

## Towards VQA Models That Can Read

Amanpreet Singh<sup>1</sup>, Vivek Natarajan, Meet Shah<sup>1</sup>, Yu Jiang<sup>1</sup>, Xinlei Chen<sup>1</sup>,  
 Dhruv Batra<sup>1,2</sup>, Devi Parikh<sup>1,2</sup>, and Marcus Rohrbach<sup>1</sup>

<sup>1</sup>Facebook AI Research, <sup>2</sup>Georgia Institute of Technology

### Abstract

Studies have shown that a dominant class of questions asked by visually impaired users on images of their surroundings involves reading text in the image. But today’s VQA models can not read! Our paper takes a first step towards addressing this problem. First, we introduce a new “TextVQA” dataset to facilitate progress on this important problem. Existing datasets either have a small proportion of questions about text (e.g., the VQA dataset) or are too small (e.g., the VizWiz dataset). TextVQA contains 45,336 questions on 28,408 images that require reasoning about text to answer. Second, we introduce a novel model architecture that reads text in the image, reasons about it in the context of the image and the question, and predicts an answer which might be a deduction based on the text and the image or composed of the strings found in the image. Consequently, we call our approach Look, Read, Reason & Answer (LoRRA). We show that LoRRA outperforms existing state-of-the-art VQA models on our TextVQA dataset. We find that the gap between human performance and machine performance is significantly larger on TextVQA than on VQA 2.0, suggesting that TextVQA is well-suited to benchmark progress along directions complementary to VQA 2.0.

### 1. Introduction

The focus of this paper is endowing Visual Question Answering (VQA) models a new capability – the ability to *read text in images and answer questions* by reasoning over the text and other visual content.

VQA has witnessed tremendous progress. But today’s VQA models fail catastrophically on questions requiring reading!<sup>1</sup> This is ironic because these are *exactly* the questions visually-impaired users frequently ask of their assistive devices. Specifically, the VizWiz study [5] found that up to 21% of these questions involve reading and reasoning about the text captured in the images of a user’s surround-

<sup>1</sup>All top entries in the CVPR VQA Challenges (2016-18) struggle to answer questions in category requiring reading correctly.

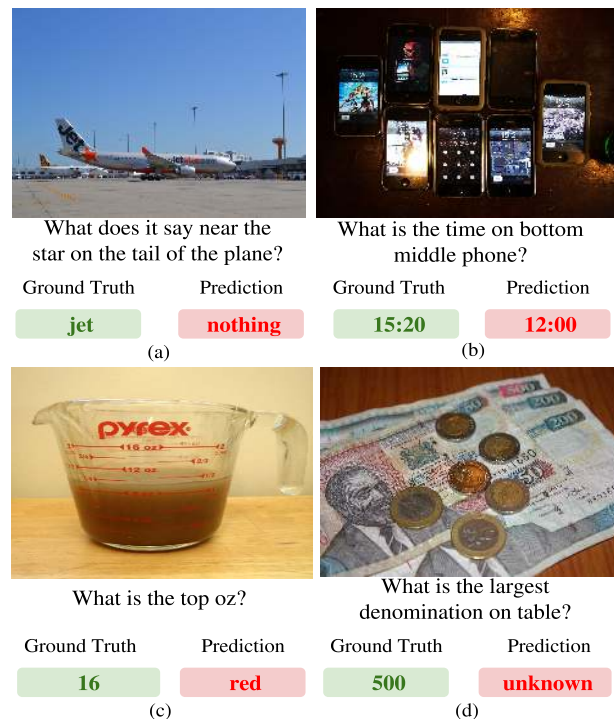


Figure 1: Examples from our TextVQA dataset that require VQA models to understand text embedded in images to answer the questions correctly. Ground truth answers are shown in green and the answers predicted by a state-of-the-art VQA model (Pythia [17]) are shown in red. Clearly, today’s VQA models fail at answering questions that involve reading and reasoning about text in images.

ings – ‘*what temperature is my oven set to?*’, ‘*what denomination is this bill?*’.

Consider the question in Fig. 1(a) – ‘*What does it say near the star on the tail of the plane?*’ from the TextVQA dataset. With a few notable exceptions, today’s state-of-art VQA models are predominantly monolithic deep neural networks (without any specialized components). Consider what we are asking such models to learn to answer this question – the model must learn to

- realize when the question is about text (‘*What ... say?*’),
- detect image regions containing text (‘*15:20*’, ‘*500*’),

- convert pixel representations of these regions (convolutional features) to symbols ('15:20') or textual representations (semantic word-embeddings),
- jointly reason about detected text and visual content, e.g. resolving spatial or other visual reference relations ('tail of the plane ... on the back') to focus on the correct regions.
- finally, decide if the detected text needs to be 'copy-pasted' as the answer (e.g. '16' in Fig. 1 (c)) or if the detected text informs the model about an answer in the answer space (e.g. answering 'jet', in Fig. 1(a))

When laid out like that, it is perhaps unsurprising why today's models have not been able to make progress on reading questions – simply put, despite all the strengths of deep learning, it seems hopelessly implausible that all of the above skills will simply *emerge* in a monolithic network all from the distant supervision of VQA accuracy.

