

Tips and Tricks for Visual Question Answering: Learnings from the 2017 Challenge

Damien Teney¹ Peter Anderson^{2*} Xiaodong He³ Anton van den Hengel¹

¹University of Adelaide, ²Australian National University, ³JD AI Research

damien.teney@adelaide.edu.au, peter.anderson@anu.edu.au
 xiaodong.he@jd.com, anton.vandenhengel@adelaide.edu.au

Abstract

Deep Learning has had a transformative impact on Computer Vision, but for all of the success there is also a significant cost. This is that the models and procedures used are so complex and intertwined that it is often impossible to distinguish the impact of the individual design and engineering choices each model embodies. This ambiguity diverts progress in the field, and leads to a situation where developing a state-of-the-art model is as much an art as a science. As a step towards addressing this problem we present a massive exploration of the effects of the myriad architectural and hyperparameter choices that must be made in generating a state-of-the-art model. The model is of particular interest because it won the 2017 Visual Question Answering Challenge. We provide a detailed analysis of the impact of each choice on model performance, in the hope that it will inform others in developing models, but also that it might set a precedent that will accelerate scientific progress in the field.

1. Introduction

The task of Visual Question Answering (VQA) involves an image and a related text question, to which the machine must determine the correct answer (see Fig. 1). The task lies at the intersection of the fields of computer vision, natural language processing, and artificial intelligence. This paper presents a relatively simple model for VQA that achieves state-of-the-art results. It is based on a deep neural network that implements the well-known joint embedding approach. The details of its architecture and hyperparameters were carefully selected for optimal performance on the VQA v2 benchmark [12]. Admittedly, a large part of such a search is necessarily guided by empirical exploration and validation. Given the limited understanding of neural networks trained on a task as complex as VQA, small variations of hyperparameters and of network architectures may have significant and sometimes unpredictable effects on final performance [20]. The aim of this paper is to share

*Work performed while interning at Microsoft Research.



What is on the coffee table ?
 candles



What color is the hydrant ?
 black and yellow



What is on the bed ?
 books



What is the long stick for ?
 whipping

Figure 1. The task of visual question answering (VQA) relates visual concepts with elements of language and, occasionally, common-sense or general knowledge. Examples of training questions and their correct answer from the VQA v2 dataset [12].

the details of a successful model for VQA. The findings reported herein may serve as a basis for future development of VQA systems and multimodal reasoning algorithms in general.

The proposed model is based on the principle of a joint embedding of the input question and image, followed by a multi-label classifier over a set of candidate answers. While this general approach forms the basis of many modern VQA methods [33], the details of the model are critical to achieving a high quality result. We also complement our model with a few key technical innovations that greatly enhance its performance. We have conducted an extensive empirical study to explore the space of architectures and hyperparameters to determine the importance of the various components.

Our key findings are summarized with the following characteristics of the proposed model, which enable its high performance (see also Table 1).

- Using a **sigmoid output** that allows multiple correct answers per question, instead of a common single-label softmax.
- Using **soft scores as ground truth targets** that cast the task as a regression of scores for candidate answers, instead of a traditional classification.
- Using **gated tanh activations** in all non-linear layers.
- Using **image features from bottom-up attention** [3] that provide region-specific features, instead of traditional grid-like feature maps from a CNN.
- Using **pretrained representations of candidate answers** to initialize the weights of the output layer.
- Using **large mini-batches** and **smart shuffling** of training data during stochastic gradient descent.

Our hope is that presenting these findings will alleviate the need for others to explore some of the intricacies of the many design and implementation decisions that need to be made in developing state-of-the-art models. In doing so we hope to further progress in the field by illuminating the distinction between the impact of a novel design from that of the other decisions made.

2. Background

The task of VQA has attracted considerable interest since the seminal paper of Antol *et al.* [6], mainly from the Computer Vision community as VQA constitutes a practical setting to evaluate deep visual understanding [25, 26]. The task of VQA is extremely challenging since it also requires the comprehension of a text question and reasoning on the basis of general and common-sense knowledge (see Fig. 1). A number of large-scale datasets for VQA have been created (*e.g.* [6, 12, 22, 38]). The *VQA-real* dataset [6] and its iteration *VQA v2* [12] have served as a benchmark to evaluate and compare methods. The latter associates two images to every question, which are chosen so as to each lead to different answers. This dataset was used for the 2017 VQA challenge [1] and for the experiments presented in this paper.

The prevailing approach to VQA is based on three components. (1) Posing question answering as a classification problem, solved with (2) a deep neural network that implements a joint embedding model, (3) trained end-to-end with supervision of example questions/answers. First, question-answering is posed as a classification over a set of candidate answers. Questions in the current VQA datasets are mostly visual in nature, and the correct answers therefore only span a small set of words and phrases. Practically, correct answers are concentrated in a small set of words and phrases (typically a few hundreds to a few thousands). Second, most VQA models are based on a deep neural network that implements a joint embedding of the image and of the

question. The two inputs are mapped into fixed-size vector representations with convolutional and recurrent neural networks, respectively. Further non-linear mappings of those representations are usually interpreted as projections into a joint “semantic” space. They can then be combined by means of concatenation of element-wise multiplication, before feeding the classifier mentioned above. Third, owing to the success of deep learning on supervised learning problems, this whole neural network is trained end-to-end from questions, images, and their ground truth answers.

The majority of recent methods for VQA have built on the basic joint embedding approach (see [33] for a survey). Interestingly, recent studies have shown that very simple models can achieve strong performance [16, 20] given a careful implementation and/or selection of hyperparameters. Our work follows a similar line. Our extensive experiments show that a some key choices in implementation (*e.g.* gated activations, regression output, smart shuffling, *etc.*) dramatically improve the performance of a relatively simple model.

