

Structured Triplet Learning with POS-tag Guided Attention for Visual Question Answering

Zhe Wang¹ Xiaoyi Liu² Liangjian Chen¹ Limin Wang³ Yu Qiao⁴ Xiaohui Xie¹ Charless Fowlkes¹

¹Dept. of CS, UC Irvine

²Microsoft

³CVL, ETH Zurich

⁴SIAT, CAS

Abstract

Visual question answering (VQA) is of significant interest due to its potential to be a strong test of image understanding systems and to probe the connection between language and vision. Despite much recent progress, general VQA is far from a solved problem. In this paper, we focus on the VQA multiple-choice task, and provide some good practices for designing an effective VQA model that can capture language-vision interactions and perform joint reasoning. We explore mechanisms of incorporating part-of-speech (POS) tag guided attention, convolutional n-grams, triplet attention interactions between the image, question and candidate answer, and structured learning for triplets based on image-question pairs¹. We evaluate our models on two popular datasets: Visual7W and VQA Real Multiple Choice. Our final model achieves the state-of-the-art performance of 68.2% on Visual7W, and a very competitive performance of 69.6% on the test-standard split of VQA Real Multiple Choice.

1. Introduction

The rapid development of deep learning approaches [9, 40, 39, 41] has resulted in great success in the areas of computer vision [11, 32, 31, 28, 29, 30] and natural language processing [24, 18]. Recently, Visual Question Answer (VQA) [3, 10] has attracted increasing attention, since it evaluates the capacity of vision systems for a deeper semantic image understanding, and is inspiring the development of techniques for bridging computer vision and natural language processing to allow joint reasoning [3, 42]. VQA also forms the basis for practical applications such as tools for education, assistance for the visually-impaired, or support for intelligence analysts to actively elicit visual information through language-based interfaces [4].

Given an image, a typical VQA task is to either gen-



Q: Why was the hand of the woman over the left shoulder of the man?

A: They were together and engaging in affection.

A: The woman was trying to get the man's attention.

A: The woman was trying to scare the man.

A: The woman was holding on to the man for balance.

Figure 1. An example of the multiple-choice VQA. The red answer is the ground-truth, and the green ones are human-generated wrong answers.

erate an answer as a free-form response to an open-ended question, or pick from a list of candidate answers (multiple-choice) [3, 42]. Similar to [13] and [27], we mainly focus on the multiple-choice task in this paper, an example of which is shown in Fig. 1. A good VQA system should be able to interpret the question semantically, extract the key information (i.e., objects, scenes, actions, etc.) presented in the given image, and then select a reasonable answer after jointly reasoning over the language-visual interactions.

In this paper, we propose a simple but effective VQA model that performs surprisingly well on two popular datasets: Visual7W Telling and VQA Real Multiple-Choice. We start with the architecture in [27], which combines word features from the question and answer sentences as well as hierarchical CNN features from the input image. Our insights on “good practice” are fourfold: (i) To precisely capture the semantics in questions and answers, we propose to exploit a part-of-speech (POS) tag guided attention model to ignore less meaningful words (e.g., coordinating conjunctions such as “for”, “and”, “or”) and put more emphasis on the important words such as nouns, verbs and adjectives. (ii) We leverage a convolutional n-gram model [16] to capture local context needed for phrase-level or even sentence-level meaning in questions and answers. (iii) To integrate the vision component (represented by the CNN

¹Code: <https://github.com/wangzheallen/STL-VQA>
Contact: buptwangzhe2012@gmail.com

features extracted from the pre-trained deep residual network (ResNet)), we introduce a triplet attention mechanism based on the affinity matrix constructed by the dot product of vector representations of each word in the question or answer and each sub-region in the image, which measures the matching quality between them. After appropriate pooling and normalization, we linearly combine the attention coefficients from questions and answers to weight relevant visual features. (iv) To encourage the learning to be more discriminative, we mine hard negative samples by sending answers corresponding to the same question to the network to learn simultaneously. By setting the margin between the correct answer and hard negative answer, we observe performance increasing.

Our proposed methods achieve state-of-the-art performance of 68.2% on Visual7W benchmark, and an competitive performance of 69.6% on test-standard split of VQA Real Multiple Choice. Our approach offers simple insights for effectively building high performance VQA systems.

2. Related Work

Models in VQA: Existing VQA solutions vary from symbolic approaches [21, 33], neural-based approaches [20, 25, 38, 23], memory-approaches [34, 14], to attention-based approaches [7, 19, 26, 17, 35]. In addition to these models, some efforts have been spent on better understanding the behavior of existing VQA models [1]. Jabri et al. [13] proposed a simple model that takes answers as input and performs binary predictions. Their model simply averages word vectors to get sentence representations, but competes well with other more complex VQA systems (e.g. LSTM). Our work proposes another language representation (i.e., Convolutional n-grams) and achieves better performances on both the Visual7W dataset [42] and the VQA dataset [3]. **Attention in VQA:** A number of recent works have explored image attention models for VQA [7, 19, 36]. Zhu et al. [42] added spatial attention to the standard LSTM model for pointing and grounded QA. Andreas et al. [2] proposed a compositional scheme that exploited a language parser to predict which neural module network should be instantiated to answer the question. Fukui et al. [7] applied multi-modal bilinear to attend images using questions. Gan [8] link the COCO segmentation and caption task, and add segmentation-like attention to their VQA system. Unlike these works, we propose a POS tag guided attention, and utilize both question and answer to drive visual attention.

