

Image-Question-Answer Synergistic Network for Visual Dialog

Dalu Guo, Chang Xu, Dacheng Tao
 UBTECH Sydney AI Centre, School of Computer Science, FEIT,
 University of Sydney, Darlington, NSW 2008, Australia
 {dguo8417@uni., c.xu@, dacheng.tao@}sydney.edu.au

Abstract

The image, question (combined with the history for de-referencing), and the corresponding answer are three vital components of visual dialog. Classical visual dialog systems integrate the image, question, and history to search for or generate the best matched answer, and so, this approach significantly ignores the role of the answer. In this paper, we devise a novel image-question-answer synergistic network to value the role of the answer for precise visual dialog. We extend the traditional one-stage solution to a two-stage solution. In the first stage, candidate answers are coarsely scored according to their relevance to the image and question pair. Afterward, in the second stage, answers with high probability of being correct are re-ranked by synergizing with image and question. On the Visual Dialog v1.0 dataset, the proposed synergistic network boosts the discriminative visual dialog model to achieve a new state-of-the-art of 57.88% normalized discounted cumulative gain. A generative visual dialog model equipped with the proposed technique also shows promising improvements.

1. Introduction

Visual dialog is an emerging research theme lying at the intersection between computer vision and natural language processing. Given the capacities of reasoning, grounding, recognition, and translating, a visual dialog agent is expected to answer questions based on an image, caption, and history. For example, in Figure 1, the agent first reasons what the word ‘their’ refers to in the current question based on history, locates the bounding boxes of the five people in the picture, recognizes the color of their jackets, and then translates the visual information into human language. Hence, a visual dialog task can also be regarded as: (i) visual grounding [33], which further converts visual information in located bounding boxes into human language; (ii) visual question answering (VQA) [2], which includes extra dialog history and caption as the input; and (iii) image captioning [25], which generates a description not only based on visual information but also the history and question.



Figure 1. A general visual dialog task. The correct answer is picked from a candidate set by investigating the given image, caption, history, and question.

A general visual dialog model has two components: an encoder to embed inputs (e.g., images and questions) into vectors and fuse them to create a unified representation, and a decoder either to translate the encoded vector directly into words for an answer or to rank the given candidate answers. Both VQA and visual dialog involve the fusion job of multiple modalities. However, as a multi-turn VQA task, in each turn, a visual dialog system must also integrate the caption and dialog history from past turns. Visual dialog systems can be grouped into two main categories according to different decoders: generative models and discriminative models. Generative models usually employ seq2seq [24] or advanced reinforcement learning [29] techniques to generate the answer set, where the highest probability answer is cho-

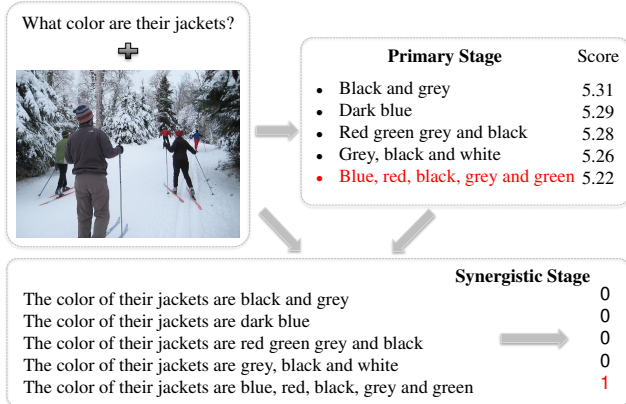


Figure 2. Candidate answers synergized with image and question are re-scored. The synergistic stage refers the answers back to the question and re-matches the image.

sen as the output. Discriminative models tend to calculate the similarity between the latent output of the encoder and the embedding of candidate answers [4, 18, 12, 29], with the correct answer expected to have the highest score. However, existing generative models aim to generate a single word of high probability at each step but omit the meaning of the whole answer sentence. The discriminative models are beneficial to understanding the answer sentence through long short-term memory (LSTM) [10], but the scoring method is insufficient to capture the similarity between inputs and answers, since the vector of inputs and answers have been separately learned without deep fusion. Furthermore, both generative and discriminative models tend to give short and safe answer, such as ‘Yes’ or ‘No’, as their fusion methods focus on the major signal in short answer but will not look into details in a longer one.

To highlight the role of answer and its integration with other ingredients (e.g., images and questions) in visual dialog, we propose an image-question-answer synergistic network. However, not all answers are plausible for the image and question. According to the experimental results of [4], recall at 10 (R@10) can already be higher than 85%, which implies that the answers ranked below 10 are unlikely to be correct. For example, in Figure 1, only answers about color in the image are related to the question, while the others are unreasonable and probably lead to network side effects. Hence, we extend the traditional one-stage model composed of an encoder and a decoder into a two-stage model containing a primary stage and a synergistic stage. The primary stage can be any existing model that coarsely scores all candidate answers or generates some high probability candidates. The synergistic stage selects answers of high probability based on certain policies, finely synergizes them with the question and then re-ranks the candidate answers according to the relevance of the synergies to the image, as shown in Figure 2. The proposed method is consistent with human practice. In multiple choice examinations,

we usually first exclude obviously wrong answers before paying more effort to compare the remaining answers that are more likely to be correct. We fill each answer into the question blank and judge whether the complete sentence is more suitable than the others. In addition, we address the class imbalance problem in the primary stage for a strong discriminative model. As a large number of easy negative candidate answers dominate the loss function, a temperature factor is considered in the loss function to discount the contribution of easy samples.