3. Proposed Model

This section presents the proposed model, based on a deep neural network. For the sake of transparency, we pragmatically describe the model with the specific choices and values of hyperparameters that lead to its best performance. Section 4 will examine variations of architecture and hyperparameters and their influence on performance.

The complete model is summarized in Fig. 2. As a one-sentence summary, it implements the well-known joint RN/CNN embedding of the question/image, with question-guided attention over the image [33, 35, 16, 20].

3.1. Question Embedding

The input for each instance – whether during training or test time – is a text question and an image. The question is tokenized, *i.e.* first split into words using spaces and punctuation. Any number or number-based word (*e.g.* *10,000* or *2:15pm*) is also considered as a word. Questions are trimmed to a maximum of 14 words for computational efficiency. The extra words are then simply discarded, but only about 0.25% of questions in the dataset are longer than 14 words. Each word is turned into a vector representation with a look-up table, whose entries are 300-dimensional vectors learned along other parameters during training. Those vectors are however initialized with pretrained *GloVe* word embeddings [28] (Global Vectors for Word Representation). We use the publicly available version of *GloVe* pretrained on the Wikipedia/Gigaword corpus [27]. The words not present in the pretrained word embedding are initialized with vectors of zeros (subsequently optimized during training). The questions shorter than 14 words are end-padded with vectors of zeros (frozen during training). The resulting se-

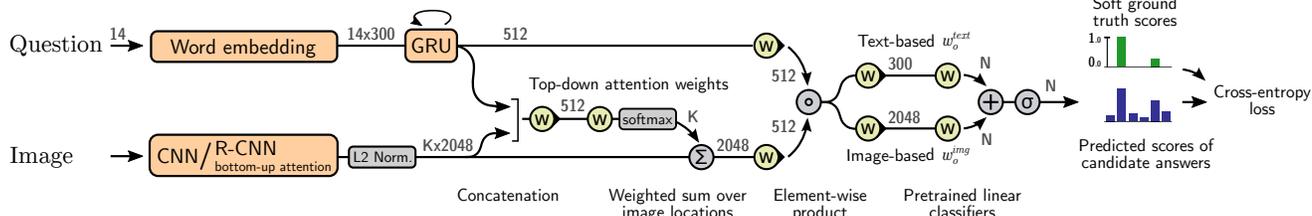


Figure 2. Overview of the proposed model. A deep neural network implements a joint embedding of the input question and image, followed by a multi-label classifier over a fixed set of candidate answers. Gray numbers indicate the dimensions of the vector representations between layers. Yellow elements use learned parameters. The elements \textcircled{W} represent linear layers, and \textcircled{W} non-linear layers (gated tanh).

quence of word embeddings is of size 14×300 and it is passed through a Recurrent Gated Unit (GRU [9]). The recurrent unit has an internal state of dimension 512, and we use its final state, *i.e.* after processing the 14 word embeddings, as our question embedding q . Note that we do not use sentinel `start` or `end` tokens, nor do we trim sequences, or process the strict number of tokens in the given sentence (also known as per-example dynamic unrolling in [20], or TrimZero in [21]). We rather found it more effective to always run the recurrent units for the same number of iterations, including entries containing zero-padding.

3.2. Image Features

The input image is passed through a Convolutional Neural Network (CNN) to obtain a vector representation of size $K \times 2048$, where K is a number of image locations. Each location is thus represented by a 2048-dimensional vector that encodes the appearance of the image in that region. Our evaluation in Section 4 compares two main options with different trade-offs in commodity and performance: a standard pretrained CNN, or a better-performing option. The first, lower-performance option is a 200-layer ResNet (Residual Network [13]) pretrained on ImageNet and publicly available [14]. This gives feature maps of size 14×14 that we resize by average pooling (*i.e.* bilinear interpolation) to 7×7 (*i.e.* $K=49$). The second, higher-performance option is to use the method proposed in [3] which provides image features using *bottom-up attention*. The method is based on a ResNet CNN within a Faster R-CNN framework [29]. It is trained to focus on specific elements in the given image, using annotations from the *Visual Genome* dataset [22]. The resulting features can be interpreted as ResNet features centered on the top- K objects in the image. Our experiments evaluate both a fixed $K=36$, and an adaptive K that uses a fixed threshold for the detected elements in the image, allowing the number of regions K to vary with the complexity of each image, up to a maximum of 100. The images used by the VQA v2 dataset yield in that case an average of about $K=60$ per image.

In all cases, the CNN is pretrained and held fixed during the training of the VQA model. The features can therefore be extracted from the input images as a preprocessing step

for efficiency.

3.3. Image Attention

Our model implements a classical question-guided attention mechanism common to most modern VQA models (see *e.g.* [38, 34, 8, 17, 5, 35]). We refer to this stage as the *top-down* attention, as opposed to the model of [3] that provides image features from bottom-up attention.

For each location $i = 1 \dots K$ in the image, the feature vector v_i is concatenated with the question embedding q (see Fig. 2). They are both passed through a non-linear layer f_a (see Section 3.7) and a linear layer to obtain a scalar attention weight $\alpha_{i,t}$ associated with that location:

$$a_i = w_a f_a([v_i, q]) \quad (1)$$

$$\alpha = \text{softmax}(a) \quad (2)$$

$$\hat{v} = \sum_{i=1}^K \alpha_i v_i \quad (3)$$

where w_a is a learned parameter vector. The attention weights are normalized over all locations with a softmax function (Eq. 2). The image features from all locations are then weighted by the normalized values and summed (Eq. 3) to obtain a single 2048-sized vector \hat{v} representing the attended image.