Learning in VQA model: VQA models use either a softmax [3] or a binary loss [13]. Softmax-loss-based models formulate VQA as a classification problem and construct a large fixed answer database. Binary-loss-based models formulate multi-choice VQA as a binary classification problem by sending image, question and candidate answer as a triplet into network. They are all based on independently

classifying (*image, question, answer*) triplets. Unlike our model, none incorporate the other triplets corresponding to the same (*image, question*) pair during training.

POS Tag Usage in Vision and Language: Our POS tag guided attention is different from [37]. They use POS tag to guide the parser to parse the whole sentence and model LSTM in a hierarchical manner and apply it to image caption task. Instead, We directly parse the POS tag of each word and utilize the POS tag as a learning mask on the word glove vector. We divide POS tag into more fine-grained 7 categories as Sec. 3.2.1. There are also other papers using POS tag in different ways. [6] calculates average precision for different POS tag. [12] discovers POS tags are very effective cues for guiding the Long Short-Term Memory (LSTM) based word generator.

3. Model Architecture

3.1. Architecture Overview

For the multiple-choice VQA, each provided training sample consists of one image I , one question Q and N candidate answers A_1, \dots, A_N , where A_1 is the ground-truth answer. We formulate this as a binary classification task by outputting a target prediction for each candidate triple $\{\text{“image”}: I, \text{“question”}: Q, \text{“candidate answer”}: A_i, \text{“target”}: t_i\}$, where, for example, $t_1 = 1$ and $t_i = 0$ for $i = 2, \dots, N$.

Different from previous work [38, 13], we adopt the architecture in Fig. 2, which combines features from question, answer and image in a hierarchical manner. For a fair comparison between our model and [13], we have implemented both and obtained the same performance of 64.8% as [13] on Visual7W, but we find that the hierarchical model training converges much faster (in less than 20 epochs).

Denote the final vector representations of the question, image and the i -th candidate answer by $\mathbf{x}_Q, \mathbf{x}_I$ and \mathbf{x}_{A_i} . In the first stage, \mathbf{x}_Q and \mathbf{x}_I are combined via the Hadamard product (element-wise multiplication) to obtain a joint representation of question and image:

$$\mathbf{x}_{QI} = \mathbf{x}_Q \odot \mathbf{x}_I.$$

In the second stage, \mathbf{x}_{QI} and \mathbf{x}_{A_i} are fused as:

$$\mathbf{x}_{QIA_i} = \tanh(\mathbf{W}_{QI}\mathbf{x}_{QI} + \mathbf{b}_{QI}) \odot \mathbf{x}_{A_i}.$$

In the last stage, a binary classifier is applied to predict the probability of the i -th candidate answer being correct:

$$p_i = \text{sigmoid}(\mathbf{W}_{QIA}\mathbf{x}_{QIA_i} + \mathbf{b}_{QIA}).$$

We jointly train all the weights, biases and embedding matrices to minimize the cross-entropy or structured loss. For inference on test samples, we calculate the probability of correctness for each candidate answer, and select the one with highest probability as the predicted correct answer: $i^* = \arg \max_{i=1, \dots, N} p_i$.