Our model is evaluated on the Visual Dialog v1.0 dataset [4]. In validation, our primary stage with a loss-balanced discriminative model improves the mean reciprocal rank (MRR) by 0.71% compared to the non-balanced model, and the synergistic stage gives additional 0.91% improvement on MRR. Furthermore, the synergistic stage in our generative model improves MRR by 4.7% and recall on top-5 responses (R@5) by 9.2% compared to the primary stage, which provides a different way to generate descriptive answers other than GAN and reinforcement learning [29]. On the test-standard dataset, our two-stage model outperforms the baselines and achieves state-of-art performance, higher than the other entries in the Visual Dialog Challenge 2018 with a 57.88% normalized discounted cumulative gain (NDCG).

2. Related Work

Visual Question Answering (VQA): VQA is the first task undertaken when querying images for text answer generation. It is a classification problem in which candidate answers are restricted to the most common answers appearing in the dataset. Current models can be classified into three main categories: early fusion models, later fusion models, and external knowledge-based models. In early fusion models, the input queries are regarded as parameters of conditional batch normalization [6, 22] in the detection network [9], which leads a pre-trained ResNet to the proposed MODERN architecture; this method affects less than 1% of parameters in the pre-trained model, which reduces the risk of over-fitting. Later fusion models mainly concentrate on how to represent the answer vector by jointing the question and global image feature [11, 2]. However, a lot of visual information is irrelevant to the input query, resulting in significant noise during reasoning. Thus, the attention mechanism is introduced to solve this problem. This starts with the linear combination of two features, such as the stacked attention network [30] by learning the attention using multi-step reasoning and the dual attention network [20] that learns both the visual and textual attention. Then, a bilinear pooling method like multi-modal compact bilinear pooling (MCB) [7] is applied by projecting the outer product of two features into the high dimension of quadratic expansion. However, MCB needs to sam-

ple features, which is computationally intractable and has very large projected dimension, so low-rank bilinear pooling (MLB) [15] and multi-factor bilinear pooling (MFB) [32] have been proposed to project the two features into a common low-rank space, and the bilinear attention network (BAN) [14] builds the interaction attention between multi-modal inputs. In contrast to early and later fusion models, external knowledge-based models suppose that common sense or information not given in the image is required to infer the right answer. [28] uses DBpedia [3] to broaden the range of answers. [21] queries the triplet (visual concept, relation, attribute) in Fvqa [26] to score the retrieved facts. However, this method is not good enough to reason complex facts and requires further development.

Visual Dialog: Extending the single turn dialog task (VQA) to a multi-turn one, we introduce visual dialog. Before the dataset proposed by [4], [5] used a dataset that located the object in an image in which one person gave a question about the image and the other provided a ‘Yes/No/NA’ answer based on the truth. The current dataset expands the range of questions and answers. Question types about image can be various including color, number, relationships, etc., while the answer could be simple as a ‘No’ or a complex description of the image. Three encoding methods were provided in [4] as baselines, namely late fusion, the hierarchical recurrent encoder, and the memory network, as well as two decoding methods: LSTM and softmax. Inspired by the generative adversarial networks (GANs) [8] and the performance gap between discriminative and generative networks, [18] transferred knowledge from a pre-trained discriminative network to a generative network with a Gumbel-softmax [13] LSTM encoder using perceptual loss. [29] combined GAN and reinforcement learning [31] to train the generator with a co-attention encoder, thereby allowing the discriminator to directly access the generated response to evaluate its quality, and used Monte Carlo (MC) searching with a roll-out policy to compute the intermediate reward for each word. [12] merged the fusion step and the scoring step into a single step, which is similar to our synergistic stage, but involved all the image information containing noise and the whole history containing topics irrelevant to current question. Furthermore, it simply arrayed the isolated vectors of the image, question, history, and answer for their representation.

3. Synergistic Network

In a candidate answer set, some of the answers are hard samples that are close or equal to the correct answer, while others are easy samples. Our framework shown in Figure 3 has two stages, the primary stage and the synergistic stage. In the primary stage, we learn representative vectors of the image, history, and question using a co-attention module and then calculates the score of each candidate answer to

separate hard answers from easy ones. In the synergistic stage, we select hard answers together with their questions to form question-answer pairs. These pairs further coordinate with the image and history to predict scores.

We first formally define the visual dialog problem and then introduce our new synergistic strategy. Given an image I and caption C , we collect historical questions and their corresponding answers as H . At turn t , our model gives a score for each answer $a_{t,i}$ in candidate set A_t based on question q_t . To describe input image I , we detect objects and their features using the Faster-RCNN model [1] and apply CNN to encode them into $V = (v_1, \dots, v_n)$, where $v_i \in R^d$ and n is the number of objects. Question q_t is a sequence of words, which can be encoded using LSTM, i.e., $m_t^q = \text{LSTM}(q_t)$. We also organize the previous dialog (including caption) as history $H = (H_0, \dots, H_{t-1})$, where $H_0 = C$, and $H_i = (q_i, a_{i,gt})$ for $i \in \{1, \dots, t-1\}$, which is the concatenation of the question and correct answer at each turn before time t . Similar to the question, we use another LSTM to extract the history features as $U = (u_0, \dots, u_{t-1})$, where $u_i = \text{LSTM}(H_i)$. $m_i^q \in R^d$ and $u_i \in R^d$ correspond to the last state of LSTM.