Fortunately, we can do more than just hope. Optical Character Recognition (OCR) is a mature sub-field of computer vision. A key thesis of this work is the following – we should bake in inductive biases and specialized components (e.g. OCR) into models to endow them with the different skills (e.g. reading, reasoning) required by the all-encompassing task of VQA.

Specifically, we propose a new VQA model that includes OCR as a module. We call it *Look, Read, Reason & Answer* (LoRRA). Our model architecture incorporates the regions (bounding boxes) in the image containing text as entities to attend over (in addition to object proposals). It also incorporates the actual text recognized in these regions (e.g. '15:20') as information (in addition to visual features) that the model learns to reason over. Finally, our model includes a mechanism to decide if the answer produced should be 'copied' over from the OCR output (in more of a generation or slot-filling flavor), or should be deduced from the text (as in a standard discriminative prediction paradigm popular among existing VQA models). Our model learns this mechanism end-to-end. While currently limited in scope to OCR, our model is as an initial step towards endowing VQA models with the ability to reason over unstructured sources of external knowledge (in this case text found in a test image) and accommodate multiple streams of information flow (in this case predicting an answer from a pre-determined vocabulary or generating an answer via copy).

One reason why there has been limited progress on VQA models that can read and reason about text in images is because such questions, while being a dominant category in real applications for aiding visually impaired users [5], are infrequent in the standard VQA datasets [3, 10, 51] because they were not collected in settings that mimic those of visually impaired users. While the VizWiz dataset [13] does contain data collected from visually impaired users, the effective size of the dataset is small due to 58% of the questions being "unanswerable". This makes it challenging to

study the problem systematically, train effective models, or even draw sufficient attention to this important skill that current VQA models lack.

To this end, we introduce the TextVQA dataset. It contains 45,336 questions asked by (sighted) humans on 28,408 images from the Open Images dataset [27] from categories that tend to contain text e.g. "billboard", "traffic sign", "whiteboard". Questions in the dataset require reading and reasoning about text in the image. Each question-image pair has 10 ground truth answers provided by humans.

Models that do well on this dataset will not only need to parse the image and the question as in traditional VQA, but also read the text in the image, identify which of the text might be relevant to the question, and recognize whether a subset of the detected text is directly the answer (e.g., in the case of 'what temperature is my oven set to?') or additional reasoning is required on the detected text to answer the question (e.g., 'which team is winning?').

We show that LoRRA outperforms existing state-of-the-art VQA models on the TextVQA dataset. Overall, our contributions are:

- We introduce a novel dataset (TextVQA) containing questions which require the model to read and reason about the text in the image to be answered.
- We propose *Look, Read, Reason & Answer* (LoRRA): a novel model architecture which explicitly reasons over the outputs from an OCR system when answering questions.
- LoRRA outperforms existing state-of-the-art VQA models on our TextVQA dataset.

## 2. Related work

**Visual Question Answering.** VQA has seen numerous advances and new datasets since the first large-scale VQA dataset was introduced by Antol *et al.* [3]. This dataset was larger, more natural, and more varied than earlier VQA datasets such as DAQUAR [31] or COCO-QA [38] but had linguistic priors which were exploited by models to answer questions without sufficient visual grounding. This issue was addressed by Goyal *et al.* [10] by adding complementary triplets  $(I_c, q, a_c)$  for each original triplet  $(I_o, q, a_o)$  where image  $I_c$  is similar to image  $I_o$  but the answer for the given question  $q$  changes from  $a_o$  to  $a_c$ . To study visual reasoning independent of language, non-photo-realistic VQA datasets have been introduced such as CLEVR [18], NLVR [42] and FigureQA [21]. Wang *et al.* [45] introduced a Fact-Based VQA dataset which explicitly requires external knowledge to answer a question.

**Text based VQA.** Several existing datasets study text detection and/or parsing in natural everyday scenes: COCO-Text [43], Street-View text [44] IIIT-5k [33] and ICDAR 2015 [22]. These do not involve answering questions about the images or reasoning about the text. DVQA [20] assesses automatic bar-chart understanding by training models to answer questions about graphs and plots. The Multi-

