

Revisiting Visual Question Answering Baselines

Allan Jabri^(✉), Armand Joulin, and Laurens van der Maaten

Facebook AI Research, New York, USA
{ajabri, ajoulin, lvdmaaten}@fb.com

Abstract. Visual question answering (VQA) is an interesting learning setting for evaluating the abilities and shortcomings of current systems for image understanding. Many of the recently proposed VQA systems include attention or memory mechanisms designed to perform “reasoning”. Furthermore, for the task of multiple-choice VQA, nearly all of these systems train a multi-class classifier on image and question features to predict an answer. This paper questions the value of these common practices and develops a simple alternative model based on binary classification. Instead of treating answers as competing choices, our model receives the answer as input and predicts whether or not an image-question-answer triplet is correct. We evaluate our model on the Visual7W Telling and the VQA Real Multiple Choice tasks, and find that even simple versions of our model perform competitively. Our best model achieves state-of-the-art performance of 65.8% accuracy on the Visual7W Telling task and compares surprisingly well with the most complex systems proposed for the VQA Real Multiple Choice task. Additionally, we explore variants of the model and study the transferability of the model between both datasets. We also present an error analysis of our best model, the results of which suggest that a key problem of current VQA systems lies in the lack of visual grounding and localization of concepts that occur in the questions and answers.

Keywords: Visual question answering · Dataset bias

1 Introduction

Recent advances in computer vision have brought us close to the point where traditional object-recognition benchmarks such as Imagenet are considered to be “solved” [1, 2]. These advances, however, also prompt the question how we can move from object recognition to *visual understanding*; that is, how we can extend today’s recognition systems that provide us with “words” describing an image or an image region to systems that can produce a deeper semantic representation of the image content. Because benchmarks have traditionally been a key driver for progress in computer vision, several recent studies have proposed methodologies to assess our ability to develop such representations. These proposals include modeling relations between objects [3], visual Turing tests [4], and visual question answering [5–8].



What color is the jacket?
 -Red and blue.
 -Yellow.
 -Black.
 -Orange.



How many cars are parked?
 -Four.
 -Three.
 -Five.
 -Six.



What event is this?
 -A wedding.
 -Graduation.
 -A funeral.
 -A picnic.



When is this scene taking place?
 -Day time.
 -Night time.
 -Evening.
 -Morning.

Fig. 1. Four images with associated questions and answers from the Visual7W dataset. Correct answers are typeset in green. (Color figure online)

The task of Visual Question Answering (VQA) is to answer questions—posed in natural language—about an image by providing an answer in the form of short text. This answer can either be selected from multiple pre-specified choices or be generated by the system. As illustrated by the examples from the Visual7W dataset [9] in Fig. 1, VQA naturally combines computer vision with natural language processing and reasoning, which makes it a good way to study progress on the path from computer vision to more general artificially intelligent systems.

VQA seems to be a natural playground to develop approaches able to perform basic “reasoning” about an image. Recently, many studies have explored this direction by adding simple memory or attention-based components to VQA systems. While in theory, these approaches have the potential to perform simple reasoning, it is not clear if they do actually reason, or if they do so in a human-comprehensible way. For example, Das *et al.* [10] recently reported that “machine-generated attention maps are either negatively correlated with human attention or have positive correlation worse than task-independent saliency”. In this work, we also question the significance of the performance obtained by current “reasoning”-based systems. In particular, this study sets out to answer a simple question: are these systems better than baselines designed to solely capture the dataset bias of standard VQA datasets? We limit the scope of our study to multiple-choice tasks, as this allows us to perform a more controlled study that is not hampered by the tricky nuances of evaluating generated text [11, 12].

We perform experimental evaluations on the Visual7W dataset [8] and the VQA dataset [5] to evaluate the quality of our baseline models. We: (1) study and model the bias in the Visual7W Telling and VQA Multiple Choice datasets, (2) measure the effect of using visual features from different CNN architectures, (3) explore the use of a LSTM as the system’s language model, and (4) study transferability of our model between datasets.