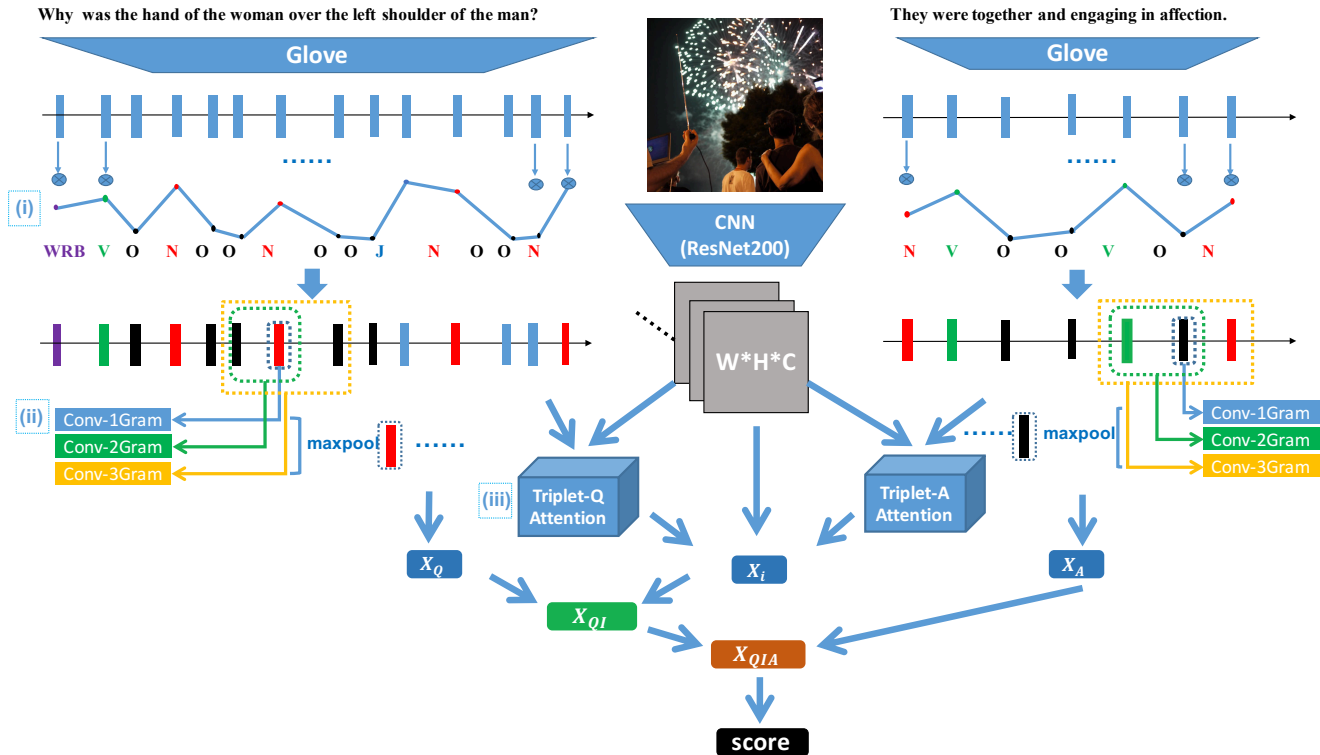


Figure 2. Illustration of our pipeline for VQA. Best viewed in color. We first extract Glove vector representations of each word in the question and answer, which are weighted by a POS tag guided attention for each word. We transform each sentence using a convolutional n-gram to encode contextual information and average to get QUESTION-vec. For visual features, we utilize a standard CNN model and conduct weighted summation to IMAGE-vec by applying triplet attention. Finally we combine QUESTION-vec, IMAGE-vec and ANSWER-vec to score the quality of the proposed answer. Structured learning on triplets is applied after getting the score for answers corresponding for each image question pair.

3.2. Model Details

In this subsection, we will describe how to obtain the final vector representations of \mathbf{x}_Q , \mathbf{x}_I and \mathbf{x}_{A_i} . We use the same mechanism for obtaining \mathbf{x}_Q and \mathbf{x}_{A_i} , and thus focus on the former. We represent each question with words q_1, q_2, \dots, q_M by $\mathbf{X}_Q = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_M]$, where \mathbf{e}_m is the corresponding vector for word q_m in the embedding matrix \mathbf{E} .

3.2.1 POS Tag Guided Attention

For each question, it is expected that some words (i.e., nouns, verbs and adjectives) should matter more than the others (e.g., the conjunctions). Hence, we propose to assign different weight to each word based on its POS tag to impose different attentions.

In practice, we find that it works better to group the original 45 pos tags into a smaller number of categories. Specifically, we consider the following seven categories:

1. CD for cardinal numbers;

2. J (including JJ, JJR and JJS) for adjectives;
3. N (including NN, NNS, NNP and NNPS) for nouns;
4. V (including VB, VBD, VBG, VBN, VBP and VBZ) for verbs;
5. WP (including WP and WP\$) for Wh-pronouns;
6. WRB for Wh-adverb;
7. O for Others.

An example of one question and its POS tag categories is shown in Fig. 2. For the question “Why was the hand of the woman over the left shoulder of the man”, the POS tags are given as “WRB, V, O, N, O, O, N, O, O, J, N, O, O, N”. Each category is assigned with one attention coefficient POS_i , which will be learned during training. In this way, each word is represented by $\hat{\mathbf{e}}_i = \mathbf{e}_i \times POS_i$, and the question is represented by $[\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \dots, \hat{\mathbf{e}}_M]$.

3.2.2 Convolutional N-Gram

We propose using a convolutional n-gram to combine contextual information over multiple words represented as vectors. Specifically, we utilize multiple window sizes for one-dimensional convolutional neural network. For a window of size L , we apply the corresponding filter F_L for each word \hat{e}_i , obtaining $F_L(\hat{e}_{i-(L-1)/2}, \dots, \hat{e}_i, \dots, \hat{e}_{i+(L-1)/2})$, when L is odd; and $F_L(\hat{e}_{i-L/2}, \dots, \hat{e}_i, \dots, \hat{e}_{i+L/2-1})$ when L is even. Therefore, we not only consider the i -th word, but also the context within the window. Practically, we apply the filters with window sizes from 1 to L , and then max-pool all of them along each word to obtain a new representation

$$\tilde{e}_i = \text{maxpool}(F_L, F_{L-1}, \dots, F_1).$$

3.2.3 Final Sentence Representation

An efficient and effective way to compute a sentence embedding is to average the embeddings of its constituent words [15]. Therefore, we let the final question vector representation to be $\mathbf{x}_Q = \frac{1}{M} \sum_{i=1}^M \tilde{e}_i$.