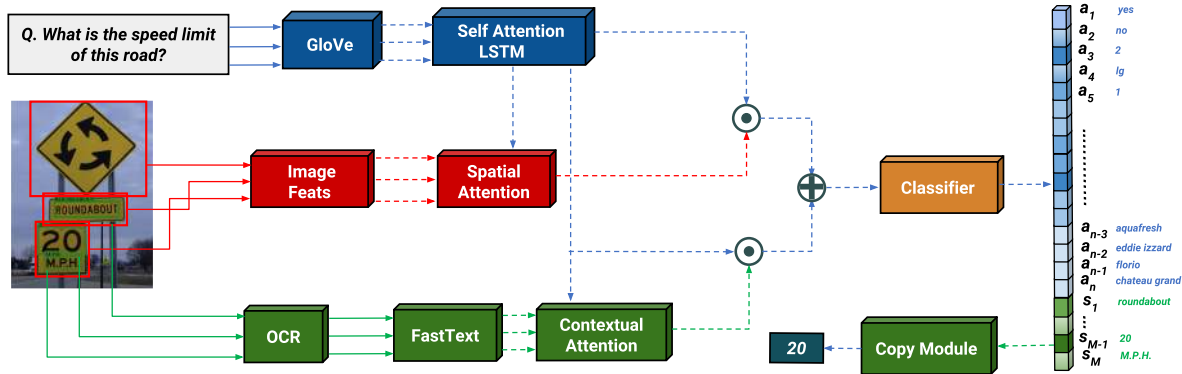


Figure 2: Overview of our approach *Look, Read, Reason & Answer* (LoRRA). Our approach looks at the image, reads its text, reasons about the image and text content and then answers, either with an answer  $a$  from the fixed answer vocabulary or by selecting one of the OCR strings  $s$ . Dashed lines indicated not jointly-trained components. The answer cubes on the right with darker color have more attention weight. The OCR token “20” has the highest attention weight in the example.

Output Model (MOM) introduced in DVQA uses an OCR module to read chart specific content. Textbook QA (TQA) [24] considers the task of answering questions from middle-school textbooks, which often require understanding and reasoning about text and diagrams. Similarly, AI2D [23] contains diagram based multiple-choice questions. Note that these works all require reasoning about text to answer questions, but in narrow domains (bar charts, textbook diagrams, etc.). The focus of our work is to reason and answer questions about text in natural everyday scenes. MemexQA [16] introduces a VQA task which involves reasoning about the time and date at which a photo/video was taken, but this information is structured and is part of the meta data.

**Visual Representations for VQA Models.** VQA models typically use some variant of attention to get a representation of the image that is relevant for answering the question [2, 7, 30, 47, 48, 51, 17]. Object region proposals and the associated features are generated by using a detection network which are then spatially attended to, conditioned on a question representation. In this work, we extend the representations that a VQA model reasons over. Specifically, in addition to attending over object proposals, our model also attends over the regions where text is detected.

**Copy Mechanism.** A core component of our proposed model is its ability to decide whether the answer to a question should be an OCR token detected in the image, or if the OCR tokens should only inform the answer to the question. The former is implemented as a “copy mechanism” – a learned slot filling approach. Our copy mechanism is based on a series of works on pointer generator networks [11, 39, 32, 12, 34]. A copy mechanism provides networks the ability to generate out-of-vocabulary words by pointing at a word in context and then copying it to the generated result. This approach has been used for a variety of tasks in NLP such as summarization [11, 34, 39], question answer-

ing [46], language modelling [32], neural machine translation [12], and dialog [37].

### 3. LoRRA: Look, Read, Reason & Answer

In this section, we introduce our novel model architecture to answer questions about images which require reading the text in the images to answer the questions.

We assume we get an image  $v$  and a question  $q$  as input, where the question consists of  $L$  words  $w_1, w_2, \dots, w_L$ . At a high level, our model contains three components: (i) a **VQA component** to reason and infer the answer based on the image  $v$  and the question  $q$  (Sec 3.3); (ii) a **reading component** which allows our model to read the text in the image (Sec 3.2); and (iii) an **answering module** which either predicts from an answer space or points to the text read by the *reading component* (Sec. 3.3). The overall model is shown in Fig. 2. Note that, the OCR module and backbone VQA model and can be any OCR model and any recent attention-based VQA model. Our approach is agnostic to the internal details of these components. We detail the exact implementation choices and hyper parameters in Sec. 3.4.

#### 3.1. VQA Component

Similar to many VQA models [7, 17], we first embed the question words  $w_1, w_2, \dots, w_L$  of the question  $q$  with a pre-trained embedding function (*e.g.* GloVe [36]) and then encode the resultant word embeddings iteratively with a recurrent network (*e.g.* LSTM [15]) to produce a question embedding  $f_Q(q)$ . For images, the visual features are represented as spatial features, either in the form of grid-based convolutions and/or features extracted from the bounding box proposals [1]. We refer to these features as  $f_I(v)$  where  $f_I$  is the network which extracts the image representation. We use an attention mechanism  $f_A$  over the spatial features [4, 7], which predicts attentions based on the  $f_I(v)$

and  $f_Q(q)$  and gives a weighted average over the spatial features as the output. We then combine the output with the question embedding. At a high level, the calculation of our VQA features  $f_{VQA}(v, q)$  can be written as:

$$f_{VQA}(v, q) = f_{comb}(f_A(f_I(v), f_Q(q)), f_Q(q)) \quad (1)$$

where  $f_{comb}$  is the combination module ( $\otimes$ ) in Fig. 2.

