

TextCaps: a Dataset for Image Captioning with Reading Comprehension

Oleksii Sidorov¹, Ronghang Hu^{1,2}, Marcus Rohrbach¹, and Amanpreet Singh¹

¹ Facebook AI Research ² University of California, Berkeley
{oleksiis,mrf,asg}@fb.com, ronghang@eecs.berkeley.edu

Abstract. Image descriptions can help visually impaired people to quickly understand the image content. While we made significant progress in automatically describing images and optical character recognition, current approaches are unable to include written text in their descriptions, although text is omnipresent in human environments and frequently critical to understand our surroundings. To study how to comprehend text in the context of an image we collect a novel dataset, TextCaps, with 145k captions for 28k images. Our dataset challenges a model to recognize text, relate it to its visual context, and decide what part of the text to copy or paraphrase, requiring spatial, semantic, and visual reasoning between multiple text tokens and visual entities, such as objects. We study baselines and adapt existing approaches to this new task, which we refer to as *image captioning with reading comprehension*. Our analysis with automatic and human studies shows that our new TextCaps dataset provides many new technical challenges over previous datasets.

1 Introduction

When trying to understand man-made environments, it is not only important to recognize objects but also frequently critical to read associated text and comprehend it in the context to the visual scene. Knowing there is “a red sign” is not sufficient to understand that one is at “Mornington Crescent” Station (see Fig. 1(a)), or knowing that an old artifact is next to a ruler is not enough to know that it is “around 40 mm wide” (Fig. 1(c)). Reading comprehension in images is crucial for blind people. As the VizWiz datasets [5] suggest, 21% of questions visually-impaired people asked about an image were related to the text in it. Image captioning plays an important role in starting a visual dialog with a blind user allowing them to ask for further information as required. In addition, text out of context (*e.g.* ‘5:43p’) may be of little help, whereas scene description (*e.g.* ‘shown on a departure tableau’) makes it substantially more meaningful.

In recent years, with the availability of large labelled corpora, progress in image captioning has seen steady increase in performance and quality [4,10,12,13,37] and reading scene text (OCR) has matured [8,16,22,24,34]. However, while OCR only focuses on written text, state-of-the-art image captioning methods focus only on the visual objects when generating captions and fail to recognize and reason about the text in the scene. For example, Fig. 1 shows predictions of a state-of-the-art model [4] on a few images that require reading comprehension. The

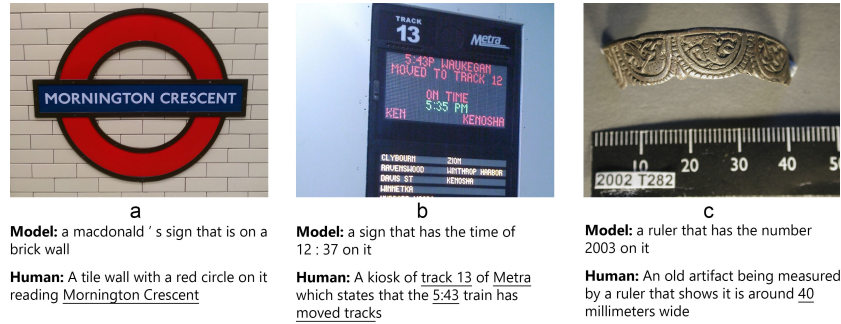


Fig. 1: **Existing captioning models cannot read!**

The *image captioning with reading comprehension* task using data from our TextCaps dataset and BUTD model [4] trained on it.

predictions clearly show an inability of current state-of-the-art image captioning methods to read and comprehend text present in images. Incorporating OCR tokens into a sentence is a challenging task, as unlike conventional vocabulary tokens which depend on the text before them and therefore can be inferred, OCR tokens often can not be predicted from the context and therefore represent independent entities. Predicting a token from vocabulary and selecting an OCR token from the scene are two rather different tasks which have to be seamlessly combined to tackle this task.

Considering the images and reference captions in Fig. 1, we can breakdown what is needed to successfully describe these images: First, detect and extract text/OCR tokens¹ (*Mornington Crescent*, *moved track*) as well the visual context such as objects in the image (*red circle*, *kiosk*). Second, generate a grammatically correct sentence which combines words from the vocabulary and OCR tokens. In addition to the challenges in normal captioning, *image captioning with reading comprehension* can include the following technical challenges:

1. Determine the relationships **between different OCR tokens** and between **OCR tokens and the visual context**, to decide if an OCR token should be mentioned in the sentence and which OCR tokens should be joined together (*e.g.* in Fig. 1b: “5:35” denotes the current time and should not be joined with “ON TIME”), based on their (a) *semantics* (Fig. 2b), (b) *spatial* relationship (Fig. 1c), and (c) *visual* appearance and context (Fig. 2d).
2. **Switching multiple times** during caption generation between the words from the model’s vocabulary and OCR tokens (Fig. 1b).
3. **Paraphrasing and inference** about the OCR tokens (Fig. 2 bold).
4. Handling of OCR tokens, including ones never seen before (**zero-shot**).