Remark 1: In our experiments, we have found that the simple average sentence embedding shows much better performance than the widely used RNNs or LSTMs in both Visual7W and VQA datasets. The reason might be that the questions and candidate answers in the majority of the two VQA datasets tend to be short and have simple dependency structure. RNNs or LSTMs should be more successful when the questions or candidate answers in VQA are more complicated, require deeper reasoning and with more data (e.g., Visual Gnome). Here we only compare models directly trained from Visual7W or VQA.

3.2.4 Triplet Attention

An effective attention mechanism should closely reflect the interactions between the question/answer and the image. Our proposed triplet attention mechanism is given as

$$\mathbf{att}_I = \text{norm}(\lambda_1 \times \mathbf{att}_{Q-I} + \mathbf{att}_{A_i-I}),$$

where $\text{norm}(\mathbf{x}) = \frac{\mathbf{x}}{\sum(\mathbf{x})}$, \mathbf{att}_{Q-I} and \mathbf{att}_{A_i-I} are the attention weights from the question or candidate answer to the image, respectively. λ_1 is a learned coefficient to balance the influences imposed from the question and candidate answer on the image features.

For a given image, the raw CNN features for each sub-region $\mathbf{X}_{I,\text{raw}} = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K]$ are transformed as $\mathbf{X}_I = \text{relu}(\mathbf{W}_I \mathbf{X}_{I,\text{raw}} + \mathbf{b}_I)$. With the previously obtained representation $\mathbf{X}_Q = [\tilde{e}_1, \dots, \tilde{e}_M]$ for the associated question, an affinity matrix is obtained as

$$\mathbf{A} = \text{softmax}(\mathbf{X}_Q^T \times \mathbf{X}_I).$$

The i -th column of \mathbf{A} reflects the closeness of matching between each word in the question and the i -th sub-region. Via max-pooling, we can find the word that matches most closely to the i -th sub-region. Thus, the attention weights from the question to the image are obtained as

$$\mathbf{att}_{Q-I} = \text{maxpool}(\mathbf{A}).$$

Similarly, \mathbf{att}_{A_i-I} is obtained as the attention weights from the candidate answer to the image. Then, the final vector representation for the image is

$$\mathbf{x}_I = \mathbf{X}_I \times \mathbf{att}_I^T.$$

Remark 2: In our experiments, we find that the combination of using `relu` only for the image features and `tanh` for others is most effective.

3.2.5 Structured Learning for Triplets

For the multiple-choice VQA, each provided training sample consists of one image I , one question Q and N candidate answers A_1, \dots, A_N , where A_1 is the ground-truth answer. We formulate this as a binary classification task by outputting a target prediction for each candidate triple $\{I, Q, A_i, t_i\}$, where, e.g., $t_1 = 1$ and $t_i = 0$ for $i = 2, \dots, N$. The output of our model for the i -th candidate answer (as discussed above) is given by:

$$p_i = \text{sigmoid}(\mathbf{W}_{QIA} \mathbf{x}_{QIA_i} + \mathbf{b}_{QIA}).$$

The standard binary classification loss seeks to minimize:

$$L_b = - \sum_{i=1}^N t_i \log p_i$$

To improve the model's discriminative ability, we introduce a structured learning loss that imposes a margin between the correct answer and any incorrect answer. We simultaneously compute scores for all candidate answers corresponding to the same image, question pair. We encourage the distance between target positive answer score and the hardest negative answer (highest scoring negative). The structured learning loss is given by:

$$L_s = \max_i(\text{margin} + p_i - p_1, 0)$$

where the *margin* is the large margin scalar we set to encourage the ability of network to distinguish right triplet and wrong triplet. Thus the final loss to minimize is:

$$L = L_b + \lambda_2 L_s,$$

where λ_2 is applied here to balance these two loss functions.

4. Experiments

In the following, we first introduce the datasets for evaluation and the implementation details of our approach. Then, we explore the proposed good practices for handling visual question answering problem and do ablation study to verify the effectiveness for each step. After this, we make comparison with the state-of-the-art. Finally, we visualize the learned POS tag guided attention for word embedding and triplet attention for images, and compare the attention when wrong answers appear and right answers appear.

4.1. Dataset and Evaluation protocol

Visual7W [42]: The dataset includes 69,817 training questions, 28,020 validation questions, 42,031 test questions, 14,366 training images 5,678 validating images and 8,609 testing images. There are a total of 139,868 QA pairs. Each question has **4** candidate answers. The performance is measured by the percentage of correctly answered questions.

VQA Real Multiple Choice [3]: The dataset includes 248,349 questions for training, 121,512 for validation, and 244,302 for testing. Each question has **18** candidate answers. We follow the VQA evaluation protocol in [3]. In the following, we mainly report our results on Visual7W, since the results on VQA are similar.

4.2. Experiment Setup

We use Tensorflow [5] to develop our model. The Adam optimizer is adopted with a base learning rate of 0.0002 for the language embedding matrix and 0.0001 for others. The momentum is set to be 0.99, and the batch size is 18 for Visual7W and 576 for VQA, since the latter is a much larger and more imbalanced dataset. Our model is trained within 20 epochs with early stopping if the validation accuracy has not improved in the last 5 epochs. Word embedding is extracted using Glove [24], and the dimension is set as 300.² The size of the hidden layer QI and QIA are both 4,096. Dimension of image embedding is 2,048. `relu` is used for image embedding, while `tanh` is used for others. Batch normalization is applied just before last full connected layer for classification, and no dropout is used.