Note that this attention mechanism is a simple one-glimpse, one-way attention, as opposed to more complex schemes of recent models (*e.g.* stacked, multi-headed, or bidirectional attention [35, 16, 20, 24]).

3.4. Multimodal Fusion

The representations of the question (q) and of the image (\hat{v}) are passed through non-linear layers and then combined with a simple Hadamard product (*i.e.* element-wise multiplication):

$$h = f_q(q) \circ f_v(\hat{v}) \quad (4)$$

The resulting vector h is referred to as the joint embedding of the question and of the image, and is then fed to the output classifier.

3.5. Output Classifier

A set of candidate answers, that we refer to as the output vocabulary, is predetermined from all the correct an-

swers in the training set that appear more than 8 times. This amounts to $N=3129$ candidate answers. We treat VQA as a multi-label classification task. Indeed, each training question in the VQA v2 dataset is associated with one or several answers, each labeled with soft accuracies in $[0, 1]$. Multiple answers and accuracies in $(0, 1)$ arise in case of disagreement between human annotators, particularly with ambiguous questions and multiple or synonymous correct answers [12]. Moreover, in our case, some training questions (about 7%) have no correct answer within the selected output vocabulary. Those questions are not discarded however. They provide a useful training signal by driving towards zero the scores predicted for all candidates of the output vocabulary, although the resulting benefit is marginal (see Section 4.1).

Our multi-label classifier passes the joint embedding \mathbf{h} through a non-linear layer f_o then through a linear mapping \mathbf{w}_o to predict a score \hat{s} for each of the N candidates:

$$\hat{s} = \sigma(\mathbf{w}_o f_o(\mathbf{h})) \quad (5)$$

where σ is a sigmoid (logistic) activation function, and $\mathbf{w}_o \in \mathbb{R}^{N \times 512}$ is a learned weight matrix initialized as described below.

The sigmoid normalizes the final scores to $(0, 1)$, which are followed by a loss similar to a binary cross-entropy, although we use *soft* target scores. This final stage can be seen as a logistic regression that predicts the correctness of each candidate answer. Our objective function is

$$L = - \sum_i^M \sum_j^N s_{ij} \log(\hat{s}_{ij}) - (1-s_{ij}) \log(1-\hat{s}_{ij}) \quad (6)$$

where the indices i and j run respectively over the M training questions and N candidate answers. The ground-truth scores s are the aforementioned soft accuracies of ground truth answers. The above formulation proved to be much more effective than a softmax classifier as commonly used in other VQA models. The advantage of the above formulation is two-fold. First, the sigmoid outputs allow optimization for multiple correct answers per question [16, 32] as is occasionally the case in the VQA v2 dataset. Second, the use of soft scores as targets provides a slightly richer training signal than binary targets, as they capture the occasional uncertainty in ground truth annotations.

3.6. Pretraining the Classifier

In the output stage described above (Eq. 5), the score of a candidate answer j is effectively determined by a dot product between our joint image-question representation $f_o(\mathbf{h})$ and the j th row of \mathbf{w}_o . During training, an appropriate representation for each candidate answer is thus learned as a row of \mathbf{w}_o . We propose to use prior information about the candidate answers from two sources to initialize the rows of \mathbf{w}_o . On the one hand, we use linguistic information in the

form of the GloVe word embeddings of the answer word (as described above for *Question embedding*). When the answer cannot be matched exactly with a pretrained embedding, we use the closest match after spell checking, removing hyphenation, or keeping a single term from multi-word expressions. The corresponding vectors are placed in the matrix $\mathbf{w}_o^{\text{text}}$.

We also exploit visual information gathered from images representing the candidate answers. We use *Google Images* to automatically retrieve 10 photographs associated with each candidate answer. Those photographs are passed through a ResNet-101 CNN pretrained on ImageNet [13]. The final mean-pooled features are extracted and averaged over the 10 photographs. The resulting 2048-sized vector of each candidate answer is placed in the corresponding row of a matrix $\mathbf{w}_o^{\text{img}}$. Those visual representations are complementary to the linguistic ones obtained through word embeddings. They can also be obtained for any candidate answer, including multi-word expressions and rare words for which no word embeddings are available. On the downside, abstract words and expressions may not lead to informative visual representations (see Section 4.6 and Fig. 5).

We combine the prior representations $\mathbf{w}_o^{\text{text}}$ and $\mathbf{w}_o^{\text{img}}$ as follows, decomposing Eq. 5 into

$$\hat{s} = \sigma(\mathbf{w}_o^{\text{text}} f_o^{\text{text}}(\mathbf{h}) + \mathbf{w}_o^{\text{img}} f_o^{\text{img}}(\mathbf{h})) \quad (7)$$

where the non-linear transformations f_o^{text} and f_o^{img} bring \mathbf{h} to the appropriate dimensions, *i.e.* 300 and 2048 respectively (see Fig. 2). The matrices $\mathbf{w}_o^{\text{text}}$ and $\mathbf{w}_o^{\text{img}}$ are fine-tuned with the remainder of the network using smaller relative learning rates, respectively 0.5 and 0.01 (determined through cross-validation).