3.1. Primary Stage

An encoder-decoder solution is adopted in the primary stage [4, 29, 18]. The encoder contains two main tasks, one is how to de-reference in the multi-turn conversations (98% of dialogs contain at least one pronoun), and the other is to locate the objects in the image mentioned in the current question. The attention mechanism [18] is commonly used to tackle the tasks. Instead of linear concatenation, we use multi-modal factorized bilinear pooling (MFB) [32], as it can overcome the difference between distributions of the two features (two LSTMs to encode question and history, respectively; LSTMs for text feature and CNNs for image feature). MFB is expected to provide a richer representation than other bilinear methods, such as MLB [15] and MCB [7]. In MFB, the fusion of two features $X, Y \in R^d$ is calculated by:

$$z = \text{MFB}(X, Y) = \sum_{i=1}^k (\mathbf{U}_i^\top X \circ \mathbf{V}_i^\top Y), \quad (1)$$

where $\mathbf{U}, \mathbf{V} \in R^{d \times l \times k}$ are the parameters to be learned, k is the number of factors, l is the hidden size, and \circ is the Hadamard product (element-wise multiplication). However, Y sometimes represents multiple channel input, e.g., detected objects or history in our model, so the formula becomes:

$$z = \text{MFB}(X, Y) = \sum_{i=1}^k ((\mathbf{U}_i^\top X \cdot \mathbf{1}^\top) \circ (\mathbf{V}_i^\top Y)), \quad (2)$$

where $\mathbf{1} \in R^\phi$ is the vector with all elements equal to one, and ϕ is the channel number of Y . To stabilize the output neurons, we use power normalization ($z \leftarrow \text{sign}(z)|z|^{0.5}$) and ℓ_2 normalization ($z \leftarrow z/\|z\|$).

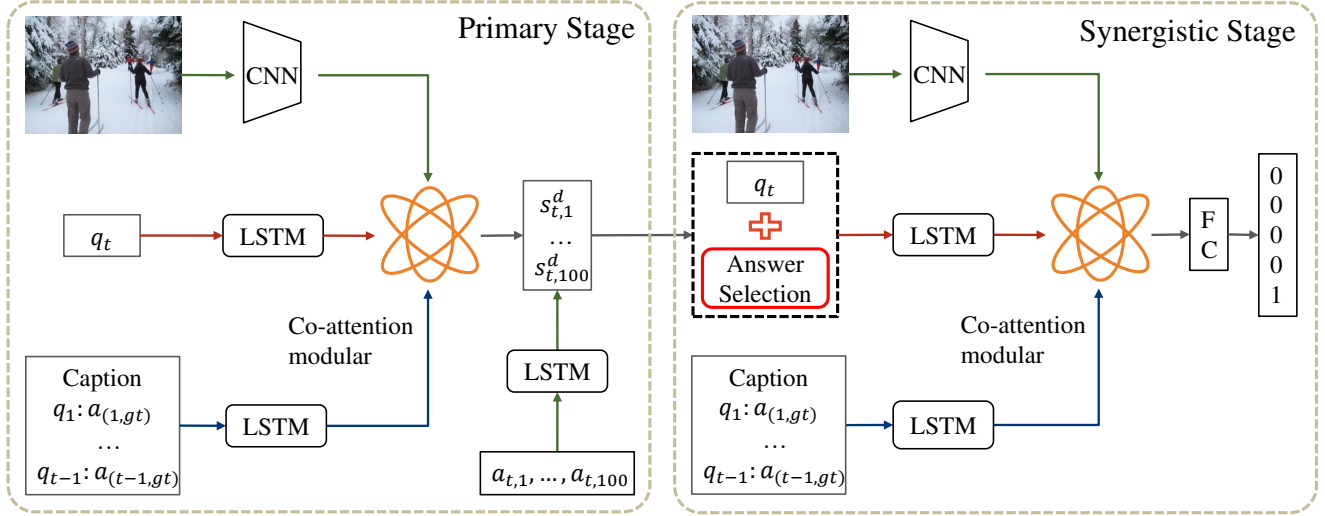


Figure 3. Architecture of our model. All candidate answers are scored in the primary stage, and some selected answers are re-scored in the synergistic stage.

We utilize MFB to learn the unified vector of the question and history, denoted as $z_t^h = \text{MFB}_h(m_t^q, U)$, where $z_t^h \in R^{l \times t}$. Then, we learn the attention weight and vector by:

$$\alpha_t^h = \text{softmax}(w_\alpha^\top z_t^h), \quad (3)$$

$$m_t^h = \sum_{i=0}^{t-1} (\alpha_{t,i}^h u_i), \quad (4)$$

$w_\alpha^\top \in R^l$ is a learned parameter, and $\alpha_t^h \in R^t$ is the calculated weight implying which history the question should refer to. The attended history vector m_t^h can be concatenated with the question vector and then fused with image features, $z_t^v = \text{MFB}_v([m_t^q : m_t^h], V)$. The image attention vector m_t^v can be obtained in a similar approach as Eq.(3) and Eq.(4) by taking z_t^v as the input. Finally, we learn the representation of text and visual features with $e_t^p = \text{MFB}_e([m_t^q : m_t^h], m_t^v)$.

The decoder encodes each candidate answer $a_{t,i} \in A_t$ to $m_{t,i}^a$ using LSTM and calculates the dot similarity score by:

$$s_{t,i}^d = e_t^p \top f_d(m_{t,i}^a), \quad (5)$$

where f_d is a one-layer MLP with activation \tanh to project the answer encoding $m_{t,i}^a \in R^d$ to the space of input embedding e_t^p .