Assuming we have a fixed answer space of  $a_1, \dots, a_N$ , we use a feed-forward MLP  $f_c$  on the combined embedding  $f_{VQA}(v, q)$  to predict probabilities  $p_1, \dots, p_N$  where the probability of  $a_i$  being the correct answer is  $p_i$ .

### 3.2. Reading Component

To add the capability of reading text from an image, we rely on an OCR model which is not jointly trained with our system. We assume that the OCR model can read and return word tokens from an image, e.g. [6, 41]. The OCR model extracts  $M$  words  $s = s_1, s_2, \dots, s_M$  from the image which are then embedded with a pre-trained word embedding,  $f_O$ . Finally, we use the same architecture as VQA component to get combined OCR-question features,  $f_{OCR}$ . Specifically,

$$f_{OCR}(s, q) = f_{comb}(f_A(f_O(s), f_Q(q)), f_Q(q)) \quad (2)$$

This is visualized in Fig. 2. Note that the parameters of the functions  $f_A$  and  $f_{comb}$  are not shared with the VQA model component above but they have the same architecture, just with different input dimensions. During weighted attention as the features are multiplied by weights and then averaged, the ordering information gets lost. To provide the answer module with the ordering information of the original OCR tokens, we concatenate the attention weights and the final weight-averaged features. This allows the answer module to know the original attention weights for each token in order.

### 3.3. Answer Module

With a fixed answer space, the current VQA models are only able to predict fixed tokens which limits the generalization to out-of-vocabulary (OOV) words. As the text in images frequently contains words not seen at training time, it is hard to answer text-based questions based on a pre-defined answer space alone. To generalize to arbitrary text, we take inspiration from pointer networks which allow pointing to OOV words in context [11, 39, 32, 12, 34]. We extend our answer space through addition of a dynamic component which corresponds to  $M$  OCR tokens. The model now has to predict probabilities  $(p_1, \dots, p_N, \dots, p_{N+M})$  for  $N+M$  items in the answer space instead of the original  $N$  items.

We pick the index with the highest probability  $p_i$  as the index of our predicted answer. If the model predicts an index larger than  $N$  (i.e., among the last  $M$  tokens in answer space), we directly “copy” the corresponding OCR token as the predicted answer. Hence, our answering module can

VQA 2.0 Accuracy	
Model	test-dev
BUTD [1]	65.32
Counter [50]	68.09
BAN [25]	69.08
Pythia v0.1 [17]	68.49
Pythia v0.3 (Ours)	68.71
Pythia v0.3 + LoRRA (Ours)	69.21
VizWiz Accuracy	
Model	test
BAN[25]	51.40
Pythia v0.3 (Ours)	54.72

Table 1: **Single model VQA 2.0 and VizWiz performance in %.** Our revised implementation of Pythia, v0.3, with LoRRA outperforms or is comparable to state-of-the-art.

be thought of as “copy if you need” module which allows answering from the OOV words using the OCR tokens.

With all of the components, the final equation  $f_{LoRRA}$  for predicting the answer probabilities can be written as:

$$f_{LoRRA}(v, s, q) = f_{MLP}([f_{VQA}(v, q); f_{OCR}(s, q)]) \quad (3)$$

where  $[\cdot]$  refers to concatenation and  $f_{MLP}$  is a two-layer feed-forward network which predicts the binary probabilities as logits for each answer. We opt for binary cross entropy using logits instead of calculating the probabilities through softmax as it allows us to handle cases where the answer can be in both the actual answer space and the OCR tokens without penalizing for predicting either one (the likelihood of logits is independent of each other). Note that if the model chooses to copy, it can only produce one of the OCR tokens as the predicted answer. 8.9% of the TextVQA questions can be answered only by combining multiple OCR tokens; we leave this as future work.


### 3.4. Implementation Details

Our VQA component is based on the VQA 2018 challenge winner entry, Pythia v0.1 [17]. Our revised implementation, Pythia v0.3, with slight changes in hyperparameters (e.g. size of question vocabulary, hidden dimensions) achieves state-of-the-art VQA accuracy for a single model (i.e. w/o ensemble) as shown in Tab. 1 on both VQA v2.0 dataset [9] and VizWiz dataset [13]. The revised design choices are discussed in [40].

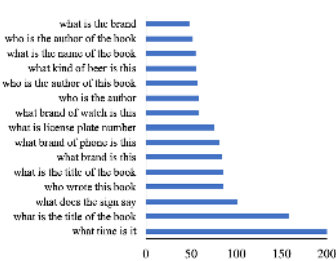
Pythia [17, 40] is inspired from the detector-based bounding box prediction of the bottom-up top-down attention network [1] (VQA winner 2017), which in turn has a multi-modal attention mechanism similar to the VQA 2016 winner [7], which relied on grid-based features.