3.7. Non-linear Layers

The network described above uses multiple learned non-linear layers (see Fig. 2). A common implementation for such a layer would be an affine transformation followed by a Rectified Linear Unit (ReLU). In our implementation, each non-linear layer uses a gated hyperbolic tangent activation. That is, each of those layers implements a function $f_a : \mathbf{x} \in \mathbb{R}^m \rightarrow \mathbf{y} \in \mathbb{R}^n$ with parameters a defined as follows:

$$\tilde{\mathbf{y}} = \tanh(W\mathbf{x} + \mathbf{b}) \quad (8)$$

$$\mathbf{g} = \sigma(W'\mathbf{x} + \mathbf{b}') \quad (9)$$

$$\mathbf{y} = \tilde{\mathbf{y}} \circ \mathbf{g} \quad (10)$$

where σ is the sigmoid activation function, $W, W' \in \mathbb{R}^{n \times m}$ are learned weights, $\mathbf{b}, \mathbf{b}' \in \mathbb{R}^n$ are learned biases, and \circ is the Hadamard (element-wise) product. The vector \mathbf{g} acts multiplicatively as a gate on the intermediate activation $\tilde{\mathbf{y}}$. That formulation is inspired by similar gating operations within recurrent units such as LSTMs and GRUs [9]. This can also be seen as a special case of highway networks [30] and has been mentioned in other work in natural language

Reference model	VQA v2 validation				Accuracy over balanced pairs
	VQA Score			Other	
	All	Yes/no	Numbers		
(1) Removing initialization of $\mathbf{w}_o^{\text{ext}}$ (randomly shuffled)	63.15 \pm 0.08	80.07	42.87	55.81	34.66
(2) Removing initialization of $\mathbf{w}_o^{\text{img}}$ (randomly shuffled)	62.67 \pm 0.07	79.75	42.64	55.14	34.22
(3) Removing image features from bottom-up attention [3]; 7 \times 7 ResNet-200 features instead	58.90 \pm 0.10	77.16	37.18	50.92	29.39
(4) Removing extra training data from the Visual Genome, only uses the VQA v2 training set	57.84 \pm 0.05	77.46	36.73	48.68	28.66
(5) Removing shuffling "by balanced pairs", standard random shuffling instead	57.68 \pm 0.14	77.25	36.46	48.58	28.53
(6) Removing GloVe embeddings to encode the question; learned from random init. instead	56.97 \pm 0.22	75.92	36.03	48.26	27.50
(7a) Replace gated tanh layers with tanh activations	52.13 \pm 7.33	75.90	34.51	38.92	24.78
(7b) Replace gated tanh layers with ReLU activations	55.38 \pm 0.05	74.65	34.83	46.33	25.24
(8a) Binary ground truth targets $s'_{ij} = (s_{ij} \stackrel{?}{>} 0.0)$; ReLU	54.34 \pm 0.19	73.81	34.50	44.95	23.75
(8b) Binary ground truth targets $s'_{ij} = (s_{ij} \stackrel{?}{=} 1.0)$; ReLU	53.41 \pm 0.20	73.23	34.38	43.54	24.32
(9) Replace sigmoid output with softmax; single binary target $s'_{ij} = (s_{ij} \stackrel{?}{=} 1.0)$; ReLU	52.52 \pm 0.31	72.79	34.33	42.09	23.81

Table 1. Cumulative ablations of a single network, evaluated on the VQA v2 validation set. The ablations of table rows are cumulative from top to bottom. The experimental setup is identical to the one used for Table 2.

processing [10, 31].

3.8. Training

We train the network using stochastic gradient descent. We use the AdaDelta algorithm [36], which does not require fixing learning rates and is very insensitive to the initialization of the parameters. The model is prone to overfitting, which we prevent by early stopping as follows. We first train leaving out the official validation set of the VQA v2 dataset for monitoring, and identify the epoch yielding the best performance (highest overall VQA score). The training is then repeated for the same number of epochs, now also using the validation set as training data (as in [12]).

We use questions/answers from the Visual Genome [22] as additional training data. We only use questions whose correct answers overlap the output vocabulary determined on the VQA v2 dataset. This amounts to only about 30% or 485,000 questions from the Visual Genome.

During training with stochastic gradient descent, we enforce the shuffling of training instance to keep *balanced pairs* of VQA v2 in the same mini-batches. Those pairs correspond to identical questions with different images and answers. Our intuitive motivation is that such pairs likely lead to gradients pulling the network parameters in different directions. Keeping an example and its balanced counter-part in a same mini-batch is expected to make the learning more stable, and encourage the network to discern the subtle differences between the paired instances [31].

4. Ablative Experiments

We present an extensive set of experiments that compare our model as presented above (referred to as the *reference* model) with alternative architecture and hyperparameter values. The objective is to show that the proposed model corresponds to a local optimum in the space of architectures and parameters, and to evaluate the sensitivity

of the final performance to each design choice. The following discussion follows the structure of Tables 2 and 1. Note the significant breadth and exhaustivity of the following experiments, which represent more than 3,000 GPU-hours of training time.

Experimental setup Each experiment in this section uses a single network (*i.e.* no ensemble) that is a variation of our *reference* model (first row of Table 2). Each network is trained on the official training set of VQA v2 and on the additional questions from the Visual Genome unless specified. Results are reported on the validation test VQA v2 at the best epoch (highest overall VQA score). Each experiment (*i.e.* each row of Tables 2 and 1) is repeated 3 times, training the same network with different random seeds. We report the average and standard deviation over those three runs. The main performance metric is the standard VQA accuracy [6], *i.e.* the average ground truth score of the answers predicted for all questions¹. We additionally report the metric of *Accuracy over pairs* [37] (last column of Tables 2 and 1). It is the ratio of balanced questions (*i.e.* questions associated with two images leading to two different answers) that are answered perfectly, *i.e.* with both predicted answers having a ground truth score of 1.0. This metric is significantly harder than the standard per-question score since it requires a correct answer to both images of the pair, discouraging blind guesses and reliance on language priors [12, 37].