Our best model outperforms the current state-of-the-art on the Visual7W telling task with a performance of 65.8%, and competes surprisingly well with the most complex systems proposed for the VQA dataset. Furthermore, our models perform competitively even with missing information (that is, missing

images, missing questions, or both). Taken together, our results suggest that the performance of current VQA systems is not significantly better than that of systems designed to exploit dataset biases.

2 Related Work

The recent surge of studies on visual question answering has been fueled by the release of several visual question-answering datasets, most prominently, the VQA dataset [5], the DAQUAR dataset [13], the Visual Madlibs Q&A dataset [7], the Toronto COCO-QA dataset [6], and the Visual7W dataset [8]. Most of these datasets were developed by annotating subsets of the COCO dataset [14]. Geman *et al.* [4] proposed a visual Turing test in which the questions are automatically generated and require no natural language processing. Current approaches to visual question answering can be subdivided into “generation” and “classification” models:

Generation Models. Malinowski *et al.* [15] train a LSTM model to generate the answer after receiving the image features (obtained from a convolutional network) and the question as input. Wu *et al.* [16] extend a LSTM generation model to use external knowledge that is obtained from DBpedia [17]. Gao *et al.* [18] study a similar model but decouple the LSTMs used for encoding and decoding. Whilst generation models are appealing because they can generate arbitrary answers (also answers that were not observed during training), in practice, it is very difficult to jointly learn the encoding and decoding models from the question-answering datasets of limited size. In addition, the evaluation of the quality of the generated text is complicated in practice [11, 12].

Classification Models. Zhou *et al.* [9] study an architecture in which image features are produced by a convolutional network, question features are produced by averaging word embeddings over all words in the question, and a multi-class logistic regressor is trained on the concatenated features; the top unique answers are treated as outputs of the classification model. Similar approaches are also studied by Antol *et al.* [5] and Ren *et al.* [6], though they use a LSTM to encode the question text instead of an average over word embeddings. Zhu *et al.* [8] present a similar method but extend the LSTM encoder to include an attention mechanism for jointly encoding the question with information from the image. Ma *et al.* [19] replace the LSTM encoder by a one-dimensional convolutional network that combines the word embeddings into a question embedding. Andreas *et al.* [20] use a similar model but perform the image processing using a compositional network whose structure is dynamically determined at run-time based on a parse of the question. Fukui *et al.* [21] propose the use of “bilinear pooling” for combining multi-modal information. Lu *et al.* [22] jointly learn a hierarchical attention mechanism based on parses of the question and the image which they call “question-image co-attention”.

Our study is most closely related to a recent study by Shih *et al.* [23], which also considers models that treat the answer as an input variable and predicts

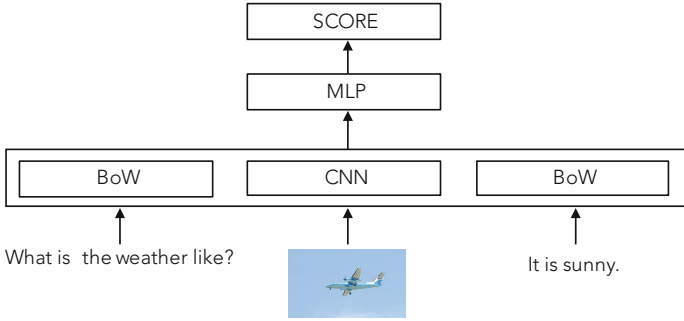


Fig. 2. Overview of our system for visual question answering. See text for details.

whether or not an image-question-answer triplet is correct. However, their study develops a substantially more complex pipeline involving image-region selection.

