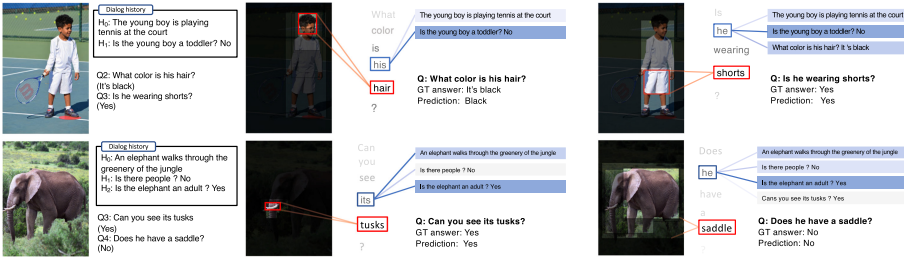


Fig. 4. Examples of visualization for the attention weights generated in our model at two Q&A rounds on two images. See Sect. 5.4 for details.

5.4 Visualization of Generated Attention

Figure 4 shows attention weights generated in our model on two rounds of Q&A on two images. We show here two types of attention. One is the self-attention weights used to compute the context vectors c_V and c_Q . For c_V , the attention weights a_V are generated over image regions (i.e., bounding boxes), as in Eq. (10). Similarly, for c_Q , the attention weights are generated over question words. These two sets of attention weights are displayed by brightness of the image bounding-boxes and darkness of question words, respectively, in the center and the rightmost columns. It can be observed from these that the relevant regions and words are properly highlighted at each Q&A round.

The other attention we visualize is the source-to-target attention computed inside the proposed block. We choose here the image-to-question attention $\bar{A}_V(Q)$ and the history-to-question attention $\bar{A}_R(Q)$. For each, we compute the average of the attention weights over all the heads computed inside the block belonging to the upper stack. In Fig. 4, the former is displayed by the red boxes connected between an image region and a question word; only the region with the largest weight is shown for the target word; the word with the largest self-attention weight is chosen for the target. The history-to-question attention is displayed by the Q&As highlighted in blue color connected to a selected question word that is semantically ambiguous, e.g., ‘its’, ‘he’, and ‘his’. It is seen that the model performs proper visual grounding for the important words, ‘hair’, ‘shorts’,

and *'tusks'*. It is also observed that the model properly resolves the co-reference for the words, *'he'* and *'its'*.

6 Summary and Conclusion

In this paper, we have proposed LTMI (Light-weight Transformer for Many Inputs) that can deal with all the interactions between multiple input utilities in an efficient way. As compared with other methods, the proposed architecture is much simpler in terms of the number of parameters as well as the way of handling inputs (i.e., their equal treatment), and nevertheless surpasses the previous methods in accuracy; it achieves the new state-of-the-art results on the VisDial datasets, e.g., high NDCG scores on the VisDial v1.0 dataset. Thus, we believe our method can be used as a simple yet strong baseline.

Acknowledgments. This work was partly supported by JSPS KAKENHI Grant Number JP15H05919 and JP19H01110.

References

1. Anderson, P., et al.: 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 (2018)
2. Antol, S., et al.: VQA: visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2425–2433 (2015)
3. Chattopadhyay, P., et al.: Evaluating visual conversational agents via cooperative human-AI games. In: Proceedings of AAAI Conference on Human Computation and Crowdsourcing (2017)
4. Chen, K., Wang, J., Chen, L.C., Gao, H., Xu, W., Nevatia, R.: ABC-CNN: an attention based convolutional neural network for visual question answering. arXiv preprint [arXiv:1511.05960](https://arxiv.org/abs/1511.05960) (2015)
5. Chen, X., et al.: Microsoft COCO captions: data collection and evaluation server. arXiv preprint [arXiv:1504.00325](https://arxiv.org/abs/1504.00325) (2015)
6. Chen, Y.C., et al.: UNITER: learning universal image-text representations. arXiv preprint [arXiv:1909.11740](https://arxiv.org/abs/1909.11740) (2019)
7. Das, A., et al.: Visual dialog. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 326–335 (2017)
8. Das, A., Kottur, S., Moura, J.M., Lee, S., Batra, D.: Learning cooperative visual dialog agents with deep reinforcement learning. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2951–2960 (2017)
9. De Vries, H., Strub, F., Chandar, S., Pietquin, O., Larochelle, H., Courville, A.: Guesswhat?! visual object discovery through multi-modal dialogue. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5503–5512 (2017)
10. Deng, C., Wu, Q., Wu, Q., Hu, F., Lyu, F., Tan, M.: Visual grounding via accumulated attention. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7746–7755 (2018)