We now discuss the results of each ablative experiment from Tables 2 and 1 in turn.

4.1. Training Data

The use of additional training questions/answers from the Visual Genome (VG) [22] increase the performance on

¹The ground truth score of a candidate answer j to a question i is a value $s_{ij} \in [0, 1]$ provided with the dataset. It accounts for possible disagreement between annotators: $s_{ij} = 1.0$ if provided by $m \geq 3$ annotators, and $s = m/3$ otherwise. Those scores are further averaged in a 10-choose-9 manner [6].

	VQA v2 validation				Accuracy over balanced pairs
	VQA Score			Other	
	All	Yes/no	Numbers		
Reference model	63.15 ±0.08	80.07	42.87	55.81	34.66
Training data (reference: with Visual Genome data; shuffling keeps balanced pairs in the same minibatches)					
Without extra training data from Visual Genome	62.48 ±0.15	80.37	42.06	54.44	34.12
Random shuffling of training data	63.16 ±0.06	80.17	42.92	55.73	34.54
Discard training questions without answers of score=1.0	63.15 ±0.12	80.15	42.90	55.73	34.62
Question embedding (reference: 300-dimensional GloVe word embeddings, 1-layer forward GRU)					
Word embeddings learned from random initialization; 1-layer forward GRU	62.28 ±0.06	78.77	41.92	55.27	33.67
100-dimensional GloVe; forward GRU	62.45 ±0.05	78.84	42.12	55.51	33.70
200-dimensional GloVe; forward GRU	62.96 ±0.12	79.76	42.49	55.75	34.34
300-dimensional GloVe; bag-of-words (sum)	62.14 ±0.04	79.54	41.88	54.41	33.73
300-dimensional GloVe; bag-of-words (average)	62.53 ±0.09	80.17	42.15	54.67	34.22
300-dimensional GloVe; 1-layer backward GRU	62.82 ±0.02	79.57	42.39	55.64	33.98
300-dimensional GloVe; 2-layer forward GRU	62.29 ±0.10	78.66	42.42	55.23	33.56
300-dimensional GloVe randomly shuffled; 1-layer forward GRU	62.16 ±0.08	78.74	41.73	55.11	33.45
Image features (reference: image features from bottom-up attention, adaptive K)					
ResNet-200 global features ($K=1$, i.e. without image attention)	56.16 ±0.14	75.89	36.39	46.57	25.93
ResNet-200 features 14×14 ($K=196$)	57.93 ±0.25	76.17	36.25	49.95	27.64
ResNet-200 features downsampled to 7×7 ($K=49$)	59.35 ±0.10	77.34	37.74	51.55	29.49
Image features from bottom-up attention, $K=36$	62.82 ±0.14	79.92	42.44	55.35	34.18
Image attention (reference: image features from bottom-up attention, 1 attention head, softmax normalization)					
ResNet-200 7×7 features, 1 head, sigmoid normalization	58.96 ±0.37	77.05	37.99	50.90	28.68
ResNet-200 7×7 features, 1 head, softmax normalization	59.35 ±0.10	77.34	37.74	51.55	29.49
ResNet-200 7×7 features, 2 heads, sigmoid normalization	58.70 ±0.30	76.42	37.95	50.87	28.40
ResNet-200 7×7 features, 2 heads, softmax normalization	59.20 ±0.13	76.97	37.52	51.58	29.43
Image features from bottom-up attention, 1 head, sigmoid normalization	62.15 ±1.41	78.83	43.44	54.57	32.98
Image features from bottom-up attention, 2 heads, sigmoid normalization	62.91 ±0.26	79.74	43.99	55.27	34.30
Image features from bottom-up attention, 2 heads, softmax normalization	63.10 ±0.05	79.80	43.17	55.82	34.52
Output vocabulary (reference: >8 occurrences as correct answers in the training data, $N = 3, 129$)					
Keep answers with >6 training occurrences, $N = 3, 793$	63.26 ±0.07	80.23	42.94	55.89	34.50
Keep answers with >10 training occurrences, $N = 2, 748$	* 63.30 ±0.06	80.33	43.02	55.86	34.69
Keep answers with >12 training occurrences, $N = 2, 418$	63.16 ±0.08	80.18	43.17	55.66	34.56
Keep answers with >14 training occurrences, $N = 2, 160$	63.27 ±0.08	80.45	43.07	55.69	34.83
Keep answers with >16 training occurrences, $N = 1, 961$	63.17 ±0.10	80.29	42.94	55.65	34.62
Keep answers with >20 training occurrences, $N = 1, 656$	63.14 ±0.05	80.21	42.90	55.66	34.67
Keep answers with >200 training occurrences, $N = 278$	59.59 ±0.04	79.90	42.11	48.93	32.64
Output classifier (reference: sigmoid output, w_o^{text} and w_o^{img} pretrained)					
Softmax output, w_o^{text} pretrained, w_o^{img} pretrained	60.47 ±0.08	78.19	40.93	52.32	32.27
Sigmoid output, w_o^{text} randomly initialized, w_o^{img} randomly initialized	62.28 ±0.99	78.50	42.50	55.32	33.14
Sigmoid output, w_o^{text} randomly shuffled, w_o^{img} randomly shuffled	62.64 ±0.15	79.82	42.26	55.13	34.17
Sigmoid output, w_o^{text} pretrained, w_o^{img} randomly initialized	62.94 ±0.14	79.99	42.72	55.47	34.26
Sigmoid output, w_o^{text} pretrained, w_o^{img} randomly shuffled	63.04 ±0.02	80.20	42.83	55.50	34.66
Sigmoid output, w_o^{text} randomly initialized, w_o^{img} pretrained	63.21 ±0.01	80.06	42.97	55.90	34.65
Sigmoid output, w_o^{text} randomly shuffled, w_o^{img} pretrained	62.97 ±0.06	79.83	42.62	55.68	34.31
General architecture (reference: gated tanh non-linear activations, hidden states of dimension 512)					
Non-linear activations: ReLU	61.63 ±0.21	78.48	41.15	54.39	32.57
Non-linear activations: tanh	61.74 ±0.44	79.20	40.92	54.13	33.10
Non-linear activations: gated ReLU	62.44 ±0.07	79.33	42.12	55.13	33.52
Hidden states of dimension 256	62.80 ±0.05	79.62	42.46	55.53	34.12
Hidden states of dimension 384	63.12 ±0.11	80.06	42.89	55.74	34.54
Hidden states of dimension 768	63.12 ±0.11	80.06	42.89	55.74	34.54
Hidden states of dimension 1024	63.02 ±0.55	79.88	42.68	55.73	34.34
Hidden states of dimension 1280	* 63.37 ±0.21	80.40	43.02	55.96	34.76
Mini-batch size (reference: 512 training questions)					
128 Training questions	62.38 ±0.08	79.70	42.29	54.67	34.01
256 Training questions	* 63.17 ±0.09	80.22	42.94	55.72	34.71
384 Training questions	* 63.20 ±0.04	80.21	43.08	55.75	34.61
768 Training questions	62.99 ±0.12	79.84	42.61	55.73	34.40