The correct answer $a_{t,gt}$ should have a higher score than the others. Thus, we use N-pair loss [23] to measure the error. Most of the 100 candidate answers are easy samples, which are irrelevant to the inputs, and contribute to no useful learning signal in this loss (score difference less than zero in Figure 5). To solve the imbalance problem, we employ temperature τ to reduce the imbalance impact:

$$\mathcal{L}_D = \log\left(\sum_{i=1}^{100} \exp\frac{s_{t,i}^d - s_{t,gt}^d}{\tau}\right), \quad (6)$$

where $\tau \leq 1$. If the candidate answer $a_{t,i}$ is correctly scored lower than that of the ground truth answer $a_{t,gt}$, loss $l_{t,i} = s_{t,i}^d - s_{t,gt}^d$ will be less than 0, and τ can reduce the contribution of answer $a_{t,i}$. For instance, with $\tau = 0.25$ and $l_{t,i} = -1$, about 20 such items make the same loss as the normal N-pair loss. Otherwise, it amplifies the loss of incorrectly scored answers.

3.2. Synergistic Stage

In the primary stage, some answers are improperly scored due to the limitations of the scoring method. In this synergistic stage, answers are coordinated with the question and image for re-rank. However, easy candidate answers are not needed in further analysis, and we want our second stage model to fully focus on hard answers within its modeling capacity. Therefore, we select the answers with higher probability of being correct based on the predicted scores from the primary stage. As seen from the recall of our best method in Table 1, the top ten answers predicted in the primary stage covered nearly 90% of the ground truth, which means that the remaining 90 answers are of lower probability and can easily be discriminated. Based on this phenomenon, we first pick the top N answers from A_t to organize a new candidate set B_t , where $B_t = (b_{t,1}, \dots, b_{t,N}), B_t \subset A_t$.

The selected answers are often ambiguous for describing the whole meaning of sentences (such as ‘No’ and ‘Black and grey’ in Figure 1), so they must work with the corresponding question to make complete sentences. Thus, we append question q_t to each answer $b_{t,j}, j \in \{1, \dots, N\}$ as a question-answer pair and encode it using LSTM to obtain a vector:

$$m_{t,j}^b = \text{LSTM}([q_t : b_{t,j}]). \quad (7)$$

Extra history is required to remedy the reference problem

of the question. Therefore, we use $m_{t,j}^b$ as a question vector combined with attended history m_t^h to learn the image’s attention parameters:

$$z_{t,j}^r = \text{MFB}_a([m_{t,j}^b : m_t^h], V) \quad (8)$$

and the attended image feature $m_{t,j}^r$ using Eq.(3) and Eq.(4) for selected answer $b_{t,j}$. Similar to the primary stage, we obtain the fusion embedding $e_{t,j}^r = \text{MFB}_r([m_{t,j}^b : m_t^h], m_{t,j}^r)$, which represents the answer vector synergized with the image, question, and history, which is directly used to calculate the score by:

$$s_{t,j}^r = f_r(e_{t,j}^r), \quad (9)$$

where f_r can be a one-layer MLP. An answer in candidate set containing more details and better matching the inputs should earn a higher score than ordinary ones.

Here, we reuse the attended history vector m_t^h from the primary stage, since the co-reference problem in the question can be resolved without knowing the answer. For example, as shown in Figure 1, we can refer ‘their’ to the ‘five people’ regardless of the colors of their jackets. The image feature V is shared with the primary stage, because we want the image feature to be represented universally in the two stages. But the attention weights are learned at each stage, since each candidate answer, as well as the question, depicts its own attention map.

We treat this stage as a classification problem, where the correct answer should have the highest probability:

$$p_t^r = \text{softmax}(s_t^r), \quad (10)$$

$$\mathcal{L}_R = \sum_{j=1}^N -y_j \log(p_{t,j}^r), \quad (11)$$

where y_{gt} is equal to 1 and the others are zeros. We note that this formula can be easily extended to soft cross entropy, where y_i is the probability marking this answer as correct if a dense annotation dataset is available in the future.

4. Extension to the Generative Model

Besides discriminative model in the primary stage, the generative model can also be used to score the candidate answers and seamlessly works with the proposed image-question-answer synergistic method. The encoder of the generative model is the same as that of the discriminative model in the primary stage in Section 3.1. Accordingly, we still use e_t^p to represent the common vector of image I , history H , and question q_t at turn t . The decoder interprets e_t^p into answers and calculates the probability of answer $a_{t,i}$ by:

$$\begin{aligned} s_{t,i}^g &= p(w_{t,i}^1, \dots, w_{t,i}^{T'} | q_t, H, I) \\ &= \prod_{j=1}^{T'} p(w_{t,i}^j | (w_{t,i}^1, \dots, w_{t,i}^{j-1}), e_t^p), \end{aligned} \quad (12)$$

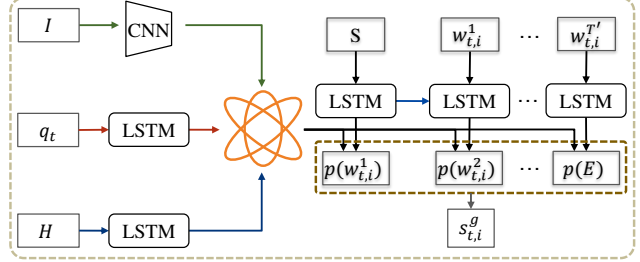


Figure 4. Primary stage of our generative model. The score of each answer is its probability of word sequence.