While this list should not suggest a temporal processing order, it explains why today’s models lack capabilities to comprehend text in images to generate

¹ The remainder of the manuscript we refer to the text in an image as “OCR tokens”, where one token is typically a word, i.e. a group of characters.

meaningful descriptions. It is unlikely that the above skills will naturally emerge through supervised deep learning on existing image captioning datasets as they are not focusing on this problem. In contrast, captions in these datasets are collected in a way that implicitly or explicitly avoids mentioning specific instances appearing in the OCR text. To study the novel task of image captioning with reading comprehension, we thus believe it is important to build a dataset containing captions which require reading and reasoning about text in images.

We find the COCO Captioning dataset [9] not suitable as only an estimated 2.7% of its captions mention OCR tokens present in the image, and in total there are less than 350 different OCRs (i.e. the OCR vocabulary size), moreover most OCR tokens are common words, such as “stop”, “man”, which are already present in a standard captioning vocabulary. Meanwhile, in Visual Question Answering, multiple datasets [6,26,33] were recently introduced which focus on text-based visual question answering. This task is harder than OCR recognition and extraction as it requires understanding the OCR extracted text in the context of the question and the image to deduce the correct answer. However, although these datasets focus on text reading, the answers are typically shorter than 5 words (mainly 1 or 2), and, typically, all the words which have to be generated are either entirely from the training vocabulary *or* OCR text, rather than requiring switching between them to build a complete sentence. These differences in task and dataset do not allow training models to generate long sentences. Furthermore and importantly, we require a dataset with human collected reference sentences to validate and test captioning models for *reading comprehension*.

Consequently, in this work, we contribute the following:

- For our novel task *image captioning with reading comprehension*, we collect a new dataset, **TextCaps**, which **contains 142,040 captions** on 28,408 images and requires models to read and reason about text in the image to generate coherent descriptions.
- We analyse our dataset, and find it has **several new technical challenges for captioning**, including the ability to switch multiple times between OCR tokens and vocabulary, zero-shot OCR tokens, as well as paraphrasing and inference about OCR tokens.
- Our evaluation shows that **standard captioning models fail on this new task**, while the state-of-the-art TextVQA [33] model, M4C [17], when trained with our dataset TextCaps, gets encouraging results. Our ablation study shows that it is important to take into account all semantic, visual, and spatial information of OCR tokens to generate high-quality captions.
- We conduct **human evaluations** on model predictions which show that there is a **significant gap between the best model and humans**, indicating an exciting avenue of future image captioning research.

2 Related work

Image Captioning. The Flickr30k [38] and COCO Captions [9] dataset have both been collected similarly via crowd-sourcing. The COCO Captions dataset is

significantly larger than Flickr30k and acts as a base for training the majority of current state-of-the-art image captioning algorithms. It includes 995,684 captions for 164,062 images. The annotators of COCO were asked “Describe all the important parts of the scene” and “Do not describe unimportant details”, which resulted in COCO being focused on objects which are more prominent rather than text. SBU Captions [27] is an image captioning dataset which was collected automatically by retrieving one million images and associated user descriptions from Flickr, filtering them based on key words and sentence length. Similarly, Conceptual Captions (CC) dataset [30] is also automatically constructed by crawling images from web pages together with their ALT-text. The collected annotations were extensively filtered and processed, e.g. replacing proper names and titles with object classes (*e.g.* man, city), resulting in 3.3 million image-caption pairs. This simplifies caption generation but at the same time removes fine details such as unique OCR tokens. Apart from conventional paired datasets there are also datasets like NoCaps [1], oriented to a more advanced task of captioning with zero-shot generalization to novel object classes.

While our TextCaps dataset also consists of image-sentence pairs, it focuses on the text in the image, posing additional challenges. Specifically, text can be seen as an additional modality, which models have to read (typically using OCR), comprehend, and include when generating a sentence. Additionally, many OCR tokens do not appear in the training set, but only in the test (zero-shot). In concurrent work, [15] collect captions on VizWiz [5] images but unlike TextCaps there isn’t a specific focus on reading comprehension.

Optical Character Recognition (OCR). OCR involves in general two steps, namely (i) detection: finding the location of text, and (ii) extraction: based on the detected text boundaries, extracting the text as characters. OCR can be seen as a subtask for our *image captioning with reading comprehension* task as one needs to know the text present in the image to generate a meaningful description of an image containing text. This makes OCR research an important and relevant topic to our task, which additionally requires to understand the importance of OCR token, their semantic meaning, as well as relationship to visual context and other OCR tokens. Recent OCR models have shown reliability and performance improvements [8,34,22,24,16]. However, in our experiments we observe that OCR is far from a solved problem in real-world scenarios present in our dataset.