3 System Overview

Figure 2 provides an overview of the architecture of our visual question answering system. The system takes an image-question-answer feature triplet as input. Unless otherwise stated (that is, in the LSTM experiment of Sect. 4), both the questions and the answers are represented by averaging word2vec embeddings over all words in the question or answer, respectively. The images are represented using features computed by a pre-trained convolutional network. Unless otherwise stated, we use the penultimate layer of Resnet-101 [2]. The word2vec embeddings are 300-dimensional and the image features are 2,048-dimensional. The three feature sets are concatenated and used to train a classification model that predicts whether or not the image-question-answer triplet is correct.

The classification models we consider are logistic regressors and multilayer perceptrons (MLP) trained on the concatenated features, and bilinear models that are trained on the answer features and a concatenation of the image and question features. The MLP has 8,192 hidden units unless otherwise specified. We use dropout [24] after the first layer. We denote the image, question, and answer features by \mathbf{x}_i , \mathbf{x}_q , and \mathbf{x}_a , respectively. Denoting the sigmoid function $\sigma(x) = 1/(1 + \exp(-x))$ and the concatenation operator $\mathbf{x}_{iq} = \mathbf{x}_i \oplus \mathbf{x}_q$, we define the models as follows:

$$\text{Linear: } y = \sigma(\mathbf{W}\mathbf{x}_{iqa} + b)$$

$$\text{Bilinear: } y = \sigma(\mathbf{x}_{iq}^\top \mathbf{W}\mathbf{x}_a + b)$$

$$\text{MLP: } y = \sigma(\mathbf{W}_2 \max(0, \mathbf{W}_1 \mathbf{x}_{iqa}) + b).$$

The parameters of the classifier are learned by minimizing the binary logistic loss of predicting whether or not an image-question-answer triplet is correct using stochastic gradient descent. During training we sampled two negative examples

Table 1. Comparison of our models with the state-of-the-art for the Visual7W telling task [8]. Human accuracy on the task is 96.0%. Higher values are better.

Method	What	Where	When	Who	Why	How	Overall
LSTM (Q + I) [15]	48.9	54.4	71.3	58.1	51.3	50.3	52.1
LSTM-Att [8]	51.5	57.0	75.0	59.5	55.5	49.8	55.6
MCB [21]	60.3	70.4	79.5	69.2	58.2	51.1	62.2
MLP (A)	47.1	57.8	73.6	63.3	57.6	36.3	50.7
MLP (A + Q)	54.9	60.0	76.0	65.9	63.8	40.2	56.1
MLP (A + I)	60.0	73.9	80.0	70.3	63.9	36.6	61.0
MLP (A + Q + I)	64.4	76.2	82.3	72.8	69.5	40.9	64.8

Table 2. Comparison of our models with the state-of-the-art single models for the VQA Real Multiple Choice task [5]. Results are reported on the test2015-standard split. Human accuracy on the task is 83.3%. Higher values are better.

Method	Yes/No	Number	Other	All
Two-layer LSTM [5]	80.6	37.7	53.6	63.1
Region selection [23]	77.2	33.5	56.1	62.4
Question-image co-attention [22]	80.0	39.5	59.9	66.1
MLP (A + Q + I)	80.8	17.6	62.0	65.2

from the multiple choices for each positive example, for a maximum of 300 epochs. The convolutional networks were pre-trained on the Imagenet dataset, following [25], and were not further finetuned. We used pre-trained word2vec [26] embeddings, which we did not finetune on VQA data either.

4 Experiments

We perform experiments on the following two datasets:

Visual7W Telling [8]. The dataset includes 69,817 training questions, 28,020 validation questions, and 42,031 test questions. Each question has four answer choices. The negative choices are human-generated on a per-question basis. The performance is measured by the percentage of correctly answered questions.

VQA Real Multiple Choice [5]. The dataset includes 248,349 questions for training, 121,512 for validation, and 244,302 for testing. Each question has 18 answer choices. The negative choices are randomly sampled from a predefined set of answers. Performance is measured following the metric proposed by [5].