which is also regarded as score in the primary stage as shown in Figure 4, where $w_{t,i}^1, \dots, w_{t,i}^{T'}$ is the word sequence of the answer $a_{t,i}$, and T' is word number. For each word, its probability is given by:

$$p(w_{t,i}^j | (w_{t,i}^1, \dots, w_{t,i}^{j-1}), e_t^p) = f_g(\text{MFB}_g(h_j, e_t^p)), \quad (13)$$

where h_0 is the last hidden state of LSTM for the question, $h_j = \text{LSTM}(h_{j-1}, w_{t,i}^{j-1})$ is the state of LSTM, and f_g maps the fusion vector to the word space. Instead of initializing the LSTM of decoder with encoded vector e_t^p [4], we regard it as a context vector [19]. This is for three main reasons: first, the encoded vector and the LSTM decoder have different distributions, and all gates and the hidden state in LSTM are learned by linear combination; second, the salient objects that should be attended by each token are already chosen by the question, and the aim of the decoder is to translate the visual information of salient objects into text, such that the context vector is fixed for each token; and third, instead of learning a joint vector with h_{j-1} , the context vector could be considered as compensatory information to the current hidden state h_j , which reduces the uncertainty for next word prediction.

To make the correct answer $a_{t,gt}$ score higher in the primary stage, we maximize the conditional probability $p(w_{t,gt}^1, \dots, w_{t,gt}^{T'} | q_t, H, I)$. Thus, the loss function is the sum of the negative log likelihood of the correct word in each step:

$$\mathcal{L}_G = - \sum_{i=1}^{T'} \log p(a_{t,gt}). \quad (14)$$

If the candidate answers are given in the generative model as in the discriminative model, we can collect the score for each answer using Eq.(12). Otherwise, we can generate some candidate answers with high probability by beam search [24]. After collecting the score for each answer in the primary stage, we follow the strategy in Section 3.2 to pick the top N answers and then synergize them with the image, question, and history to learn a better representation for re-scoring.

5. Experiments

In this section, we evaluate our synergistic strategy on a visual dialog dataset. We introduce the dataset and evalua-

tion metric, before describing our experimental setting and results, and finally the qualitative analysis.

5.1. Dataset and Evaluation Metric

Our model is trained on the Visual Dialog v1.0 dataset [4], which contains about 120k images from COCO-trainval [17]. Each image has one caption and 10 turn dialogs, i.e., about 1.2 million question-answer pairs. To organize this dataset, two people chatted on the Amazon Mechanical Turk. The questioner could not see the image and asked a question based on the given caption and previous context to better understand the scene, while the answerer could see both the image and caption and replied to the question as naturally and conversationally as possible. Each question has 100 candidate answers containing one correct answer, 50 answers to similar questions, 30 popular answers, and some randomly picked answers from the dataset. For the validation and test datasets, 10K COCO-like images were collected from Flickr. The test dataset was densely annotated in v1.0 by four people to allow application for the more robust evaluation metric NDCG [27] rather than traditional retrieval metrics, such as mean rank, R@1,5,10, and MRR. There could be more than one correct answer to each question in the candidate set, such as ‘yeah’ and ‘yes’. In this situation, NDCG is invariant to the order of options with identical relevance. For each candidate answer, its relevance is $\frac{\text{number of annotators who marked answer as relevant}}{4}$. The metric is given by:

$$\text{DCG@k} = \sum_{i=1}^k \frac{\text{relevance}_i}{\log_2(i+1)}, \quad (15)$$

$$\text{NDCG@k} = \frac{\text{DCG@k for submitted ranking}}{\text{DCG@k for ideal ranking}}, \quad (16)$$

where k is the number of answer options whose relevance is greater than zero. Of these metrics, a higher score is better for NDCG, MRR, and R@1,5,10, but a lower score is better for mean rank.

5.2. Implementation Details

We first constructed the vocabulary, which contains the words appearing in the questions, correct answers, and captions more than four times in the train dataset. This made 11,213 words with the padding word ‘PAD’, out-of-vocabulary word ‘UNK’, start symbol ‘START’, and end symbol ‘END’. Then, each word was embedded within a 300-dimension vector shared across caption, history, question, and answer. The maximum lengths of the caption, question, answer, and history were 40, 20, 20, and 40, respectively. For each candidate answer, we inserted ‘START’ at the head and appended ‘END’ at the tail. The LSTMs of the question and history are two layered, while it is one layered for the answer in the primary stage and the question-answer pair in the synergistic stage. The hid-

Model	NDCG	MRR	R@1	R@5	R@10	Mean
LF[4]	45.31	55.42	40.95	72.45	82.83	5.95
HRE[4]	45.46	54.16	39.93	70.45	81.50	6.41
MN[4]	47.50	55.49	40.98	72.30	83.30	5.92
MN-att[4]	49.58	56.90	42.43	74.00	84.35	5.59
LF-att[4]	49.76	57.07	42.08	74.83	85.05	5.41
Technion	54.46	67.25	53.40	85.28	92.70	3.55
MS AI	55.35	63.27	49.53	80.40	89.60	4.15
USTC-YTH	56.47	61.44	47.65	78.13	87.88	4.65
Single (ours)	57.32	62.20	47.90	80.43	89.95	4.17
Ensemble (ours)	57.88	63.42	49.30	80.77	90.68	3.97

Table 1. Performance of discriminative models on test-standard server of Visual Dialog Challenge 2018. We ensemble 10 models with different seeds and varying M from 25 to 35.

den state dimension d for all LSTMs and CNN is 512. For bilinear pooling, we set k to 5 and l to 1000 as Yu [32].