Visual Question Answering with Text Reading Ability. Recently, three different text-oriented datasets were presented for the task of Visual Question Answering. TextVQA [33] consists of 28,408 images from selected categories of Open Images v3 dataset, corresponding 45,336 questions, and 10 answers for each question. Scene Text VQA (ST-VQA) dataset [6] has a similar size of 23,038 images and 31,791 questions but only one answer for each question. Both these datasets were annotated via crowd-sourcing. OCR-VQA [26] is a larger dataset (207,572 images) collected semi-automatically using photos of book covers and corresponding metadata. The rule generated questions were paraphrased by human annotators. These three datasets require reading and reasoning about the text in the image while considering the context for answering a question,

which is similar in spirit to TextCaps. However, the image, question and answer triplet is not directly suitable for generation of descriptive sentences. We provide additional quantitative comparisons and discussion between our and existing captioning and VQA datasets in Section 3.2.

3 TextCaps Dataset

We collect TextCaps with the goal of studying the novel task of *image captioning with reading comprehension*. Our dataset allows us to test captioning models’ reading comprehension ability and we hope it will also enable us to teach image captioning models how “to read”, *i.e.*, allow us to design and train image captioning algorithms which are able to process and include information from the text in the image. In this section, we describe the dataset collection and analyze its statistics. The dataset is publicly available at textvqa.org/textcaps.

3.1 Dataset collection

With the goal of having a diverse set of images, we rely on images from Open Images v3 dataset (CC 2.0 license). Specifically, we use the same subset of images as in the TextVQA dataset [33]; these images have been verified to contain text through an OCR system [8] and human annotators [33]. Using the same images as TextVQA additionally allows multi-task and transfer learning scenarios between OCR-based VQA and image captioning tasks. The images were annotated by human annotators in two stages.²

Annotators were asked to describe an image in one sentence which would require reading the text in the image.³

Evaluators were asked to vote yes/no on whether the caption written in the first step satisfies the following requirements: requires reading the text in the image; is true for the given image; consists of one sentence; is grammatically correct; and does not contain subjective language. The majority of 5 votes was used to filter captions of low quality. The quality of the work of evaluators was controlled using gold captions of known good/bad quality.

Five independent captions were collected for each image. An additional 6th caption was collected for the test set only to estimate human performance on the dataset. The annotators did not see previously collected captions for a particular image and did not see the same image twice. In total, we collected 145,329 captions for 28,408 images. We follow the same image splits as TextVQA for training (21,953), validation (3,166), and test (3,289) sets. An estimation performed using ground-truth OCR shows that on average, 39.5% out of all OCR tokens present in the image are covered by the collected human annotations.

² The full text of the instructions as well as screenshots of the user interface are presented in the Supplemental (Sec. F).

³ Apart from direct copying, we also allowed indirect use of text, *e.g.* inferring, paraphrasing, summarizing, or reasoning about it (see Fig. 2). This approach creates a fundamental difference from OCR datasets where alteration of text is not acceptable. For captioning, however, the ability to reason about text can be beneficial.

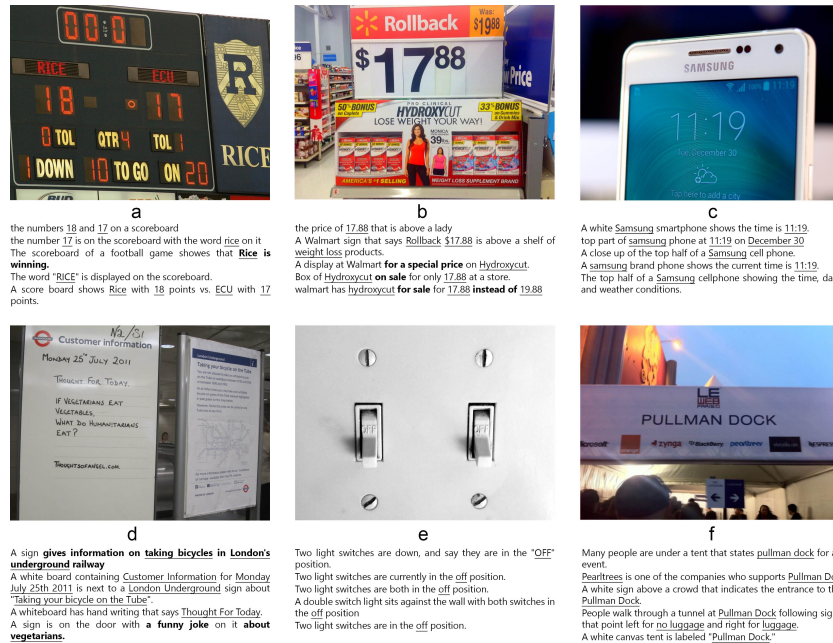


Fig. 2: **Illustration of TextCaps captions.** The bold font highlights instances which do not copy the text directly but require paraphrasing or some inference beyond copying. Underlined font highlights copied text tokens.

3.2 Dataset analysis

We first discuss several properties of the TextCaps qualitatively and then analyse and compare its statistics to other captioning and OCR-based VQA datasets.