4.1 Comparison with State-of-the-Art

We first compare the MLP variant of our model with the state-of-the-art. Table 1 shows the results of this comparison on Visual7W, using three variants of our

Table 3. Accuracy of models using either a *softmax* or a *binary* loss. Results are presented for different models using answer, question and image. On VQA, we use the test2015-dev split. Higher values are better.

Dataset	Model	Softmax	Binary
Visual7W	Linear	41.6	42.7
	Bilinear	–	61.8
	MLP	50.2	64.8
VQA	MLP	61.1	64.9

Table 4. The five most similar answers in the Visual7W dataset for three answers appearing in that dataset (in terms of cosine similarity between their feature vectors).

During the daytime	On the bus stop bench	On a tree branch
During daytime	Bus bench	On the tree branch
Outside, during the daytime	In front of the bus stop	The tree branch
Inside, during the daytime	The bus stop	Tree branch
In the daytime	At the bus stop	A tree branch
In the daytime	The sign on the bus stop	Tree branches

baseline with different inputs: (1) answer and question (A + Q); (2) answer and image (A + I); (3) and all three inputs (A + Q + I). The model achieves state-of-the-art performance when it has access to all the information. Interestingly, as shown by the results with the A + Q variant of our model, simply exploiting the most frequent question-answer pairs obtains competitive performance. Surprisingly, even a variant of our model that is trained *on just the answers* already achieves a performance of 50.7%, simply by learning biases in the answer distribution.

In Table 2, we also compare our models with the published state-of-the-art on the VQA dataset. Despite its simplicity, our baseline achieves comparable performance with state-of-the-art models. We note that recent work [21] obtained 70.1%, but used an ensemble of 7 models trained on additional data (the Visual Genome dataset [3]). Nonetheless, [21] performs only 5% better than our model whilst being substantially more complex.

4.2 Additional Experiments

In the following, we present the results of additional experiments to understand why our model performs relatively well, and when it fails. All evaluations are conducted on the Visual7W Telling dataset unless stated otherwise.

Does It Help to Consider the Answer as an Input? In Table 3, we present the results of experiments in which we compare the performance of our (binary) baseline model with variants of the model that predict softmax probabilities over

Table 5. Accuracy on the Visual7W Telling task using visual features produced by five different convolutional networks. Higher values are better.

Model	AlexNet	GoogLeNet	ResNet-34	ResNet-50	ResNet-101
(dim.)	(4,096)	(1,792)	(512)	(2,048)	(2,048)
Linear	42.7	42.9	42.7	42.7	42.8
Bilinear	54.9	57.0	58.5	60.7	61.8
MLP	61.6	62.1	63.9	64.2	64.8

a discrete set of the 5,000 most common answers, as is commonly done in most prior studies, for instance, [9].

The results in the table show a substantial advantage of representing answers as inputs instead of outputs for the Visual7W Telling task and the VQA Real Multiple Choice task. Taking the answer as an input allows the system to model the similarity between different answers. For example, the answers “two people” and “two persons” are modeled by disjoint parameters in a softmax model; instead, the binary model will generally assign similar scores to these answers because they have similar bag-of-words word2vec representations.

To illustrate this, Table 4 shows examples of the similarities that can be captured by the binary model. For a given answer, the table shows the five most similar answers in the dataset based on cosine similarity between the feature vectors. The binary model can readily exploit these similarities, whereas a softmax model has to learn them from the (relatively small) Visual7W training set.

Interestingly, the gap between the binary and softmax models is smaller on the VQA datasets. This result may be explained by the way the incorrect-answer choices were produced in both datasets: the choices are human-generated for each question in the Visual7W dataset, whereas in the VQA dataset, the choices are randomly chosen from a predefined set that includes irrelevant correct answers.