Table 2. Ablations of a single network, evaluated on the VQA v2 validation set. We evaluate variations of our best “reference” model (first row), and show that it corresponds to a local optimum in the space of architectures and hyperparameters. Every row corresponds to one variation of that reference model. We train each variation with three different random seeds and report averages and standard deviations (\pm).

VQA v2 [12] in all question types. As mentioned above, we only use VG instances with a correct answer appearing the output vocabulary determined on VQA v2, and which use an image also used in VQA v2. Note that the +0.67% increase in performance is modest relative to the amount of additional training questions (an additional 485,000 over the 650,000 of VQA v2). Note that including VG questions relating to MS COCO images not used in VQA v2 resulted in slightly lower final performance (not reported in Table 2). We did not further investigate this issue.

We compare the proposed **shuffling of training data** which keeps balanced pairs in the same mini-batches, with a standard, arbitrary random shuffling. The results in Table 2 are inconclusive: the overall VQA score is virtually identical either way. The accuracy over pairs however improves with the proposed shuffling. This is to be expected, as the purpose of the proposed method is to improve the learning of differentiating between balanced pairs. The cumulative ablation (row (5) in Table 1) confirms this advantage more clearly.

We evaluate discarding the training questions that do not have their ground truth answer within the selected candidates. Early experiments have shown that those instances do still carry a useful training signal by drawing the predicted scores of the selected candidate answers towards zero. In Table 2, the VQA score remains virtually identical, but a very slight benefit is observed on the accuracy over pairs by retaining those instances.

4.2. Question Embedding

Our reference model uses pretrained GloVe word embeddings of dimension 300 followed by a one-layer GRU processing words in forward order. We compare this choice to a series of more simple and more advanced options. Learning the word embeddings from scratch, *i.e.* from random initializations reduces performance by 0.87%. The gap with pretrained embeddings is even greater as the model is trained on less training data (Fig. 3 and Section A.2). This is consistent with findings previously reported in [32]. This experiment shows, on the one hand, a benefit of leveraging non-VQA training data. On the other hand, it suggests that a sufficiently large VQA training set may remove this benefit altogether. Using GloVe vectors of smaller dimension (100 or 200) also give lower performance than those in the reference model (300).

We investigate whether the pretrained word embeddings (GloVe vectors) really capture word-specific information. The alternative, null hypothesis is that the simple spreading of vectors in the word embedding space is a benefit in itself. To test this, we randomly shuffle the same GloVe vectors, *i.e.* we associate them with words from the input dictionary chosen at random. The shuffled GloVe vectors perform even worse than word embeddings learned from scratch. This

shows that GloVe vectors indeed capture word-specific information.

Replacing our GRU with advanced options (backward, bidirectional, or two-layer GRU) gives lower performance. Simpler options (bag-of-words, simply summing or averaging word embeddings) also give lower performance but are still a surprisingly strong baseline, as previously reported in [16, 32].

4.3. Image Features

Our best model uses the **images features from bottom-up attention** of [3]. These are obtained through a Faster R-CNN framework and an underlying 101-layer ResNet that focuses on specific image regions. The method uses a fixed threshold on object detections, and the number of features K is therefore adaptive to the contents of the image. It is capped at a maximum of 100, and yields an average in the order of $K=60$. We experiment with a fixed number of features $K=36$. The performance degrades only slightly. This option may be a reasonable alternative considering the lower implementation and computational costs.

We also experiment with mainstream image features from a 200-layer ResNet [13, 14]. As per the common practice, we use the last feature map of the network, of dimensions 14×14 and 2,048 channels. The performance with ResNet features drops dramatically from 63.15% to 57.52%. A global average pooling of those features, *i.e.* collapsing to a 1×1 map and discarding the attention mechanism, is expectedly even worse. A surprising finding however is that **coarser 7×7 ResNet feature maps** give a reasonable performance of 59.24. Those are obtained by linear interpolation of the 14×14 maps, equivalent to a 2×2 average pooling of stride 2. As an explanation, we hypothesize that the resolution of the coarser maps better correspond to the scale of objects in our images.