Qualitative observations. Examples of our collected dataset in Fig. 2 demonstrate that our image captions combine the textual information present in the image with its natural language scene description. We asked the annotators to read and use text in the images but we did not restrict them to directly copy the text. Thus, our dataset also contains captions where OCR tokens are not present directly but were used to infer a description, *e.g.* in Fig. 2a “Rice is winning” instead of “Rice has 18 and Ecu has 17”. In a human evaluation of 640 captions we found that about 20% of images have at least one caption (8% of captions) which require more challenging reasoning or paraphrasing rather than just direct copying of visible text. Nevertheless, even the captions which require copying text directly can be complex and may require advanced reasoning as illustrated in multiple examples in Fig. 2. The collected captions are not limited to trivial template “Object *X* which says *Y*”. We have observed various types of relations between text and other objects in a scene which are impossible to formulate without reading comprehension. For example, in Fig. 2: “A *score board* shows




Fig. 3: **Distribution of caption/answer lengths** in Image Captioning (left) and VQA (right) datasets. VQA answers are significantly shorter than image captions and mostly concentrated within 5 words limit.

Rice with 18 points vs. ECU with 17 points” (a), “Box of Hydroxycut on sale for only 17.88 at a store” (b), “Two light switches are both in off position” (e).

Dataset statistics. To situate TextCaps properly w.r.t. other image captioning datasets, we compare TextCaps with other prominent image captioning datasets, namely COCO [9], SBU [27], and Conceptual Captions [30], as well as reading-oriented VQA datasets TextVQA [33], ST-VQA [6], and OCR-VQA [26].

The average caption length is 12.0 words for SBU, 9.7 words for Conceptual Captions, and 10.5 words for COCO, respectively. The average length for TextCaps is 12.4, slightly larger than the others (see Fig. 3). This can be explained by the fact that captions in TextCaps typically include both scene description as well as the text from it in one sentence, while conventional captioning datasets only cover the scene description. Meanwhile, the average answer length is 1.53 for TextVQA, 1.51 for ST-VQA and 3.31 for OCR-VQA – much smaller than the captions in our dataset. Typical answers like ‘yes’, ‘two’, ‘coca cola’ may be sufficient to answer a question but insufficient to describe the image comprehensively.

Fig. 4 compares the percentage of captions with a particular number of OCR tokens between COCO and TextCaps datasets.⁴ TextCaps has a much larger number of OCR tokens in the captions as well as in the images compared to COCO (note the high percentage at 0). A small part (2.7%) of COCO captions which contain OCR tokens is mostly limited to one token per caption; only 0.38% of captions contain two or more tokens. Whereas in TextCaps, multi-word reading is much more common (56.8%) which is crucial for capturing real-world information (e.g. authors, titles, monuments, etc.). Moreover, while COCO Captions contain less than 350 unique OCR tokens, TextCaps contains 39.7k of them.

We also measured the frequency of OCR tokens in the captions. Fig. 5a illustrates the number of times a particular OCR token appears in the captions.

⁴ Note that OCR tokens are extracted using Rosetta OCR system [8] which cannot guarantee exhaustive coverage of all text in an image and presents just an estimation.





Fig. 4: **Distribution of OCR tokens** in COCO and TextCaps captions (left) and images (right). In total, COCO contains **2.7%** of captions and **12.7%** of images with at least one OCR token, whereas TextCaps – **81.3%** and **96.9%**.



(a) **OCR frequency distribution** shows how many OCR tokens occur once, twice, *etc.* TextCaps has the largest amount of unique and rare (< 5) OCR tokens. Note that TextVQA has 10 answers for each question which are often identical. (b) **Number of switches between $\text{OCR} \rightleftharpoons \text{Vocab}$** illustrates the technical complexity of the datasets. An approach which cannot make switches will be sufficient for most of COCO captions and TextVQA but not for TextCaps.

Fig. 5: Analysis of OCR in our dataset vs. others

More than 9000 tokens appear only once in the whole dataset. The curve drops rapidly after 5 occurrences and only a small part of tokens occur more than 10 times. Quantitatively, 75.7% of tokens are presented less than 5 times, and only 12.9% are presented more than 10 times. The distribution specifically demonstrates the large variance in text occurring in natural images which is challenging to model using a fixed word vocabulary.

In addition to this long-tailed distribution, we find that an impressive number of 2901 of 6329 unique OCR tokens appearing in the test set captions, have neither appeared in the training nor validation set (i.e. they are “zero-shot”) which makes it necessary for models to be able to read new text in images.

TextCaps dataset also creates new technical challenges for the models. Figure 5b illustrates that due to the common use of OCR tokens in the captions, models required to switch between OCR and vocabulary words often. The majority of the TextCaps captions require to switch twice or more, whereas most COCO and TextVQA outputs can be generated even without any switches.

4 Benchmark Evaluation

4.1 Baselines

Our baselines aim to illustrate the gap between performance of conventional state-of-the-art image captioning models (BUTD [4], AoANet[18]) in comparison to recent architectures which incorporate reading (M4C [17]).