In Pythia, for spatial features  $f_I(v)$ , we rely on both grid and region based features for an image. The grid based features are obtained by average pooling 2048D features







(a) Number of questions with a particular question length. We see that average question length (7.16) is higher in TextVQA compared to others. (b) Number of majority answers with a particular length. Average answer length (1.7) is high and answer can contain long paragraph and quotes. (c) Number of images with a particular number of OCR tokens. Average number of tokens is around 3.14. In TextVQA, 10x more images contain OCR text than others.



(d) Text of top 15 most occurring questions in TextVQA. Most of the top questions start with “what”.



(e) Total occurrences for 500 most common questions among 23184 unique questions with markers for particular ranks.



(f) Similar to 6e, plot shows total occurrences for 500 most common majority answers with markers for particular ranks.

Figure 6: Question, Answer and OCR statistics for TextVQA. We show comparisons with VQA 2.0 [10] and VizWiz [13].

gories that tend to have images with text in them, we select 100 random images for each category (or all images if max images for that category is less than 100). We run a state-of-the-art OCR model Rosetta [6] on these images and compute the average number of OCR boxes in a category. The average number of OCR boxes per-category were normalized and used as per-category weights for sampling the images from the categories.

We collect TextVQA’s training and validation set from Open Images’ training set while test set is collected from Open Images’ test set. We set up a three stage pipeline for crowd-sourcing our data. In the first stage, annotators were asked to identify images that did not contain text (using a forced-choice “yes”/“no” flag). Filtering those (and noisy data from annotators) out resulted in 28,408 images, which from the basis of our TextVQA dataset.

### 4.2. Questions and Answers

In the second stage, we collect 1-2 questions for each image. For the first question, we show annotators an image and ask them to provide a question which requires reading the text to answer. Specifically, they were told to ‘Please ensure that answering the question requires reading of the text in the image. It is OK if the answer cannot be directly copied from the text but needs to be inferred or paraphrased.’

To collect a second question that is different from the first, we show annotators the first question and ask them

to come up with a question that requires reasoning about the text in the image and has a different answer. Following VQA [3, 10] and VizWiz [13] datasets, we collect 10 answers for each question.

To ensure answer quality, we gave annotators instructions similar to those used in [3, 13] when collecting the VQA and VizWiz datasets. In addition, to catch any poor quality data from earlier steps, we give annotators these four options: (i) no text in image; (ii) not a question; (iii) answering the question doesn’t require reading any text in image; and (iv) unanswerable, e.g. questions involving speculation about the meaning of text. We remove the questions where a majority of workers marked any of these flags. Additionally, we use hand-crafted questions for which we know the correct answers to identify and filter out bad annotators.

### 4.3. Statistics and Analysis

We first analyze the diversity of the questions that we have in the dataset. TextVQA contains 45,336 questions of which 37,912 (83.6%) are unique. Fig. 6a shows the distribution of question length along with the same statistics for the VQA 2.0 and the VizWiz datasets for reference. The average question length in TextVQA is 7.18 words which is higher than in VQA 2.0 (6.29) and VizWiz (6.68). We also note that the minimum question length is 3 words. Workers often form questions which are longer to disambiguate the response (e.g. specifying where exactly the text

is in the image, see Fig. 3). Fig. 6d shows top 15 most occurring questions in the dataset with their count while Fig. 6e shows top 500 most occurring questions with their counts. We can see the uniform shift from common questions about “time” to questions occurring in specific situations like “team names”. Fig. 4 shows sunburst for first 4 words in questions. We also observe that most questions involve reasoning about common things (e.g. figuring out brand names, cities and temperature). Questions often start with “what”, frequently inquiring about “time”, “names”, “brands” or “authors”.

In total there are 26,263 (49.2%) unique majority answers in TextVQA. The percentage of unique answers in TextVQA is quite high compared to VQA 2.0 (3.4%) and VizWiz (22.8%). All 10 annotators agree on the most common answer for 22.8% questions, while 3 or more annotators agree on most common answer for 97.9% questions. Fig. 5 (left) shows a word cloud plot for the majority answers in the dataset. The answer space is diverse and involves brand names, cities, people’s names, time, and countries. Note that this diversity makes it difficult to have a fixed answer space – a challenge that most existing VQA datasets do not typically pose. The most common answer (“yes”) is the majority answer for only 4.71% of the dataset and “yes/no” (majority answer) questions in total only make up 5.55% of the dataset. The average answer length is 1.58 (Fig. 6b). In a few occurrences where the text in the image is long (e.g., a quote or a paragraph), the answer length is high. Fig. 6f shows the frequency of top 500 most common answers. The gradual shift from brands to rare cities is depicted. We also note that the drop in TextVQA for number of answers of a particular answer length is more gradual than in VQA 2.0 which drops sharply after answer length

Finally, we analyze the OCR tokens produced by the Rosetta OCR system [6]. In Fig. 6c, we plot number of images containing “x” number of OCR tokens. The peak between 4 and 5 shows that a lot of images in our dataset contain a good number of OCR tokens. In some cases, when the system is unable to detect text we get 0 tokens but those cases are restricted to  $\sim 1.5k$  images and we manually verified that the images actually do contain text. Fig. 5 (right) shows a word cloud of OCR tokens which shows they do contain common answers such as brand names and cities.