We start training the primary stage with loss $\mathcal{L}_D(\mathcal{L}_G)$ for 7 epochs to arrange the top ranked answers relative to the inputs. The synergistic stage follows with loss $\mathcal{L}_D(\mathcal{L}_G) + \mathcal{L}_R$ for another 15 epochs. When training the synergistic stage, our policy is to randomly sample $N - 1$ answers from the top M ranked in the first stage combined with the right answer $a_{t,gt}$ to organize B_t . During testing, we only select the top N answers. For this dataset, we choose $N = 10$ or 15 and M varying from 10 to 40 for the discriminative model, then $N = 10, 20, 30$ and M fixed to 30 for the generative model, and we analyze the effect of different N and M on performance in Section 5.4. Our model is trained using the Adam solver [16] with $\beta_1 = 0.9$, $\beta_2 = 0.99$, initial learning rate 10^{-3} , and decay every 7 epochs with an exponential rate of 0.25.

5.3. Comparison with the State-of-the-art

We compare our discriminative model with baselines [4] and other methods: Later Fusion (LF), which encodes the question, image, and history respectively and projects their concatenation into a joint embedding; Hierarchical Recurrent Encoder (HRE), which uses a hierarchical architecture to encode the dialog history; Memory Network (MN), which maintains a memory bank storing previous dialog, which is attended by a question; and MN-att and LF-att, which add an attention mechanism for image to their base methods. From Table 1, it can be seen that our best single model improves NDCG by 7.56% and MRR by 5.13% compared with LF-att. To improve accuracy, we ensemble 10 models with different seeds and M . We rank the answers by summing scores of all models and achieve the highest NDCG on the test-standard server of Visual Dialog Challenge 2018.

5.4. Ablation Study

We conduct several ablation studies to verify the contribution of each module to our discriminative model. The first three lines in Table 2 show the performance of the primary stage with different τ varying from 1 to 0.25. The

N	M	τ	MRR	R@1	R@5	R@10	Mean
-	-	1.00	61.92	47.53	79.78	89.28	4.42
-	-	0.50	62.51	48.30	80.05	89.57	4.27
-	-	0.25	62.63	48.31	80.39	89.65	4.23
10	10	0.25	62.31	48.18	79.45	90.21	4.22
10	20	0.25	62.83	48.45	80.70	90.28	4.11
10	30	0.25	63.54	49.21	81.01	90.32	4.09
10	40	0.25	63.14	48.77	80.97	90.02	4.13
15	30	0.25	63.16	48.91	80.75	90.22	4.12

Table 2. Performance of our discriminative model on validation dataset.

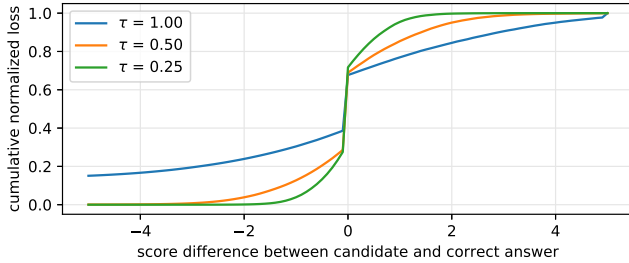


Figure 5. Cumulative normalized loss for different τ .

performance improves as τ decays, because most candidate answers are easily discriminated. The summed loss of these easy negative answers whose scores are lower than the correct answer by one consumes almost 30% of the model’s energy at $\tau = 1.00$, as shown in Figure 5. This kind of loss is reduced as τ decreases and is nearly zero when $\tau = 0.25$, saturating our primary stage. This makes the model pay more attention to incorrect answers scored near or higher than the correct answer. The last five lines show the results of the synergistic stage with different settings. MRR of this stage drops at $N = 10$ and $M = 10$, as the top answers of the primary stage become stable after several epochs, resulting in the synergistic stage learning bias. Feeding more samples by increasing M improves performance to produce the best model at $M = 30$, showing that synergy can learn a better representation of the image, question, and answer. What of interest is that the performance drops at $M = 40$ or $N = 15$, possibly because the selected answers become less relevant to the inputs, and the second stage model’s ability is to score nearly correct answers but is sensitive to unrelated ones. Conversely, it also addresses the importance of the primary stage to the synergistic stage. And the primary stage is also necessary to balance the memory cost, since each answer learns its own attention map and fusion with the image in the synergistic stage, while in the primary stage, only one attention map is required for the question and image.

For the generative model, the first line in Table 3 shows the performance of the memory network initialized by the common vector of inputs, and the second line shows the result of our model using only the primary stage, leaving other three lines for the synergistic stage. It can be seen

Model	N	M	MRR	R@1	R@5	R@10	Mean
MN-att [4]	-	-	47.94	37.48	58.56	65.57	17.61
Primary (ours)	-	-	49.01	38.54	59.82	66.94	16.69
Synergistic (ours)	10	30	51.62	40.77	63.58	67.00	16.51
Synergistic (ours)	20	30	53.23	41.42	67.22	72.91	15.87
Synergistic (ours)	30	30	53.73	41.28	69.01	75.85	15.12

Table 3. Performance of our generative model on validation dataset with candidate set.

that the primary stage outperforms baseline by 1.1% with respect to MRR, which generates a strong candidate set for the next stage. Our synergistic stage further improves MRR by 2.6% when $N = 10$, since the primary stage focuses on each word but lacks the understanding of whole answer sequence. In contrast to the discriminative model in which about 90% of correct answers are top ten ranked in the primary stage, only 2 out of 3 correct answers are top sorted by the generative model, so we increase the selected answer number N from 10 to 30 in the second stage, which further increases MRR by 2.1%. Furthermore, our model improves R@5 by 9.2% and R@10 by 8.9%.