Bottom-Up Top-Down Attention model (BUTD) [4] is a widely used image captioning model based on Faster R-CNN [29] object detection features (Bottom-Up) in conjunction with attention-weighted LSTM layers (Top-Down).

Attention on Attention model (AoANet) [18] is a current SoTA captioning algorithm which uses the attention-on-attention module (AoA) to create a relation between attended vectors in both encoder and decoder.

M4C-Captioner. M4C [17] is a recent model with state-of-the-art performance on the TextVQA task. The model fuses different modalities by embedding them into a common semantic space and processing them with a multimodal transformer. Apart from that, unlike conventional VQA models where a prediction is made via classification, it enables iterative answer decoding with a dynamic pointer network [25,36], allowing the model to generate a multi-word answer, which is not limited to a fixed vocabulary. This feature makes it also suitable for reading-based caption generation. We adapt M4C to our task by removing the question input and directly use its multi-word answer decoder to generate a caption conditioned on the detected objects and OCR tokens in the image (we refer to this model as **M4C-Captioner** and illustrate it in Figure 6).

M4C-Captioner ablations. In comparison to its full version, we also evaluate a restricted version of this model without access to OCR results (referred to as **M4C-Captioner w/o OCRs**), where we use an empty OCR token list as input to the model. Additionally, we experiment with removing the pointer network (described in details in [17]) from M4C-Captioner, so that the model still has access to OCR features but cannot directly copy OCR tokens, and must use its fixed vocabulary for caption generation (referred to as **M4C-Captioner w/o copying**). As multiple types of features are used for OCR tokens in M4C-Captioner by default (same as in [17]), we further study the impact of each OCR feature type and use only **spatial** information (4-dimensional relative bounding box coordinates $[x_{\min}, y_{\min}, x_{\max}, y_{\max}]$ of OCR tokens), **semantic** information

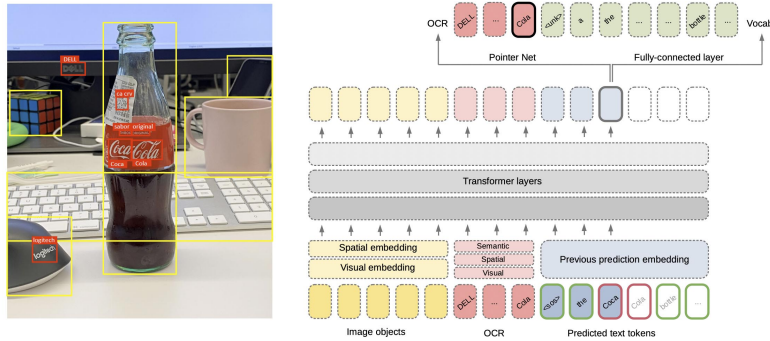


Fig. 6: **M4C-Captioner architecture** for the image captioning with reading comprehension task.

(FastText [7] and PHOC [2]), and **visual** (Faster R-CNN [29]) features in different experiments. Additionally, we use ground truth OCR tokens annotated by humans (referred to as **M4C-Captioner w/ GT OCRs**) for training and prediction⁵ to study the influence of mistakes of automatic OCR methods.

Human performance. In addition to our baselines, we provide an estimate of human performance by using the same metrics on the TextCaps test set to benchmark the progress that models still need to make. As discussed in Section 4.3, we collected one more caption for each image in the test set. The metrics are then calculated by averaging the results over 6 runs, each time leaving out one caption as a prediction, similar to [14]. On the test set, we use the same approach to evaluate machine-generated captions, so numbers are comparable.

4.2 Experimental setup

⁶ We follow the default configurations and hyper-parameters for training and evaluation of each baseline. For AoANet we use original implementation and feature extraction technique. For BUTD [4], we use the implementation and hyper-parameters from MMF [32,31]. For M4C-Captioner [17], we follow the same implementation details as used for TextVQA task [17]. We train both models for the same number of iterations on the TextCaps training set. During caption generation, we remove the `<unk>` token (for unknown words).

Datasets. We first evaluate the models trained using COCO dataset on TextCaps to demonstrate how existing datasets and models lack reading comprehension. Then we train and evaluate each baseline using TextCaps.

Metrics. Apart from automatic captioning metrics including BLEU [28], METEOR [11], ROUGE.L [23], SPICE [3], and CIDEr [35], we also perform human evaluation. We collect 5000 human scores on a Likert scale from 1 to 5

⁵ This includes a small number of images without GT-OCRs (Supplemental Sec. A).

⁶ Code for experiments is available at <https://git.io/JJGuG>

Table 1: **Performance of our baselines on our TextCaps dataset.** M4C-Captioner significantly benefits from OCR inputs and achieves the highest CIDEr score, suggesting that it is important to copy text from image on this task. However, there is still a large gap between the current machine performance and human performance, which we hope can be closed by future work.