In our experiments, when attention is not applied, we resize each image to be 256×256 , crop the center area of 224×224 , and take the activation after last pooling layer of the pre-trained ResNet 200-layer model [11] as the extracted CNN feature (size 2048). When the proposed triplet attention is applied (ablation study), we rescale the image to be 224×224 , and take the activation from the last pooling layer as the extracted features of its sub-regions (size $7 \times 7 \times 2048$). For the full model we rescale the image

²In our experiments, Glove achieves better performance than Word2Vec.

to be 448×448 , and take the activation from the last pooling layer as the extracted features of its sub-regions (size $14 \times 14 \times 2048$).

4.3. Evaluation on Good Practices

Exploration of Network Training: As most datasets for VQA are relatively small and rather imbalanced, training deep networks is challenging due to the risk of learning biased-data and over-fitting. To mitigate these problems, we devise several strategies for training as follows:

Handling Imbalanced Data. As shown in Section 3.1, we reformulate the learning as a binary classification problem. The positive and negative examples for each question are quite unbalanced, i.e., 1:3 for Visual7W and 1:17 for VQA. Thus during training, for each epoch, we sample a certain number of negative examples corresponding each question. Figure 3(a) illustrates our exploration on Visual7W and it suggests two negative samples for each positive one.

Appropriate Training Batch Size. Handling data imbalance results in changes that require further adjustment of optimization parameters such as batch size. We explore the effect of training batch size after handling data imbalance and show the results in Figure 3(b). We can see training batch size of 18 achieves best performance for Visual7W.

Batch normalization. Batch Normalization (BN) is an important component that can speed up the convergence of training. We explore where to add BN, such as the source features (i.e., $\mathbf{x}_Q, \mathbf{x}_I$ or \mathbf{x}_{A_i}), after the first stage (i.e., \mathbf{x}_{QI}), or after the second stage (i.e., \mathbf{x}_{QIA_i}). Figure 3(d) shows that adding BN to \mathbf{x}_{QIA_i} not only maintains efficient training speed, but also improves the performance.

Evaluation on POS Tag guided Attention: POS tags provide an intuitive mechanism to guide word-level attention. But to achieve the optimal performance, a few practical concerns have to be taken care of. Specifically, we found weight initialization was important. And the best performance is gained when the POS tag guided attention weights are initialized with a uniform distribution between 0 and 2. The performance comparison before and after adding POS tag guided attention is in Table.1. POS tag guided attention alone helps improve performance by 0.7% on Visual7W.

Evaluation on Convolutional N-Gram: We propose the convolutional n-gram mechanism to incorporate the contextual information in each word location. Specifically, we apply 1D convolution with window size of 1 word, 2 words, 3 words, etc, and use max pooling along each dimension between different convolutional n-grams. The exploration of window size is illustrated in Figure 3(c). We adopt a convolutional 3-gram of window size 1, 2 and 3 for its efficiency and effectiveness. As shown in Table. 1, convolutional 3-gram alone helps improve performance by 0.6% on Visual7W.




Figure 3. Good practices of histograms on data imbalance, appropriate training batch size, batch normalization (with time and performance) and convolutional n-gram. λ_2 and margin is also explored in this picture.

Method	Visual7W	VQA validation
Our Baseline	65.6	58.3
+POS tag guided attention (POS-Att)	66.3	58.7
+Convolutional N-Gram (Conv N-Gram)	66.2	59.3
+POS-Att +Conv N-Gram	66.6	59.5
+POS-Att +Conv N-Gram +Triplet attention-Q	66.8	60.1
+POS-Att +Conv N-Gram +Triplet attention-A	67.0	60.1
+POS-Att +Conv N-Gram +Triplet attention-Q+A	67.3	60.2
+POS-Att +Conv N-Gram +Triplet attention-Q+A + Structured Learning Triplet	67.5	60.3

Table 1. Performance of our model on Visual7W and VQA validation set, for fast ablation study, we use the $7*7$ image feature instead of $14*14$ image feature by feed the $1/2$ height and $1/2$ width of original size image to the network

Evaluation on Triplet Attention: Our triplet attention model spatially attends to image regions based on both question and answer. We try to add only question attention att_{Q-I} , only answer attention att_{A-I} , and both question and answer attention att_I (initial λ_1 is set as 0.5). Resulting comparisons are shown in Table. 1. Answer attention alone improves more than question alone while our proposed triplet attention mechanism improves the performance by 0.7%.

Evaluation on Structured Learning on Triplets: Our structured learning on triplets mechanism not only uses the traditional binary classification loss but also encourages large margin between positive and hard negative answers. It helps to improve performance on Visual7W from 67.3 to 67.5 and on VQA validation set from 60.2 to 60.3. We explore the parameter λ_2 , which balances the two losses and *margin*, to mine hard negative samples. And they are illustrated in Figure 3(e) and Figure 3(f). In our final model, λ_2 is set as

0.5 for Visual7W and 0.05 for VQA dataset while margin is set as 0.2 for both dataset.