5.5. Qualitative Analysis

To further demonstrate the effects of our synergistic model, we present some examples from the validation dataset. Figure 6 shows the results of the discriminative model ranked by only the primary stage and our two-stage network. The answer in red is correct, while the others are top-ranked candidate answers. It can be seen that the one-stage model tries to give a safe answer such as ‘No’ to the top-left and middle-left images, while our model depicts more details by adding ‘only the giraffe’ and ‘another pair of feed’. Biased answers are given in middle-right image, since for a binary question with choices, the answer word is always contained in the question. There is also bias in bottom-right image, because only humans tend to wear hats. Our model gives unbiased answers based on the images, and furthermore, can detect the discrepancies between similar words in the bottom-left image.

In order to apply our method to realistic applications, we abandon the prepared candidate set in the first stage and generate another within our primary model by beam search [24], which maintains a partial sequence list of size B . At each step, all partial sequences are extended with the whole vocabulary, and only the top B sequences with the highest probability are retained for next step. Sentences meeting the ‘END’ symbol are moved from the partial list to the complete one. Starting with the ‘START’ symbol and iterating for a maximum 20 steps, we obtain a candidate set of size B with all complete sentences complemented by some partial ones. In Figure 7, we show the top generated answers with $B = 15$ for the primary and synergistic stage, the pre-set answer is below the question. The primary stage always ranks short answers having one or two words higher than the long sequences that depict images with more informa-

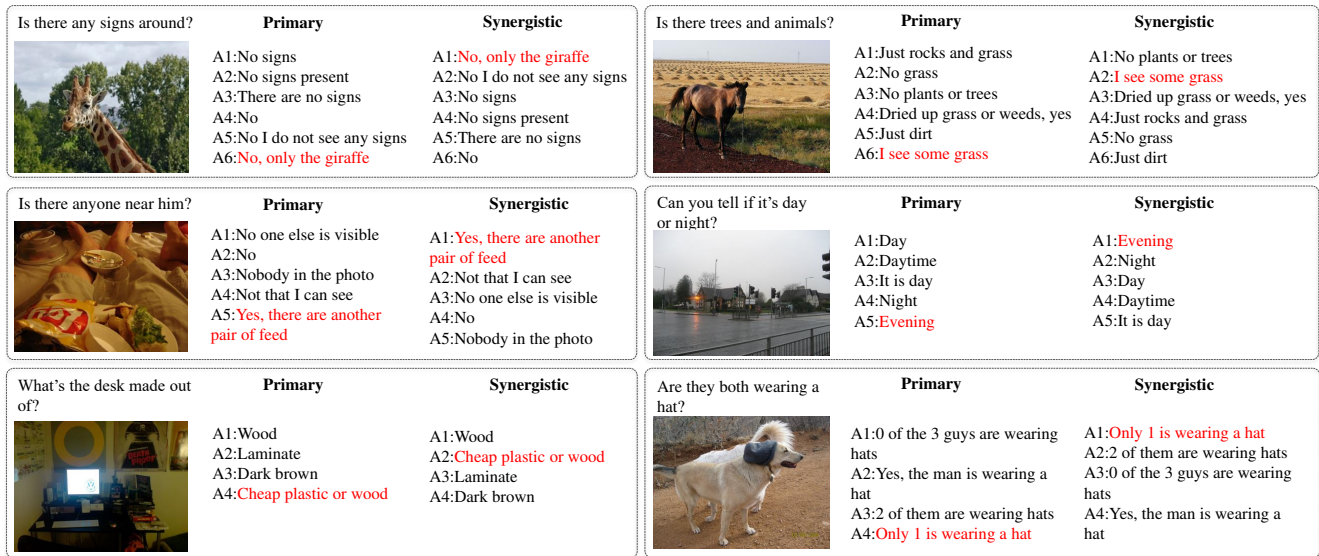


Figure 6. Qualitative comparison for discriminative model.

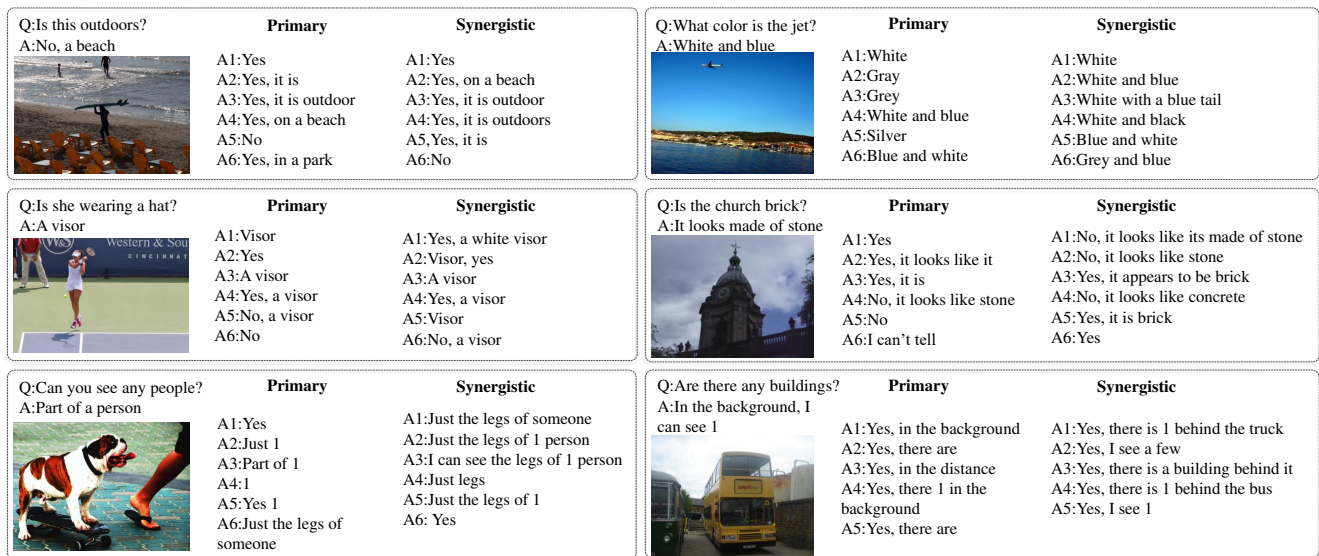


Figure 7. Qualitative comparison for generative model without candidate answers.