11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
12. Gan, Z., Cheng, Y., Kholy, A.E., Li, L., Liu, J., Gao, J.: Multi-step reasoning via recurrent dual attention for visual dialog. In: Proceedings of the Conference of the Association for Computational Linguistics, pp. 6463–6474 (2019)
13. Gao, P., et al.: 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 (2019)
14. Guo, D., Xu, C., Tao, D.: Image-question-answer synergistic network for visual dialog. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 10434–10443 (2019)
15. Hori, C., et al.: End-to-end audio visual scene-aware dialog using multimodal attention-based video features. In: ICASSP 2019–2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2352–2356 (2019)
16. Ilievski, I., Yan, S., Feng, J.: A focused dynamic attention model for visual question answering. arXiv preprint [arXiv:1604.01485](https://arxiv.org/abs/1604.01485) (2016)
17. Jain, U., Lazebnik, S., Schwing, A.G.: Two can play this game: visual dialog with discriminative question generation and answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5754–5763 (2018)
18. Kang, G.C., Lim, J., Zhang, B.T.: Dual attention networks for visual reference resolution in visual dialog. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, pp. 2024–2033 (2019)
19. Kim, H., Tan, H., Bansal, M.: Modality-balanced models for visual dialogue. arXiv preprint [arXiv:2001.06354](https://arxiv.org/abs/2001.06354) (2020)
20. Kim, J.H., Jun, J., Zhang, B.T.: Bilinear attention networks. In: Advances in Neural Information Processing Systems, pp. 1564–1574 (2018)
21. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
22. Kottur, S., Moura, J.M.F., Parikh, D., Batra, D., Rohrbach, M.: Visual coreference resolution in visual dialog using neural module networks. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11219, pp. 160–178. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01267-0_10
23. Kottur, S., Moura, J.M., Parikh, D., Batra, D., Rohrbach, M.: CLEVR-dialog: a diagnostic dataset for multi-round reasoning in visual dialog. arXiv preprint [arXiv:1903.03166](https://arxiv.org/abs/1903.03166) (2019)
24. Krishna, R., et al.: Visual genome: connecting language and vision using crowd-sourced dense image annotations. *Int. J. Comput. Vis.* **123**(1), 32–73 (2017)
25. Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., Soricut, R.: ALBERT: a lite BERT for self-supervised learning of language representations. arXiv preprint [arXiv:1909.11942](https://arxiv.org/abs/1909.11942) (2019)
26. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: VisualBERT: a simple and performant baseline for vision and language. arXiv preprint [arXiv:1908.03557](https://arxiv.org/abs/1908.03557) (2019)
27. Lu, J., Batra, D., Parikh, D., Lee, S.: ViLBERT: pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. arXiv preprint [arXiv:1908.02265](https://arxiv.org/abs/1908.02265) (2019)
28. Lu, J., Kannan, A., Yang, J., Parikh, D., Batra, D.: Best of both worlds: transferring knowledge from discriminative learning to a generative visual dialog model. In: Advances in Neural Information Processing Systems, pp. 314–324 (2017)

29. Lu, J., Yang, J., Batra, D., Parikh, D.: Hierarchical question-image co-attention for visual question answering. In: *Advances in Neural Information Processing Systems*, pp. 289–297 (2016)
30. Murahari, V., Batra, D., Parikh, D., Das, A.: Large-scale pretraining for visual dialog: a simple state-of-the-art baseline. arXiv preprint [arXiv:1912.02379](https://arxiv.org/abs/1912.02379) (2019)
31. Nguyen, D.K., Okatani, T.: 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 (2018)
32. Niu, Y., Zhang, H., Zhang, M., Zhang, J., Lu, Z., Wen, J.R.: Recursive visual attention in visual dialog. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6679–6688 (2019)
33. Pennington, J., Socher, R., Manning, C.: Glove: global vectors for word representation. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 1532–1543 (2014)
34. Qi, J., Niu, Y., Huang, J., Zhang, H.: Two causal principles for improving visual dialog. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10860–10869 (2020)
35. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems*, pp. 91–99 (2015)
36. Schwartz, I., Yu, S., Hazan, T., Schwing, A.G.: Factor graph attention. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2039–2048 (2019)
37. Seo, P.H., Lehmann, A., Han, B., Sigal, L.: Visual reference resolution using attention memory for visual dialog. In: *Advances in Neural Information Processing Systems*, pp. 3719–3729 (2017)
38. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: a cleaned, hypemymed, image alt-text dataset for automatic image captioning. In: *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 2556–2565 (2018)
39. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: *Advances in Neural Information Processing Systems*, pp. 3104–3112 (2014)
40. Tan, H., Bansal, M.: LXMERT: learning cross-modality encoder representations from transformers. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (2019)
41. Vaswani, A., et al.: Attention is all you need. In: *Advances in Neural Information Processing Systems*, pp. 5998–6008 (2017)
42. Wang, Y., Joty, S., Lyu, M.R., King, I., Xiong, C., Hoi, S.C.: VD-BERT: a unified vision and dialog transformer with BERT. arXiv preprint [arXiv:2004.13278](https://arxiv.org/abs/2004.13278) (2020)
43. Wu, Q., Wang, P., Shen, C., Reid, I., van den Hengel, A.: Are you talking to me? Reasoned visual dialog generation through adversarial learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6106–6115 (2018)
44. Yang, T., Zha, Z.J., Zhang, H.: Making history matter: history-advantage sequence training for visual dialog. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2561–2569 (2019)
45. Yang, Z., He, X., Gao, J., Deng, L., Smola, A.: Stacked attention networks for image question answering. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 21–29 (2016)

46. Yu, L., et al.: MAttNet: modular attention network for referring expression comprehension. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1307–1315 (2018)
47. Yu, Z., Yu, J., Cui, Y., Tao, D., Tian, Q.: Deep modular co-attention networks for visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6281–6290 (2019)
48. Yu, Z., Yu, J., Fan, J., Tao, D.: Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1821–1830 (2017)
49. Yu, Z., Yu, J., Xiang, C., Fan, J., Tao, D.: Beyond bilinear: generalized multimodal factorized high-order pooling for visual question answering. *IEEE Trans. Neural Netw. Learn. Syst.* **29**(12), 5947–5959 (2018)
50. Zhang, H., et al.: Generative visual dialogue system via weighted likelihood estimation. In: Proceedings of the International Joint Conference on Artificial Intelligence, pp. 1025–1031 (2019)
51. Zheng, Z., Wang, W., Qi, S., Zhu, S.C.: Reasoning visual dialogs with structural and partial observations. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6669–6678 (2019)
52. Zhuang, B., Wu, Q., Shen, C., Reid, I., van den Hengel, A.: Parallel attention: a unified framework for visual object discovery through dialogs and queries. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4252–4261 (2018)