State-of-the-art performance: After verifying the effectiveness of all the components above, we compare our full model with the state-of-the-art performance in Table. 2. For Visual7W, our model improves the state-of-the-art performance by 1.1%. And we compare our model on VQA with other state-the-art performance based on the same training schema(no extra data, no extra language embedding, single model performance). We also get the best performance (69.6%) on VQA test-standard set.

4.4. Visualization and Analysis

We visualize the POS tag guided attention for questions/answers and the triplet attention for images in Fig 4. The first column is original image. The second column is the visualization based on question with wrong answer attention. It should be compared with the fourth column,

Method	Visual7W	VQA Test Standard	VQA Test Dev all	Y/N	Num	Other
Co-attention [19]	-	66.1	65.8	79.7	40.0	59.8
RAU [22]	-	67.3	67.7	81.9	41.1	61.5
Attention-LSTM [42]	55.6	-	-	-	-	-
MCB + Att [7]	62.2	-	68.6	-	-	-
Zero-shot [27]	65.7	-	-	-	-	-
MLP [13]	67.1	68.9	65.2	80.8	17.6	62.0
VQS [8]	-	-	68.9	80.6	39.4	65.3
Full model	68.2	69.6	69.7	81.9	44.3	64.7

Table 2. Quantitative results on Visual7W [42], the test2015-standard split on VQA Real Multiple Choice [3] and the test2015-develop split with each division (Y/N, number, and others). For fair comparison, we only compare with the single model performance without using other large dataset (e.g. Visual Gnome) for pre-training. Our full model outperforms the state-of-the-art performance by a large margin on Visual7W and obtains a competitive result on VQA test-standard set.

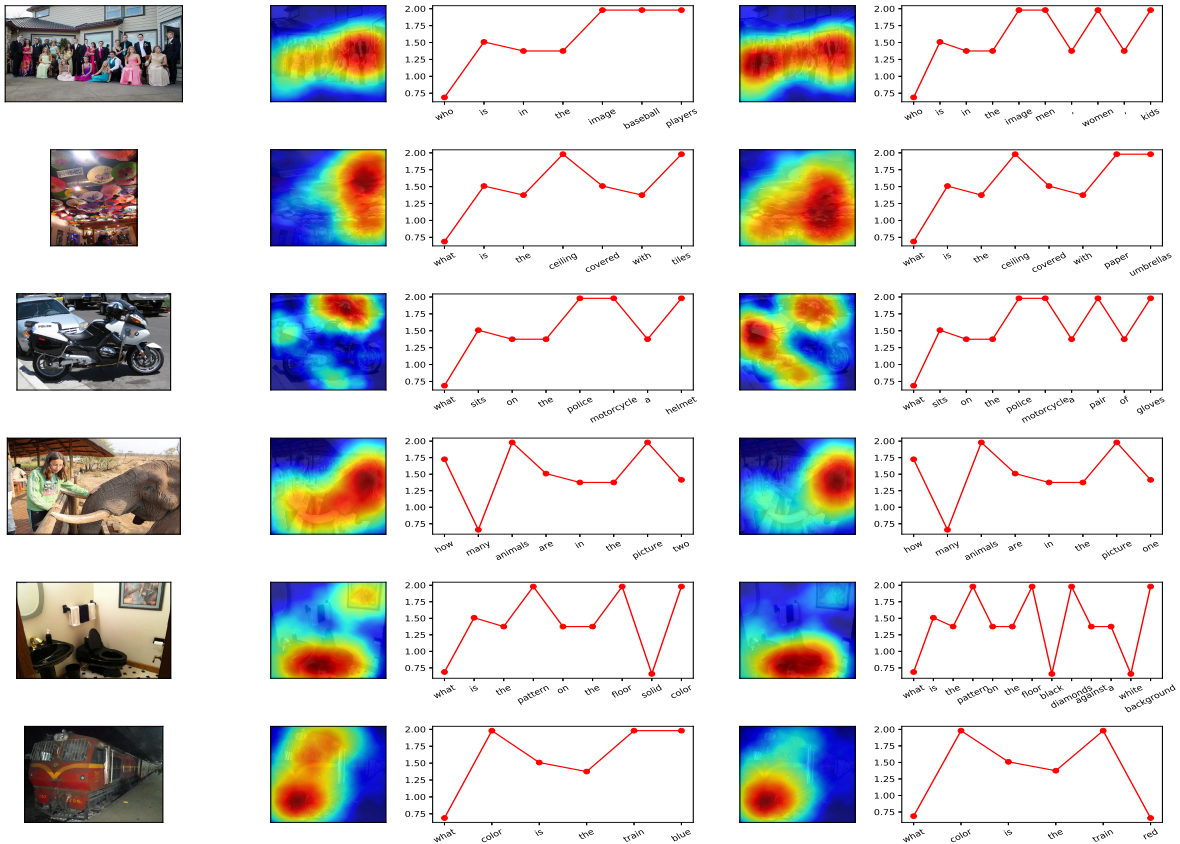


Figure 4. Best viewed in color. The first column is original image. The second column is question and wrong answer attention visualization. It should be compared with the fourth column, which is the question and right answer attention visualization. The third column and the fifth column is the question with wrong answers and question with right answers.