What is the Influence of Convolutional Network Architectures? Nearly all prior work on VQA uses features extracted using a convolutional network that is pre-trained on Imagenet to represent the image in an image-question pair. Table 5 shows to what extent the quality of these features influences the VQA performance by comparing five different convolutional network architectures: AlexNet [27], GoogLeNet [1], and residual networks with three different depths [2]. While the performance on Imagenet is correlated with performance in visual question answering, the results show this correlation is quite weak: a reduction in the Imagenet top-5 error of 18% corresponds to an improvement of only 3% in question-answering performance. This result suggests that the performance on VQA tasks is limited by either the fact that some of the visual concepts in the questions do not appear in Imagenet, or by the fact that the convolutional networks are only trained to recognize object presence and not to predict higher-level information about the visual content of the images.

Table 6. Accuracy on Visual7W Telling dataset of a bag-of-words (*BoW*) and a *LSTM* model. We did not use image features to isolate the difference between language models. Higher values are better.

Model	BoW	LSTM
Bilinear	51.5	52.5
MLP	56.1	51.0

Table 7. Accuracy on Visual7W of models (1) trained from *scratch*, (2) *transferred* from the VQA dataset, and (3) *finetuned* after transferring. Higher values are better.

Model	Method	What	Where	When	Who	Why	How	Overall
MLP (A + Q)	Scratch	54.9	60.0	76.0	65.9	63.8	40.2	56.1
	Transfer	44.7	38.9	32.9	49.6	45.0	27.3	41.1
MLP (A + I)	Scratch	60.0	73.9	80.0	70.3	63.9	36.6	61.0
	Transfer	28.4	26.6	44.1	37.0	31.7	25.2	29.4
MLP (A + Q + I)	Scratch	64.4	76.2	82.3	72.8	69.5	40.9	64.8
	Transfer	58.7	61.7	41.7	60.2	53.2	29.1	53.8
	Finetune	66.4	76.3	81.6	73.1	68.7	41.7	65.8

Do Recurrent Networks Improve Over Bag of Words? Our baseline uses a simple bag-of-words (BoW) model to represent the questions and answers. Recurrent networks (in particular, LSTMs [28]) form a popular alternative for BoW models. We perform an experiment in which we replace our BoW representations by a LSTM model. The LSTM was trained on the Visual7W Telling training set, using a concatenation of one-hot encodings and pre-trained word2vec embeddings as input for each word in the question.

We experimented with using the average over time of the hidden states as feature representation for the text, as well as using only the last hidden state. We observed little difference between the two; here, we report the results using the last-state representation.

Table 6 presents the results of our experiment comparing BoW and LSTM representations. To study just the difference between the language models, we did not use images features as input in this experiment. The results show that despite their greater representation power, LSTMs actually do not outperform BoW representations on the Visual7W Telling task, presumably, because the dataset is quite small and the LSTM overfits easily. This may also explain why attentional LSTM models [8] perform poorly on the Visual7W dataset.

Can We Transfer Knowledge from VQA to Visual7W? An advantage of the presented model is that it can readily be transferred between datasets: it does not suffer from out-of-vocabulary problems nor does it require the set of answers to be known in advance. Table 7 shows the results of a transfer-learning experiment in which we train our model on the VQA dataset, and

use it to answer questions in the Visual7W dataset. We used three different variants of our model, and experimented with three different input sets. The table presents three sets of results: (1) baseline results in which we trained on Visual7W from *scratch*, (2) *transfer* results in which we train on VQA but test on Visual7W, and (3) results in which we train on VQA, *finetune* on Visual7W, and then test on Visual7W.

The poor performance of the A + I transfer-learning experiment suggests that there is a substantial difference in the answer distribution between both datasets, especially since both use images from [14]. Transferring the full model from VQA to Visual7W works surprisingly well: we achieve 53.8% accuracy, which is less than 2% worse than LSTM-Att [8], even though the model never learns from Visual7W training data. If we finetune the transferred model on the Visual7W dataset, it actually outperforms a model trained from scratch on that same dataset, obtaining an accuracy of **65.8%**. This additional boost likely stems from the model adjusting to the biases in the Visual7W dataset.