tion, because the generative method calculates the score for each answer by product probability of its words. In the synergistic stage, this problem is overcome, since the extra attribute information, e.g., ‘a white visor’ (middle-left image) and ‘white with a blue tail’ (top-right image), can have a higher score than a simple answer. Surprisingly, our model can sometimes even generate better answers than those provided, such as ‘Just the legs of someone’ vs. ‘Part of a person’ in the bottom-left image and ‘behind the truck’ vs. ‘In the background’ in the bottom-right image.

6. Conclusions

The limitations of previous input-answer fusion methods mean that they cannot correctly represent the common vector of these features. As a result, they omit detailed in-

formation and focus on short and safe answers. In this paper, we develop a synergistic network that jointly learns the representation of the image, question, answer, and history in a single step. We also improve the N-pair loss function to solve the class-imbalanced problem in the discriminative model. Our final proposed discriminative model achieves state-of-the-art performance on the Visual Dialog v1.0 test-standard server with the robust NDCG evaluation metric. The results of our generative model are also encouraging.

7. Acknowledgement

This work was supported in part by the Australian Research Council Projects FL-170100117, DP-180103424, IH-180100002, and DE-180101438.

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, volume 3, page 6, 2018. 3
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. 1, 2
- [3] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A nucleus for a web of open data. In *The semantic web*, pages 722–735. Springer, 2007. 3
- [4] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, 2017. 2, 3, 5, 6, 7
- [5] Harm De Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron C Courville. Guess-what?! visual object discovery through multi-modal dialogue. In *CVPR*, volume 1, page 3, 2017. 3
- [6] Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems*, pages 6594–6604, 2017. 2
- [7] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*, 2016. 2, 3
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 3
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [11] Allan Jabri, Armand Joulin, and Laurens van der Maaten. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer, 2016. 2
- [12] Unnat Jain, Svetlana Lazebnik, and Alexander G Schwing. Two can play this game: Visual dialog with discriminative question generation and answering. In *Proc. CVPR*, volume 1, 2018. 2, 3
- [13] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. 3
- [14] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. *arXiv preprint arXiv:1805.07932*, 2018. 3
- [15] Jin-Hwa Kim, Kyoung-Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Hadamard product for low-rank bilinear pooling. *arXiv preprint arXiv:1610.04325*, 2016. 3
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [18] Jiasen Lu, Anitha Kannan, Jianwei Yang, Devi Parikh, and Dhruv Batra. Best of both worlds: Transferring knowledge from discriminative learning to a generative visual dialog model. In *Advances in Neural Information Processing Systems*, pages 314–324, 2017. 2, 3
- [19] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. 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 (CVPR)*, volume 6, page 2, 2017. 5
- [20] Hyeonseob Nam, Jung-Woo Ha, and Jeonghee Kim. Dual attention networks for multimodal reasoning and matching. *arXiv preprint arXiv:1611.00471*, 2016. 2
- [21] Medhini Narasimhan and Alexander G Schwing. Straight to the facts: Learning knowledge base retrieval for factual visual question answering. *arXiv preprint arXiv:1809.01124*, 2018. 3
- [22] Ethan Perez, Harm De Vries, Florian Strub, Vincent Dumoulin, and Aaron Courville. Learning visual reasoning without strong priors. *arXiv preprint arXiv:1707.03017*, 2017. 2
- [23] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, pages 1857–1865, 2016. 4
- [24] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014. 1, 5, 7
- [25] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015. 1
- [26] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Fvqa: Fact-based visual question answering. *IEEE transactions on pattern analysis and machine intelligence*, 2017. 3
- [27] Yining Wang, Liwei Wang, Yuanzhi Li, Di He, Wei Chen, and Tie-Yan Liu. A theoretical analysis of ndcg ranking measures. In *Proceedings of the 26th Annual Conference on Learning Theory (COLT 2013)*, volume 8, 2013. 6
- [28] Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Ask me anything: Free-form visual question answering based on knowledge from external

- sources. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4622–4630, 2016. 3
- [29] Qi Wu, Peng Wang, Chunhua Shen, Ian Reid, and Anton van den Hengel. Are you talking to me? reasoned visual dialog generation through adversarial learning. *arXiv preprint arXiv:1711.07613*, 2017. 1, 2, 3
- [30] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 21–29, 2016. 2
- [31] Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*, pages 2852–2858, 2017. 3
- [32] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. *IEEE International Conference on Computer Vision (ICCV)*, pages 1839–1848, 2017. 3, 6
- [33] Zhou Yu, Jun Yu, Chenchao Xiang, Zhou Zhao, Qi Tian, and Dacheng Tao. Rethinking diversified and discriminative proposal generation for visual grounding. *arXiv preprint arXiv:1805.03508*, 2018. 1