## 5. Experiments

We start by explaining our baselines including both heuristics and end-to-end trained models which we compare with LoRRA. We divide TextVQA into train, validation and test splits with size 34,602, 5,000, and 5,734, respectively. The TextVQA questions collected from Open Images v3’s training set were randomly split into training and validation sets. There is no image overlap between the sets. For our approach, we use a vocabulary SA of size 3996, which

Model	Accuracy(%)		Model	Vocab	Accuracy(%)	
	Val	Test			Val	Test
Human	85.01	86.79	<b>Q</b>	LA	8.09	8.70
<b>OCR UB</b>	37.12	36.52	<b>I</b>	LA	6.29	5.58
<b>LA UB</b>	48.46	48.16	<b>Pythia (I+Q)</b>	LA	13.04	14.0
<b>LA+OCR UB</b>	67.56	68.24	<b>+O</b>	LA	18.35	–
<b>Rand 100</b>	0.22	0.20	<b>+O+C</b>	n/a	20.06	–
<b>Wt. Rand 100</b>	0.27	0.26	<b>+LoRRA</b>	LA	26.23	–
<b>Majority Ans</b>	4.48	2.63	<b>+LoRRA</b>	SA	<b>26.56</b>	<b>27.63</b>
<b>Random OCR</b>	7.72	9.12	<b>BAN (I+Q)</b>	LA	12.30	–
<b>OCR Max</b>	9.76	11.60	<b>+LoRRA</b>	SA	18.41	–

Table 2: Evaluation on TextVQA. (Left) Accuracies for various heuristics baselines, which show that using **OCR** can help in achieving a good accuracy on TextVQA. **LA+OCR UB** refers to maximum accuracy achievable by models using **LoRRA** with our OCR tokens. (Right) Accuracies of our trained baselines and ablations in comparison with our model **LoRRA**. **I** denotes usage of image features, **Q** question features, **O** OCR tokens’ features, and **C** copy mechanism. **LA** and **SA** refer to use of large and short vocabulary, respectively. Models with **LoRRA** outperform VQA SoTA (Pythia, BAN) and other baselines.

contains answers which appear at least twice in the training set. For the baselines that don’t use the copy mechanism, this vocabulary turns out to be too limited. To give them a fair shot, we also create a larger vocabulary (**LA**), containing the 8000 most frequent answers. **Upper Bounds and Heuristics.** These mainly evaluate the upper bounds of what can be achieved using the OCR tokens detected by our OCR module and benchmark biases in the dataset. We test (i) **OCR UB**: the upper bound accuracy one can get if the answer can be build directly from OCR tokens (and can always be predicted correctly). **OCR UB** considers combinations of OCR tokens upto 4-grams. (ii) **LA UB**: the upper bound accuracy by always predicting the correct answer if it is present in **LA**. (iii) **LA+OCR UB**: (i) + (ii) - the upper bound accuracy one can get by predicting the correct answer if it is present in either LA or OCR tokens. (iv) **Rand 100**: the accuracy one can get by selecting a random answer from top 100 most frequent answers (v) **Wt. Rand 100**: the accuracy of baseline (iv) but with weighted random sampling using 100 most occurring tokens’ frequencies as weights. (vi) **Majority Ans**: the accuracy of always predicting the majority answer “yes” (vii) **Random OCR token**: the accuracy of predicting a random OCR token from the OCR tokens detected in an image (viii) **OCR Max**: accuracy of always predicting the OCR token that is detected maximum times in the image (e.g., “crayola” in Fig. 3 (b)).

**Baselines.**<sup>2</sup> We make modifications to the implementation discussed in Sec. 3.4 for our baselines which include (i) **Question Only (Q)**: we only use the  $f_Q(q)$  module of LoRRA to predict the answer and the rest of the features are zeroed out. (ii) **Image Only (I)**: similar to **Q**, we only

<sup>2</sup>Code for experiments is available at <https://textvqa.org/code>

use image features  $f_I(v)$  to predict answers. **Q** and **I** do not have access to OCR tokens and predict from **LA**.

**Ablations.** We create several ablations of our approach LoRRA by using the reading component and answering module in conjunction and alternatively. (i) **I+Q**: This ablation is state-of-the-art for VQA 2.0 and doesn't use any kind of OCR features; we provide results on Pythia v0.3 and BAN [25] in Tab. 1; (ii) **Pythia+O**: Pythia with OCR features as input but no copy module or dynamic answer space; (iii) **Pythia+O+C**: (ii) with the copy mechanism but no fixed answer space *i.e.* the model can only predict from the OCR tokens. Abbreviation **C** is used when we add the copy module and dynamic answer space to a model.

Our full model corresponds to **LoRRA** attached to Pythia. We also compare **Pythia+LoRRA** with small answer space (**SA**) to a version with large answer space (**LA**). We also provide results on **LoRRA** attached to BAN [25].