5 Error Analysis

To better understand the shortcomings and limitations of our models, we performed an error analysis of the best model we obtained in Sect. 4 on six types of questions, which are illustrated in Figs. 3, 4 and 5.



Fig. 3. Examples of good and bad predictions by our visual question answering model on color and shape questions. Correct answers are typeset in green; incorrect predictions by our model are typeset in red. See text for details. (Color figure online)

Colors and Shapes. Approximately 5,000 questions in the Visual7W test set are about colors and approximately 200 questions are about shapes. While colors and shapes are fairly simple visual features, our models only achieve around 55% accuracy on these types of questions. For reference, our (A + Q) baseline already achieves 50% in accuracy. This means that our models primarily learn the bias in the dataset. For example, for shape, it predicts either “circle”, “round”, or “octagon” when the question is about a “sign”. For color questions, even though the

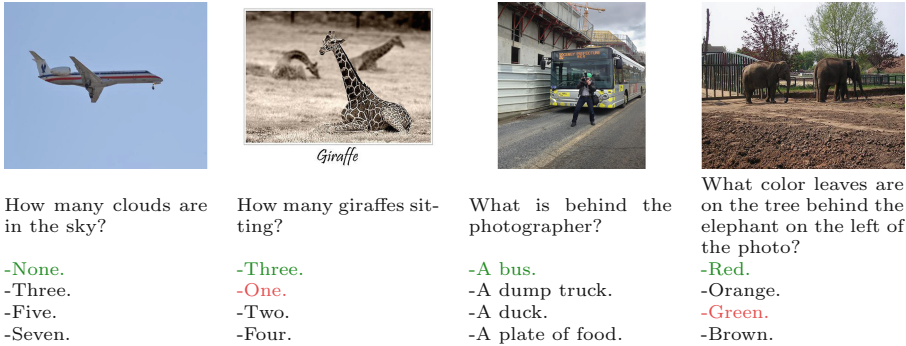


Fig. 4. Examples of good and bad predictions by our visual question answering model on counting and spatial reasoning. Correct answers are typeset in green; incorrect predictions by our model are typeset in red. See text for details. (Color figure online)

performances are similar, it appears that the image-based models are able to capture additional information. For example, Fig. 3 shows that the model tends to predict the most salient color, but fails to capture color coming from small objects, which constitute a substantial number of questions in the Visual7W dataset. This result highlights the limits of using global image features in visual question answering.

Counting. There are approximately 5,000 questions in the Visual7W test set that involve counting the number of objects in the image (“how many ...?”). On this type of questions, our model achieves an accuracy of 36%. This accuracy is hardly better than that the 35% achieved by the (Q + A) baseline. Again, this implies that our model does not really extract information from the image that can be used for counting. In particular, our model has a strong preference for answers such as: “none”, “one”, or “two”.

Spatial Reasoning. We refer to any question that refers to a relative position (“left”, “right”, “behind”, *etc.*) as questions about “spatial reasoning”. There are approximately 1,500 such questions in the Visual7W test set. On questions requiring spatial reasoning, our models achieve an accuracy of approximately 50%, whereas a purely text-based model achieves an accuracy 40%. This suggests that our models, indeed, extract some information from the images that can be used to make inferences about spatial relations.

Actions. We refer to any question that asks what an entity is “doing” as an “action” question. There are approximately 1,200 such questions in the Visual7W test set. Our models achieve an accuracy of roughly 75% on action questions. By contrast, a purely text-based model achieves an accuracy of around 65%. This result suggests that our model does learn to exploit image features in recognizing actions. This result is in line with results presented in earlier studies that show image features transfer well to simple action-recognition tasks [29, 30].



What is the man doing?

- Surfing.
- Singing.
- Working.
- Playing.



What is the man doing?