Note that the proposed model was initially developed with standard ResNet features, and is not specifically optimized for the features of [3]. On the contrary, we found that optimal choices were stable across types of image features (including for the attention mechanism, see Section 4.4). In all cases, we also observed that an $L2$ normalization of the image features was crucial for good performance – at least with our choices of architecture and optimizer.

4.4. Image Attention

Our reference model uses a single set of K attention weights normalized with a softmax. Some previous studies have reported better performance with multiple sets of attention weights (also known as multiple “glimpses” or “attentions heads”) [11] and/or with a sigmoid normalization over the weights [31] (which allows multiple glimpses in *one* pass). In our case, none of these options proved beneficial, whether with ResNet features or with features from bottom-up attention. This confirms that the latter image fea-

tures are suited to a simple drop-in replacement in lieu of traditional feature maps, as argued in [3].

4.5. Output Vocabulary

We determine the output vocabulary, *i.e.* the set of candidate answers from those appearing in the training set more than ℓ times. We cross-validate this threshold and find a relatively broad optimum around $\ell = 8\text{--}12$ occurrences. This corresponds to about $N=2,400$ to $3,800$ candidate answers, which is of the same order as the values typically reported in the literature (without cross-validation) of $2,000$ or $3,000$. Note however a higher ℓ (*i.e.* lower N) still gives reasonable performance with a lower number of parameters and computational complexity.

4.6. Output classifier

Our reference model uses sigmoid outputs and the soft scores s_{ij} as ground truth targets. We first compare the soft scores with two binarized versions $s'_{ij} = (s_{ij} \stackrel{?}{>} 0.0)$, and $s''_{ij} = (s_{ij} \stackrel{?}{=} 1.0)$. The proposed **soft scores perform significantly better** than either of those binarized versions (Table 1).

We then compare the sigmoid outputs of our reference model to the common choice of a softmax. Both use a cross-entropy loss. The softmax uses the single ground truth answer provided in the dataset, whereas the sigmoid uses the complete annotation data, occasionally with multiple correct answers marked for one question, due to disagreement between multiple annotators. The **sigmoid performs significantly better than a softmax**. This observation is confirmed in the cumulative ablations (Table 1)

We now evaluate the proposed pretraining of the parameters w_o^{text} and/or w_o^{img} of the classifiers. We consider two baselines for those two matrices: random initialization and random shuffling. The former simply initializes and learns the weights like any others in the network. The latter uses the proposed pretraining as initializations but assigns then to the candidate answers randomly by shuffling the rows. This should reveal whether the proposed initialization provides answer-specific information, or whether it simply helps the numerical optimization because of better conditioned initial values. The **proposed initialization performs consistently better than those baselines**. The random initialization of w_o^{text} may seem surprisingly good as it even surpasses the reference model on the overall VQA score. It does not, however, on the metric of *Accuracy over pairs*. The experiments with reduced training data (Section A.2 and Fig. 3) confirm even more clearly the benefits of the proposed initialization.

4.7. General Architecture

All of our non-linear layers are implemented as **gated tanh** (Section 3.7). These show a clear benefit over the

gated ReLU, and even more so over simple ReLU or tanh activations. Note that we also experimented, without success, with various other combinations of highway [30], residual [13] and gating connections (not reported in Table 2). One benefit of gated layers is to double the number of learned parameters without increasing the dimension of the hidden states.

We cross-validate the **dimension of hidden states** among the values $\{256, 384, 512, 768, 1024, 1280\}$. We settled on 512 as a reasonable sweet spot. Larger dimensions (*e.g.* 1280) can be better but without guarantees. The variance across repeated experiments is larger, likely due to overfitting and unstable training. We use a simple **element-wise product to combine** the question and image representations. This proved far superior to a concatenation. We did not experiment with advanced forms of pooling [11, 7].

Please consult the supplementary material for additional experiments including a comparison with competing methods on the VQA v2 benchmark.

5. Discussion and Conclusions

This paper presented a model for VQA based on a deep neural network that significantly outperforms all other approaches proposed to date. Importantly, we reported an extensive suite of experiments that identify the contribution of each design choice and the performance of alternative designs. The general take-away from this study is that the performance is very dependent on design choices and on various details of the implementation. We attribute the success of our model to a number of points. Some are seemingly minor and easily implemented (*e.g.* large mini-batch size, sigmoid output) while others are clearly non-trivial (*e.g.* gated non-linear activations, image features from bottom-up attention, pretraining the output classifier).

This paper does not claim to make breakthrough advances in the field of VQA, which remains a largely unsolved problem. We hope that our model may however serve as a solid basis on which to make future progress. Our extensive analysis and exploration of designs is unprecedented in scale for VQA, and is a significant contribution to the field. It provides indicators on the importance of the various components of a VQA model. It also allows us to point at promising directions for future developments.

Our evaluation of simple baselines have shown surprisingly strong performances. For example, encoding questions as a simple bag-of-words performs almost as well as state-of-the-art recurrent encoders. It suggests that the word ordering in questions may not convey much information – which is plausible for the most basic ones – or, more realistically, that our current models are still unable to understand and make effective use of language structure. Recent works on compositional models for VQA are a promising direction to address this issue [4, 15, 18, 19].