#	Method	Trained on	TextCaps validation set metrics				
			B-4	M	R	S	C
1	BUTD [4]	COCO	12.4	13.3	33.7	8.7	24.2
2	BUTD [4]	TextCaps	20.1	17.8	42.9	11.7	41.9
3	AoANet [18]	COCO	18.1	17.7	41.4	11.2	32.3
4	AoANet [18]	TextCaps	20.4	18.9	42.9	13.2	42.7
5	M4C-Captioner	COCO	12.3	14.2	34.8	9.2	30.3
6	M4C-Captioner	TextVQA	0.1	4.4	11.3	2.8	16.9
7	M4C-Captioner w/o OCRs	TextCaps	15.9	18.0	39.6	12.1	35.1
8	M4C-Captioner w/o copying	TextCaps	18.2	19.2	41.5	13.1	49.2
9	M4C-Captioner (OCR semantic)	TextCaps	21.4	20.4	44.0	14.1	69.0
10	M4C-Captioner (OCR spatial)	TextCaps	21.7	20.6	44.6	13.7	72.0
11	M4C-Captioner (OCR visual)	TextCaps	22.5	21.3	45.3	14.4	84.0
12	M4C-Captioner (OCR semantic & visual)	TextCaps	23.4	21.5	45.8	14.9	86.0
13	M4C-Captioner	TextCaps	23.3	22.0	46.2	15.6	89.6
14	M4C-Captioner (w/ GT OCRs)	TextCaps	26.0	23.2	47.8	16.2	104.3

#	Method	Trained on	TextCaps test set metrics					
			B-4	M	R	S	C	H
15	BUTD [4]	TextCaps	14.9	15.2	39.9	8.8	33.8	1.4
16	AoANet [18]	TextCaps	15.9	16.6	40.4	10.5	34.6	1.4
17	M4C-Captioner	TextCaps	18.9	19.8	43.2	12.8	81.0	3.0
18	M4C-Captioner (w/ GT OCRs)	TextCaps	21.3	21.1	45.0	13.5	97.2	3.4
19	Human	-	24.4	26.1	47.0	18.8	125.5	4.7

B-4: BLEU-4; M: METEOR; R: ROUGE_L; S: SPICE; C: CIDEr; H: human evaluation

for a random sample of 200 images and compute median score for each caption. Fig. 7 shows that ranking of the sentence quality is the same as for automatic metrics. Moreover, all the metrics show very high correlation with human scores but CIDEr and METEOR have the highest. For comparison between different methods, we focus on the CIDEr, which puts more weight on informative n-grams in the captions (such as OCR tokens) and less weight on commonly occurring words with TF-IDF weighting.

4.3 Results

TextCaps Dataset. It can be observed in results (Table 1) that the BUTD model trained




Fig. 7: **Human evaluation in comparison to automatic metrics.**



Fig. 8: Illustration of positive and negative predictions from different models on TextCaps validation set. For M4C-Captioner, square brackets indicate tokens copied from OCR. While most of the time OCR tokens are very important for correct copying of the text from the images, for common terms such as “pepsi” or “pence”, the model sometimes prefer to select them from the vocabulary.

on the COCO captioning dataset (line 1) achieves the lowest CIDEr score, indicating that it fails to describe text in the image. When trained on the TextCaps dataset (line 2), the BUTD model has higher scores as expected, since there is no longer a domain shift between training and evaluation. AoANet (line 3, 4), which is a stronger captioning model, outperforms BUTD but still cannot handle reading comprehension and largely underperforms M4C-Captioner. For the M4C-Captioner model, there is a large gap (especially in CIDEr scores) between training with and without OCR inputs (line 13 vs. 7). Moreover, “M4C-Captioner w/o copying” (line 8) is worse than the full model (line 13) but better than the more restricted “M4C-Captioner w/o OCRs” (line 7). The results indicate that it is important to both encode OCR features **and** be able to directly copy OCR tokens. We also observe (in line 13 vs. 9-12) that it is important for a model to use spatial, visual, and semantic features of OCR tokens together, especially in the complex combinations of OCR tokens where both spatial relation and semantics play an important role in finding a connection between words. However, on the

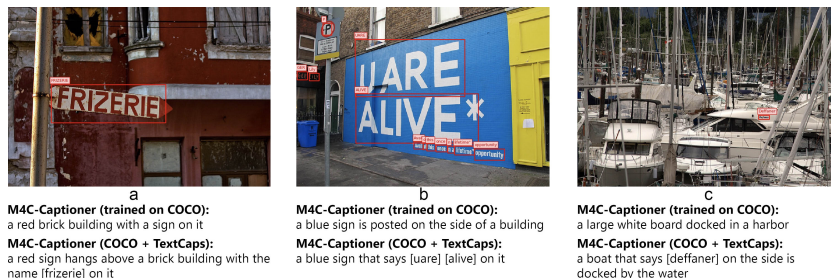


Fig. 9: **Examples of M4C-Captioner’s predictions on COCO data** when trained on COCO and TextCaps. It can be observed that *despite of availability of OCR module in both cases*, using TextCaps pushes model to read the text. Square brackets indicate tokens copied from OCR.

test set, we still notice a large gap between the best machine performance (line 17) and the human performance (line 19) on this task. Also, using ground-truth OCRs (line 18) reduces this gap but still does not close it, suggesting that there is room for future improvement in both better reasoning and better text recognition.