**Experimental Setup.** We develop our model in PyTorch [35]. We use AdaMax optimizer [26] to perform back-propagation [29]. We predict logits and train using binary cross-entropy loss. We train all of our models for 24000 iterations with a batch size of 128 on 8 GPUs. We set the maximum question length to 14 and maximum number of OCR tokens to 50. We pad rest of the sequence if it is less than the maximum length. We use a learning rate of  $5e-2$  for all layers except the *fc7* layers used for fine-tuning which are trained with  $5e-3$ . We uniformly decrease the learning rate to  $5e-4$  after 14k iterations. We calculate val accuracy using VQA accuracy metric [10] at every 1000th iteration and use the model with the best validation accuracy to calculate the test accuracy. All validation accuracies are averaged over 5 runs with different seeds.

**Results.** Tab. 2 shows accuracies on both heuristics (left) and trained baselines and models (right). Despite collecting open-ended answers from annotators, we find that human accuracy is 85.01%, consistent with that on VQA 2.0 [10] and VizWiz [13]. While the OCR system we used is not perfect, the upper-bound on the validation set that one can achieve by correctly predicting the answer using these OCR tokens is 37.12%. This is higher than our best model, suggesting room for improvement to reason about the OCR tokens. **LA UB** is quite high as they contain most commonly occurring questions. This accuracy on VQA 2.0 validation set with 3129 most common answers is 88.9% which suggests uniqueness of answers in TextVQA and limits of a fixed answer space. The difference between LoRRA and **LA+OCR UB** of 41% represents the room for improvement in modelling with current OCR tokens and **LA**. Majority answer (“yes”) gets only 4.48% on test set. Random baselines, even the weighted one, are rarely correct. **Random OCR** token selection and maximum occurring OCR token selection (**OCR Max**) yields better accuracies compared to other heuristics baselines.

Question only (**Q**) and Image only (**I**) baseline get 8.09% and 6.29% validation accuracies, respectively, which shows that the dataset does not have significant biases w.r.t. images and questions. **I+Q** models - Pythia v0.3 [40] and BAN [25], which are state-of-the-art on VQA 2.0 and VizWiz only achieve 13.04% and 12.3% validation accuracy on TextVQA, respectively. This demonstrates the inability of current VQA models to read and reason about text in images. A jump in accuracy to 18.35% is observed by feeding OCR tokens (**Pythia+O**) into the model; this supports the hypothesis that OCR tokens do help in predicting correct answers. Validation accuracy of 20.06 achieved by **Pythia+O+C** by only predicting answers from OCR tokens, further bolsters OCR importance as it is quite high compared to our Pythia v0.3 [40]. Our **LoRRA (LA)** with Pythia model outperforms all of the ablations. Finally, a slight modification which allows the model to predict from the OCR tokens more often by changing the fixed answer space **LA** to **SA** further improves performance. Validation accuracy for BAN [25] also improves to 18.41% by adding LoRRA. This suggests that LoRRA can help state-of-the-art VQA models to perform better on TextVQA.

While LoRRA can reach up to 26.56% accuracy on the TextVQA’s validation set, there is a large gap to human performance of 85.01% and LA+OCR UB of 67.56%.

Interestingly, when adding LoRRA to Pythia it improves accuracy from 68.71 to 69.21 on VQA 2.0 [9] (see Tab. 1), indicating the ability of our model to also exploit reading and reasoning in this more general VQA benchmark.

## 6. Conclusion

We explore a specific skill in Visual Question Answering that is important for applications involving aiding visually impaired users – answering questions about everyday images that involve reading and reasoning about text in these images. We find that existing datasets do not support a systematic exploration of the research efforts towards this goal. To this end, we introduce the TextVQA dataset which contains questions which can only be answered by reading and reasoning about text in images. We also introduce *Look, Read, Reason & Answer* (LoRRA), a novel model architecture for answering questions based on text in images. LoRRA reads the text in images, reasons about it based on the provided question, and predicts an answer from a fixed vocabulary or the text found in the image. LoRRA is agnostic to the specifics of the underlying OCR and VQA modules. LoRRA significantly outperforms the current state-of-the-art VQA models on TextVQA. Our OCR model, while mature, still fails at detecting text that is rotated, a bit unstructured (e.g., a scribble) or partially occluded. We believe TextVQA will encourage research both on improving text detection and recognition in unconstrained environments as well as on enabling the VQA models to read and reason about text in images.