References

- [1] VQA Challenge leaderboard. <http://visualqa.org> and <http://evalai.cloudev.org>. 2, 12, 13
- [2] L. A. Adamic and B. A. Huberman. Zipf's law and the internet. *Glottometrics*, 3(1):143–150, 2002. 11
- [3] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. Bottom-up and top-down attention for image captioning and vqa. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2018. 2, 3, 5, 7, 8, 12
- [4] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein. Learning to compose neural networks for question answering. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2016. 8
- [5] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein. Neural Module Networks. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 3
- [6] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. VQA: Visual Question Answering. In *Proc. IEEE Int. Conf. Comp. Vis.*, 2015. 2, 5
- [7] H. Ben-younes, R. Cadène, M. Cord, and N. Thome. MUTAN: multimodal tucker fusion for visual question answering. In *Proc. IEEE Int. Conf. Comp. Vis.*, 2017. 8, 13
- [8] K. Chen, J. Wang, L.-C. Chen, H. Gao, W. Xu, and R. Nevatia. ABC-CNN: An Attention Based Convolutional Neural Network for Visual Question Answering. *arXiv preprint arXiv:1511.05960*, 2015. 3
- [9] K. Cho, B. van Merriënboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proc. Conf. Empirical Methods in Natural Language Processing*, 2014. 3, 4
- [10] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier. Language modeling with gated convolutional networks. *arXiv preprint arXiv:1612.08083*, 2016. 5
- [11] 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. *arXiv preprint arXiv:1606.01847*, 2016. 7, 8, 13
- [12] 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. *arXiv preprint arXiv:1612.00837*, 2016. 1, 2, 4, 5, 7, 13
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 3, 4, 7, 8
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. *arXiv preprint arXiv:1603.05027*, 2016. 3, 7
- [15] R. Hu, J. Andreas, M. Rohrbach, T. Darrell, and K. Saenko. Learning to reason: End-to-end module networks for visual question answering. *arXiv preprint arXiv:1704.05526*, 2017. 8
- [16] A. Jabri, A. Joulin, and L. van der Maaten. Revisiting visual question answering baselines. 2016. 2, 3, 4, 7
- [17] A. Jiang, F. Wang, F. Porikli, and Y. Li. Compositional Memory for Visual Question Answering. *arXiv preprint arXiv:1511.05676*, 2015. 3
- [18] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. B. Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. *arXiv preprint arXiv:1612.06890*, 2016. 8
- [19] J. Johnson, B. Hariharan, L. van der Maaten, J. Hoffman, F. Li, C. L. Zitnick, and R. B. Girshick. Inferring and executing programs for visual reasoning. *arXiv preprint arXiv:1705.03633*, 2017. 8
- [20] V. Kazemi and A. Elqursh. Show, ask, attend, and answer: A strong baseline for visual question answering. *arXiv preprint arXiv:1704.03162*, 2017. 1, 2, 3
- [21] J.-H. Kim, J. Kim, J.-W. Ha, and B.-T. Zhang. TrimZero: A Torch Recurrent Module for Efficient Natural Language Processing. In *Proceedings of KIIS Spring Conference*, volume 26, pages 165–166, 2016. 3
- [22] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. Bernstein, and L. Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332*, 2016. 2, 3, 5
- [23] J. Lu, X. Lin, D. Batra, and D. Parikh. Deeper lstm and normalized cnn visual question answering model. https://github.com/VT-vision-lab/VQA_LSTM_CNN, 2015. 13
- [24] J. Lu, J. Yang, D. Batra, and D. Parikh. Hierarchical question-image co-attention for visual question answering. *arXiv preprint arXiv:1606.00061*, 2016. 3
- [25] M. Malinowski and M. Fritz. Towards a Visual Turing Challenge. *arXiv preprint arXiv:1410.8027*, 2014. 2
- [26] M. Malinowski, M. Rohrbach, and M. Fritz. Ask Your Neurons: A Neural-based Approach to Answering Questions about Images. In *Proc. IEEE Int. Conf. Comp. Vis.*, 2015. 2
- [27] J. Pennington, R. Socher, and C. Manning. Glove: Global vectors for word representation. <http://nlp.stanford.edu/projects/glove/>. 2
- [28] J. Pennington, R. Socher, and C. Manning. Glove: Global Vectors for Word Representation. In *Conference on Empirical Methods in Natural Language Processing*, 2014. 2
- [29] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proc. Advances in Neural Inf. Process. Syst.* 2015. 3
- [30] R. K. Srivastava, K. Greff, and J. Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387v1*, 2015. 4, 8
- [31] D. Teney, L. Liu, and A. van den Hengel. Graph-structured representations for visual question answering. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 5, 7
- [32] D. Teney and A. van den Hengel. Zero-shot visual question answering. 2016. 4, 7, 11, 12
- [33] Q. Wu, D. Teney, P. Wang, C. Shen, A. Dick, and A. van den Hengel. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 2017. 1, 2
- [34] H. Xu and K. Saenko. Ask, Attend and Answer: Exploring Question-Guided Spatial Attention for Visual Question Answering. *arXiv preprint arXiv:1511.05234*, 2015. 3

- [35] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked Attention Networks for Image Question Answering. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 2, 3
- [36] M. D. Zeiler. ADADELTA: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012. 5
- [37] P. Zhang, Y. Goyal, D. Summers-Stay, D. Batra, and D. Parikh. Yin and yang: Balancing and answering binary visual questions. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 5
- [38] Y. Zhu, O. Groth, M. Bernstein, and L. Fei-Fei. Visual7W: Grounded Question Answering in Images. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016. 2, 3