Figure 8 shows qualitative examples from different methods. It can be seen that BUTD and M4C-Captioner without OCR inputs rarely mention text in the image except for common brand logos such as “pepsi” that are easy to recognize visually. On the other hand, the full M4C-Captioner approach learns to read text in the image and mention it in its generated captions.⁷ Moreover, M4C-Captioner learns and recognizes relations between objects and is able to combine multiple OCR tokens into one complex description. For e.g., in Fig. 8(d) the model uses a OCR token to correctly name a player who is blocking another player; in Fig. 8(e) the model attempts to include and combine multiple tokens into a single message (“the *track* is *moved* in *Kenosha*” instead of “the word *moved*, the word *track*, and the word *Kenosha* are on the sign”). In Fig. 8(b) prediction is constructed fully from vocabulary, and even then the model counts similar objects and returns “two pepsi bottles” instead of “pepsi bottle and pepsi bottle”. We also observe a large amount of mistakes in model predictions. Many mistakes are due to wrong scene understanding and object identification, which is a common problem in captioning algorithms. We also observe placing OCR tokens in the wrong object or semantic context in the caption (Fig. 8(c, e)), incorrect repetition of an OCR token in a caption (Fig. 8(a, e)), or insufficient use of them (Fig. 8(f)) by the model. Some mistakes (as “number 3” in Fig. 8(d) are due to the errors of OCR detection algorithm and not the captioning model. This points to many potential directions for future development on this challenging generative task, which requires visual and textual understanding, requiring new model designs, conceptually different from previously existing captioning models.

Transferring to COCO. We further qualitatively show that when integrated with other datasets such as COCO [9], our dataset also enables text-based

⁷ More predictions from M4C-Captioner are presented in Supplemental (Fig. F.1).

captioning on other datasets. In this setting, we experiment training M4C-Captioner (Table 1’s best) on both TextCaps dataset and COCO dataset together. We balance the number of samples seen by the model from both COCO and TextCaps during training, and apply the trained model on the COCO validation set. COCO Captions mostly focus on visual objects but we show several examples where reading is necessary to describe the scene in Fig 9. When trained on the union of our dataset and COCO, the M4C-Captioner learns to generate captions containing text present in the images. On the other hand, the same model only describes visual objects without mentioning any text when trained on COCO alone. Quantitative results can be found in Supplemental (Sec. C).

5 Conclusion

Image captioning with reading comprehension is a novel challenging task requiring models to read text in the image, recognize the image content, and comprehend both modalities jointly to generate a succinct image caption. To enable models to learn this ability and study this task in isolation, we collected TextCaps with 142k captions. The captions include a mix of objects and/or visual scene descriptions in relation to OCR tokens copied or rephrased from the images. In most cases, OCR tokens have to be copied and related to the visual scene, but sometimes the OCR tokens have to be understood, and sometimes spatial or visual reasoning between text and objects in the image is required, as shown in our ablation study. Our analysis also points out several challenges of this dataset: Different from other captioning datasets, nearly all our captions require integration of OCR tokens, many are unseen (“zero-shot”). In contrast to TextVQA datasets, TextCaps requires generating long sentences and involves new technical challenges, including many switches between OCR and vocabulary tokens.

We find that current state-of-the-art image captioning models cannot read when trained on existing captioning dataset. However, when adapting the recent M4C VQA model to our task and training it on our TextCaps dataset, we are able to generate impressive captions on both TextCaps and COCO, which involve copying multiple OCR tokens and correctly integrating them in the captions. Our human evaluation confirms the result of the automatic metrics with very high correlation, and also shows that human captions are still significantly better than automatically generated ones, leaving room for many advances in future work, including better semantic understanding between image and text content, missing reasoning capabilities, and reading long text or single characters.

We hope our dataset with challenge server, available at textvqa.org/textcaps, will encourage the community to design better image captioning models for this novel task and address its technical challenges, especially increasing their usefulness for assisting visually disabled people.

Acknowledgments. We would like to thank Guan Pang and Mandy Toh for helping us with OCR ground-truth collection. We would also like to thank Devi Parikh for helpful discussions and insights.