- Golfing.
- Playing tennis.
- Walking.
- Biking.



Why is the ground white?

- Snow.
- Sand.
- Stones.
- Concrete.



Why is his arm up?

- To serve the tennis ball.
- About to hit the ball.
- Reaching for the ball.
- Swinging his racket.

Fig. 5. Examples of good and bad predictions by our visual question answering model on action and causality. Correct answers are typeset in green; incorrect predictions by our model are typeset in red. See text for details. (Color figure online)

Causality. “Why” questions test the model’s ability to capture a weak form of causality. There are around 2,600 of them. Our model has an accuracy of 68% on such questions, but a simple text-based model already obtains 62%. This means that most “why” questions can be answered by looking at the text. This is unsurprising, as many of these questions refer to common sense that is encoded in the text. For example, in Fig. 5, one hardly needs the image to correctly predict that the ground is “white” because of “snow” instead of “sand”.

6 Discussion and Future Work

This paper presented a simple alternative model for visual question answering multiple choice, explored variants of this model, and experimented with transfer between VQA datasets. Our study produced stronger baseline systems than those presented in prior studies. In particular, our results demonstrate that featurizing the answers and training a binary classifier to predict correctness of an image-question-answer triplet leads to substantial performance improvements over the current state-of-the-art on the Visual7W Telling task: our best model obtains an accuracy of 64.8% when trained from scratch, and 65.8% when transferred from VQA and finetuned on the Visual7W. On the VQA Real Multiple Choice task, our model outperforms models that use LSTMs and attention mechanisms, and is close to the state-of-the-art despite being very simple.

Our error analysis demonstrates that future work in visual question answering should focus on grounding the visual entities that are present in the images, as the “difficult” questions in the Visual7W dataset cannot be answered without such grounding. Whilst global image features certainly help in visual question answering, they do not provide sufficient grounding of concepts of interest. More precise grounding of visual entities, as well as reasoning about the relations between these entities, is likely to be essential in making further progress.

Furthermore, in order to accurately evaluate future models, we need to understand the biases in VQA datasets. Many of the complex methods in prior work

perform worse than the simple model presented in this paper. We hypothesize that one of two things (or both) may explain these results: (1) it may be that, currently, the best-performing models are those that can exploit biases in VQA datasets the best, *i.e.*, models that “cheat” the best; (2) it may be that current, early VQA models are unsuitable for the difficult task of visual question answering, as a result of which all of them hit roughly the same ceiling in experiments and evaluations. In some of our experiments, we have seen that a model that appears qualitatively better may perform worse quantitatively, because it captures dataset biases less well. To address such issues, it may be necessary to consider alternative evaluation criteria that are less sensitive to dataset bias.

Finally, the results of our transfer-learning experiments suggest that exploring the ability of VQA systems to generalize across datasets may be an interesting alternative way to evaluate such systems in future work.