## References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [2] Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural module networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015.
- [4] Dzimtry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, 2015.
- [5] Jeffrey P Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, et al. Vizwiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pages 333–342. ACM, 2010.
- [6] Fedor Borisjuk, Albert Gordo, and Viswanath Sivakumar. Rosetta: Large scale system for text detection and recognition in images. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 71–79. ACM, 2018.
- [7] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016.
- [8] Ross Girshick, Ilija Radosavovic, Georgia Gkioxari, Piotr Dollár, and Kaiming He. Detectron. <https://github.com/facebookresearch/detectron>, 2018.
- [9] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- [10] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [11] J Gu, Z Lu, H Li, and VOK Li. Incorporating copying mechanism in sequence-to-sequence learning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016. Association for Computational Linguistics., 2016.
- [12] Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. Pointing the unknown words. In *ACL*, 2016.
- [13] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [16] Lu Jiang, Junwei Liang, Liangliang Cao, Yannis Kalantidis, Sachin Farfadi, and Alexander G Hauptmann. Memexqa: Visual memex question answering. *arXiv:1708.01336*, 2017.
- [17] Yu Jiang, Vivek Natarajan, Xinlei Chen, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia v0. 1: the winning entry to the vqa challenge 2018. *arXiv preprint arXiv:1807.09956*, 2018.
- [18] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Computer Vision and Pattern Recognition (CVPR)*, 2017 IEEE Conference on, pages 1988–1997. IEEE, 2017.
- [19] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. In *European Chapter of the Association for Computational Linguistics*, 2017.
- [20] Kushal Kafle, Scott Cohen, Brian Price, and Christopher Kanan. Dvqa: Understanding data visualizations via question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5648–5656, 2018.
- [21] Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Akos Kadar, Adam Trischler, and Yoshua Bengio. Figureqa: An annotated figure dataset for visual reasoning. In *ICLR workshop track*, 2018.
- [22] Dimosthenis Karatzas, Luis Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *Document Analysis and Recognition (ICDAR)*, 2015 13th International Conference on, pages 1156–1160. IEEE, 2015.
- [23] Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images. In *European Conference on Computer Vision*, pages 235–251. Springer, 2016.
- [24] Aniruddha Kembhavi, Min Joon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension. In *Computer Vision and Pattern Recognition (CVPR)*, volume 2, page 3, 2017.
- [25] Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. Bilinear attention networks. In *Neural Information Processing Systems*, 2018.
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- [27] Ivan Krasin, Tom Duerig, Neil Alldrin, Andreas Veit, Sami Abu-El-Haija, Serge Belongie, David Cai, Zheyun Feng, Vittorio Ferrari, Victor Gomes, et al. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2(6):7, 2016.

- [28] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017.
- [29] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- [30] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *Advances In Neural Information Processing Systems*, pages 289–297, 2016.
- [31] Mateusz Malinowski and Mario Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Advances in neural information processing systems*, pages 1682–1690, 2014.
- [32] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. In *International Conference on Learning Representations (ICLR)*, 2017.
- [33] Anand Mishra, Karteek Alahari, and CV Jawahar. Scene text recognition using higher order language priors. In *BMVC-British Machine Vision Conference*. BMVA, 2012.
- [34] Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *The SIGNLL Conference on Computational Natural Language Learning (CoNLL)*, 2016.
- [35] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. *NIPS AutoDiff Workshop*, 2017.
- [36] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014.
- [37] Dinesh Raghu, Nikhil Gupta, et al. Hierarchical pointer memory network for task oriented dialogue. *arXiv preprint arXiv:1805.01216*, 2018.
- [38] Mengye Ren, Ryan Kiros, and Richard Zemel. Exploring models and data for image question answering. In *Advances in neural information processing systems*, pages 2953–2961, 2015.
- [39] Abigail See, Peter J Liu, and Christopher D Manning. Get to the point: Summarization with pointer-generator networks. In *Association for Computational Linguistics*, 2017.
- [40] Amanpreet Singh, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia-a platform for vision & language research. *SysML Workshop, NeurIPS 2019*, 2018.
- [41] Ray Smith. An overview of the tesseract ocr engine. In *Document Analysis and Recognition, 2007. ICDAR 2007. Ninth International Conference on*, volume 2, pages 629–633. IEEE, 2007.
- [42] Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. A corpus of natural language for visual reasoning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 217–223, 2017.
- [43] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*, 2016.
- [44] Kai Wang and Serge Belongie. Word spotting in the wild. In *European Conference on Computer Vision*, pages 591–604. Springer, 2010.
- [45] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Fvqa: Fact-based visual question answering. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [46] Caiming Xiong, Victor Zhong, and Richard Socher. Dynamic coattention networks for question answering. In *International Conference on Learning Representations (ICLR)*, 2016.
- [47] Huijuan Xu and Kate Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *European Conference on Computer Vision*, pages 451–466. Springer, 2016.
- [48] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [49] Zhou Yu, Jun Yu, Chenchao Xiang, Jianping Fan, and Dacheng Tao. Beyond bilinear: Generalized multimodal factorized high-order pooling for visual question answering. *IEEE Transactions on Neural Networks and Learning Systems*, 2018.
- [50] Yan Zhang, Jonathon Hare, and Adam Prügel-Bennett. Learning to count objects in natural images for visual question answering. In *International Conference on Learning Representations (ICLR)*, 2018.
- [51] Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4995–5004, 2016.