which is the visualization based on question with right answer attention. The third column and the fifth column is the question with wrong answers and question with right answers. The wrong question answer pairs from top to bottom in the third column are: "Who is in the image, basketball players", "what is the ceiling covered with, ties", "what sits on the police motorcycle, a helmet", "how many animals

are in the picture, two", "what is the pattern on the floor, solid color", "what color is the train, blue". And the right question answer pairs from top to bottom in the fifth column are: "Who is in the image, women kids", "what is the ceiling covered with, paper umbrellas", "what sits on the police motorcycle, a pair of gloves", "how many animals are in the picture, one", "what is the pattern on the floor,

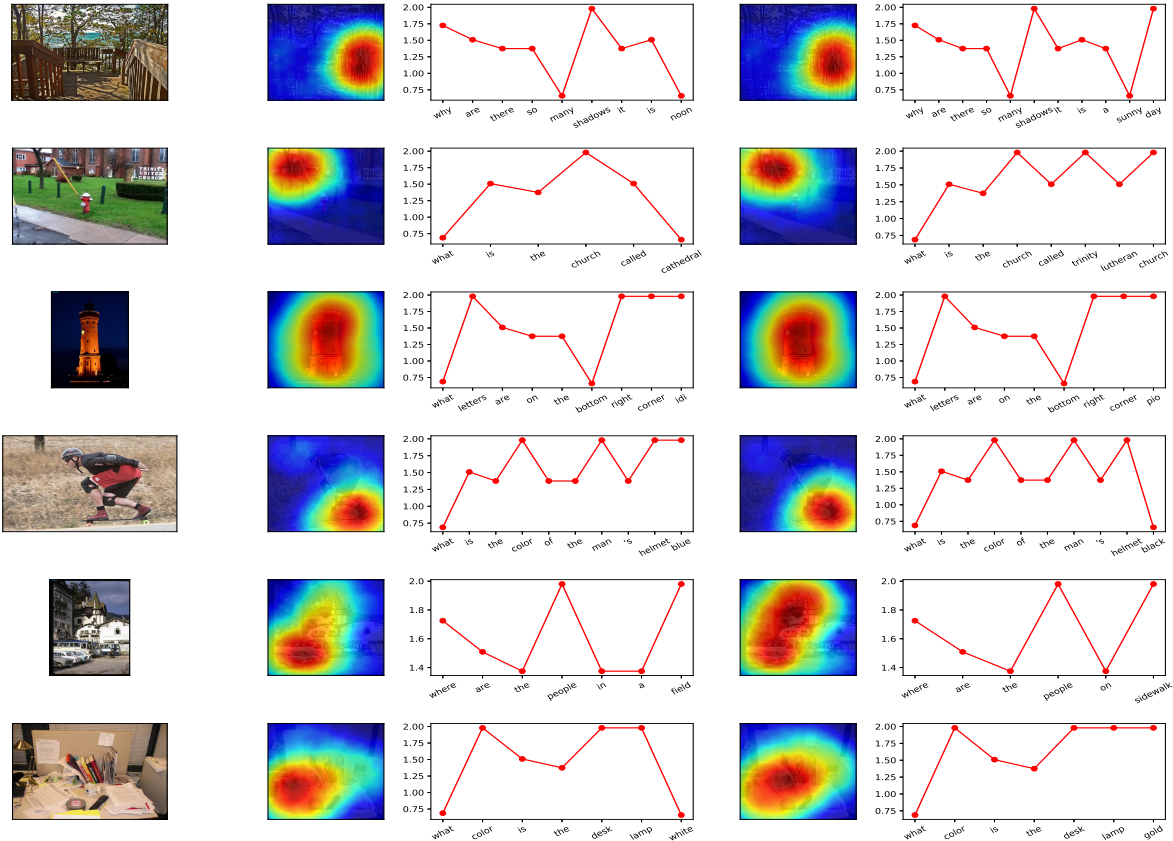


Figure 5. Failure case visualization. The layout is the same as Fig 4.

diamonds against a white background”, ”what color is the train, red”. The first and second row shows our VQA system are able to link the attention to certain enough objects in images instead of part of them while the third row shows our VQA system’s ability to distinguish what really represent ”A policeman” and what relates to them. The fourth to sixth row shows our VQA system can spot what should be spotted without redundant area.

We can conclude that both our POS tag guided and triplet attentions help the model focus on the desired part, and is thus beneficial for final reasoning.

We also visualize some failure case in Fig 5. In these cases, right question answer pairs either attend the same wrong image part with wrong question answer pair, or attend to some unreasonable place.

5. Conclusion

This paper presents a simple yet highly effective model for visual question answering. We attribute this good performance to the novel POS tag guided and triplet attention mechanisms, as well as a series of good practices including structured learning for triplets. The former provides a justification and interpretation of our model, while the latter

makes it possible for our model to achieve fast convergence and achieve better discriminative ability during learning.

Acknowledgement

This work was supported in part by NSF grants IIS-1618806, IIS-1253538 and a hardware donation from NVIDIA. Zhe Wang personally thanks Mr. Jianwei Yang for the helpful discussion.