References

1. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2015)
2. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2016)
3. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalanditis, Y., Li, L.J., Shamma, D., Bernstein, M., Fei-Fei, L.: Visual genome: connecting language and vision using crowdsourced dense image annotations. [arXiv:1602.07332](https://arxiv.org/abs/1602.07332) (2016)
4. Geman, D., Geman, S., Hallonquist, N., Younes, L.: Visual turing test for computer vision systems. *Proc. Natl. Acad. Sci.* **112**(12), 3618–3623 (2015)
5. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C., Parikh, D.: VQA: visual question answering. In: Proceedings of the International Conference on Computer Vision (2015)
6. Ren, M., Kiros, R., Zemel, R.: Exploring models and data for image question answering. In: Advances in Neural Information Processing Systems (2015)
7. Yu, L., Park, E., Berg, A., Berg, T.: Visual madlibs: fill in the blank image generation and question answering. [arXiv:1506.00278](https://arxiv.org/abs/1506.00278) (2015)
8. Zhu, Y., Groth, O., Bernstein, M., Fei-Fei, L.: Visual7W: grounded question answering in images. [arXiv:1511.03416](https://arxiv.org/abs/1511.03416) (2015)
9. Zhou, B., Tian, Y., Sukhbataar, S., Szlam, A., Fergus, R.: Simple baseline for visual question answering. [arXiv:1512.02167](https://arxiv.org/abs/1512.02167) (2015)
10. Das, A., Agrawal, H., Zitnick, C.L., Parikh, D., Batra, D.: Human attention in visual question answering: do humans and deep networks look at the same regions? [arXiv:1606.03556](https://arxiv.org/abs/1606.03556) (2016)
11. Koehn, P.: Statistical significance tests for machine translation evaluation. In: EMNLP, pp. 388–395 (2004)
12. Callison-Burch, C., Osborne, M., Koehn, P.: Re-evaluation the role of BLEU in machine translation research. In: EACL, vol. 6, pp. 249–256 (2006)
13. Malinowski, M., Fritz, M.: A multi-world approach to question answering about real-world scenes based on uncertain input. *CoRR abs/1410.0210* (2014)

14. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.: Microsoft COCO: common objects in context. In: Proceedings of the European Conference on Computer Vision (2014)
15. Malinowski, M., Rohrbach, M., Fritz, M.: Ask your neurons: a neural-based approach to answering questions about images. In: Proceedings of the International Conference on Computer Vision (2015)
16. Wu, Q., Shen, C., van den Hengel, A., Wang, P., Dick, A.: Image captioning and visual question answering based on attributes and their related external knowledge. [arXiv:1603.02814](https://arxiv.org/abs/1603.02814) (2016)
17. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: DBpedia: a nucleus for a web of open data. In: Aberer, K., Choi, K.-S., Noy, N., Allemang, D., Lee, K.-I., Nixon, L., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., Schreiber, G., Cudré-Mauroux, P. (eds.) ASWC/ISWC -2007. LNCS, vol. 4825, pp. 722–735. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-76298-0_52](https://doi.org/10.1007/978-3-540-76298-0_52)
18. Gao, H., Mao, J., Zhou, J., Huang, Z., Wang, L., Xu, W.: Are you talking to a machine? Dataset and methods for multilingual image question answering. In: Advances in Neural Information Processing Systems (2015)
19. Ma, L., Lu, Z., Li, H.: Learning to answer questions from image using convolutional neural network. [arXiv:1506.00333](https://arxiv.org/abs/1506.00333) (2015)
20. Andreas, J., Rohrbach, M., Darrell, T., Klein, D.: Deep compositional question answering with neural module networks. [arXiv:1511.02799](https://arxiv.org/abs/1511.02799) (2015)
21. Fukui, A., Huk Park, D., Yang, D., Rohrbach, A., Darrell, T., Rohrbach, M.: Multi-modal compact bilinear pooling for visual question answering and visual grounding (2016)
22. Lu, J., Yang, J., Batra, D., Parikh, D.: Hierarchical question-image co-attention for visual question answering (2016)
23. Shih, K.J., Singh, S., Hoiem, D.: Where to look: Focus regions for visual question answering. [arXiv:1511.07394](https://arxiv.org/abs/1511.07394) (2016)
24. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014)
25. Gross, S., Wilber, M.: Training and investigating residual nets (2016)
26. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. [arXiv:1301.3781](https://arxiv.org/abs/1301.3781) (2013)
27. Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems (2012)
28. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
29. Joulin, A., van der Maaten, L., Jabri, A., Vasilache, N.: Learning visual features from large weakly supervised data. [arXiv:1511.0225](https://arxiv.org/abs/1511.0225) (2015)
30. Razavian, A.S., Azizpour, H., Sullivan, J., Carlsson, S.: CNN features off-the-shelf: an astounding baseline for recognition. [arXiv:1403.6382](https://arxiv.org/abs/1403.6382) (2014)