References

- [1] A. Agrawal, D. Batra, and D. Parikh. Analyzing the behavior of visual question answering models. In *EMNLP*, 2016.
- [2] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein. Deep compositional question answering with neural module networks. In *CVPR*, 2016.
- [3] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. Lawrence Zitnick, and D. Parikh. VQA: Visual question answering. In *ICCV*, 2015.
- [4] J. P. Bigham, C. Jayant, H. Ji, G. Little, A. Miller, R. C. Miller, R. Miller, A. Tatarowicz, B. White, S. White, , and T. Yeh. VizWiz: Nearly realtime answers to visual questions. In *User Interface Software and Technology*, 2010.
- [5] M. A. .et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. 2016. arXiv:1603.04467.

- [6] H. Fang et al. From captions to visual concepts and back. In *CVPR*, 2015.
- [7] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016.
- [8] C. Gan, Y. Li, H. Li, C. Sun, and B. Gong. Vqs: Linking segmentations to questions and answers for supervised attention in vqa and question-focused semantic segmentation. 2017. *ICCV*.
- [9] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [10] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. 2016. arXiv.org:1612.00837.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [12] X. He et al. Image caption generation with part of speech guidance. *Pattern Recognition Letters*, 2017.
- [13] A. Jabri, A. Joulin, and L. van der Maaten. Revisiting visual question answering baselines. In *ECCV*, 2016.
- [14] A. Jiang, F. Wang, F. Porikli, and Y. Li. Compositional memory for visual question answering. 2015. arXiv.1511.05676.
- [15] T. Kenter, A. Borisov, and M. de Rijke. Siamese CBOW: Optimizing word embeddings for sentence representations. In *Dutch-Belgian Information Retrieval Workshop*, 2016.
- [16] Y. Kim. Convolutional neural networks for sentence classification. In *EMNLP*, 2014.
- [17] R. Li and J. Jia. Visual question answering with question representation update. In *NIPS*, 2016.
- [18] J. Lu, A. Kannan, J. Yang, D. Parikh, and D. Batra. Best of both worlds: Transferring knowledge from discriminative learning to a generative visual dialog model. In *NIPS*, 2017.
- [19] J. Lu, J. Yang, D. Batra, and D. Parikh. Hierarchical question-image co-attention for visual question answering. In *NIPS*, 2016.
- [20] L. Ma, Z. Lu, and H. Li. Learning to answer questions from image using convolutional neural network. In *AAAI*, 2016.
- [21] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *NIPS*, 2014.
- [22] H. Noh and B. Han. Training recurrent answering units with joint loss minimization for vqa. 2016. arXiv:1606.03647.
- [23] D. S.-S. D. B. D. P. Peng Zhang, Yash Goyal. Yin and yang: Balancing and answering binary visual questions. In *CVPR*, 2016.
- [24] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014.
- [25] M. Ren, R. Kiros, and R. Zemel. Exploring models and data for image question answering. In *NIPS*, 2015.
- [26] K. J. Shih, S. Singh, and D. Hoiem. Where to look: Focus regions for visual question answering. In *CVPR*, 2016.
- [27] D. Teney and A. van den Hengel. Zero-shot visual question answering. 2016. arXiv.1611.05546.
- [28] L. Wang, Z. Wang, S. Guo, and Y. Qiao. Better exploiting os-cnns for better event recognition in images. In *ICCVW*, 2015.
- [29] L. Wang, Z. Wang, Y. Qiao, and L. V. Gool. Transferring deep object and scene representations for event recognition in still images. *IJCV*, 2017.
- [30] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. V. Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016.
- [31] Z. Wang, L. Wang, W. Du, and Y. Qiao. Exploring fisher vector and deep networks for action spotting. In *CVPRW*, 2015.
- [32] Z. Wang, L. Wang, Y. Wang, B. Zhang, and Y. Qiao. Weakly supervised patchnets: Describing and aggregating local patches for scene recognition. *TIP*, 26(4):2028–2041, 2017.
- [33] Q. Wu, C. S. Peng Wang, A. Dick, and A. van den Hengel. Ask me anything: Free-form visual question answering based on knowledge from external sources. In *CVPR*, 2016.
- [34] C. Xiong, S. Merity, and R. Socher. Dynamic memory networks for visual and textual question answering. In *ICML*, 2016.
- [35] H. Xu and K. Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *ECCV*, 2016.
- [36] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *CVPR*, 2016.
- [37] N. Zhenxing et al. Hierarchical multimodal lstm for dense visual-semantic embedding. In *ICCV*, 2017.
- [38] B. Zhou, Y. Tian, S. Sukhbaatar, A. Szlam, and R. Fergus. Simple baseline for visual question answering. 2015. arXiv:1512.02167.
- [39] W. Zhu et al. Deep multi-instance networks with sparse label assignment for whole mammogram classification. In *MICCAI*, 2017.
- [40] W. Zhu et al. Adversarial deep structured nets for mass segmentation from mammograms. *IEEE ISBI*, 2018.
- [41] W. Zhu et al. Deeplung: Deep 3d dual path nets for automated pulmonary nodule detection and classification. In *IEEE WACV*, 2018.
- [42] Y. Zhu, O. Groth, M. Bernstein, and L. FeiFei. Visual7W: Grounded Question Answering in Images. In *CVPR*, 2016.