References

1. Agrawal*, H., Desai*, K., Wang, Y., Chen, X., Jain, R., Johnson, M., Batra, D., Parikh, D., Lee, S., Anderson, P.: nocaps: novel object captioning at scale. In: International Conference on Computer Vision (ICCV) (2019) [4](#)
2. Almazán, J., Gordo, A., Fornés, A., Valveny, E.: Word spotting and recognition with embedded attributes. *IEEE transactions on pattern analysis and machine intelligence* **36**(12), 2552–2566 (2014) [10](#)
3. Anderson, P., Fernando, B., Johnson, M., Gould, S.: Spice: Semantic propositional image caption evaluation. In: European Conference on Computer Vision. pp. 382–398. Springer (2016) [10](#)
4. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6077–6086 (2018) [1](#), [2](#), [9](#), [10](#), [11](#)
5. Bigham, J.P., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R.C., Miller, R., Tatarowicz, A., White, B., White, S., et al.: Vizwiz: nearly real-time answers to visual questions. In: Proceedings of the 23rd annual ACM symposium on User interface software and technology. pp. 333–342. ACM (2010) [1](#), [4](#)
6. Biten, A.F., Tito, R., Mafla, A., Gomez, L., Rusiñol, M., Valveny, E., Jawahar, C., Karatzas, D.: Scene text visual question answering. *arXiv preprint arXiv:1905.13648* (2019) [3](#), [4](#), [7](#)
7. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* **5**, 135–146 (2017) [10](#)
8. Borisjuk, F., Gordo, A., Sivakumar, V.: Rosetta: Large scale system for text detection and recognition in images. In: ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 71–79. ACM (2018) [1](#), [4](#), [5](#), [7](#), [19](#)
9. Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C.L.: Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325* (2015) [3](#), [7](#), [13](#)
10. Chen, Y.C., Li, L., Yu, L., Kholy, A.E., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740* (2019) [1](#)
11. Denkowski, M., Lavie, A.: Meteor universal: Language specific translation evaluation for any target language. In: Proceedings of the ninth workshop on statistical machine translation. pp. 376–380 (2014) [10](#)
12. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL-HLT (2019) [1](#)
13. Goyal, P., Mahajan, D.K., Gupta, A., Misra, I.: Scaling and benchmarking self-supervised visual representation learning. *International Conference on Computer Vision* **abs/1905.01235** (2019) [1](#)
14. Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: Vqa 2.0 evaluation. <https://visualqa.org/evaluation.html> [10](#)
15. Gurari, D., Zhao, Y., Zhang, M., Bhattacharya, N.: Captioning images taken by people who are blind. *arXiv preprint arXiv:2002.08565* (2020) [4](#)
16. He, T., Tian, Z., Huang, W., Shen, C., Qiao, Y., Sun, C.: An end-to-end textspotter with explicit alignment and attention. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5020–5029 (2018) [1](#), [4](#)

17. Hu, R., Singh, A., Darrell, T., Rohrbach, M.: Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. arXiv preprint arXiv:1911.06258 (2019) [3](#), [9](#), [10](#), [19](#)
18. Huang, L., Wang, W., Chen, J., Wei, X.Y.: Attention on attention for image captioning. In: IEEE International Conference on Computer Vision. pp. 4634–4643 (2019) [9](#), [11](#)
19. Karatzas, D., Shafait, F., Uchida, S., Iwamura, M., i Bigorda, L.G., Mestre, S.R., Mas, J., Mota, D.F., Almazan, J.A., De Las Heras, L.P.: Icdar 2013 robust reading competition. In: 2013 12th International Conference on Document Analysis and Recognition. pp. 1484–1493. IEEE (2013) [19](#)
20. Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Haija, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Veit, A., Belongie, S., Gomes, V., Gupta, A., Sun, C., Chechik, G., Cai, D., Feng, Z., Narayanan, D., Murphy, K.: Openimages: A public dataset for large-scale multi-label and multi-class image classification. Dataset available from <https://github.com/openimages> (2017) [22](#)
21. Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Haija, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Veit, A., Belongie, S., Gomes, V., Gupta, A., Sun, C., Chechik, G., Cai, D., Feng, Z., Narayanan, D., Murphy, K.: Openimages: A public dataset for large-scale multi-label and multi-class image classification. (2017), https://storage.googleapis.com/openimages/web/factsfigures_v4.html [22](#)
22. Li, H., Wang, P., Shen, C.: Towards end-to-end text spotting with convolutional recurrent neural networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5238–5246 (2017) [1](#), [4](#)
23. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out. pp. 74–81 (2004) [10](#)
24. Liu, X., Liang, D., Yan, S., Chen, D., Qiao, Y., Yan, J.: Fots: Fast oriented text spotting with a unified network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5676–5685 (2018) [1](#), [4](#)
25. Lu, J., Yang, J., Batra, D., Parikh, D.: Neural baby talk. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7219–7228 (2018) [9](#)
26. Mishra, A., Shekhar, S., Singh, A.K., Chakraborty, A.: Ocr-vqa: Visual question answering by reading text in images. In: ICDAR (2019) [3](#), [4](#), [7](#)
27. Odonez, V., Kulkarni, G., Berg, T.L.: Im2text: Describing images using 1 million captioned photographs. In: Neural Information Processing Systems (NIPS) (2011) [4](#), [7](#)
28. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th annual meeting on association for computational linguistics. pp. 311–318. Association for Computational Linguistics (2002) [10](#)
29. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. pp. 91–99 (2015) [9](#), [10](#)
30. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2556–2565 (2018) [4](#), [7](#)
31. Singh, A., Goswami, V., Natarajan, V., Jiang, Y., Chen, X., Shah, M., Rohrbach, M., Batra, D., Parikh, D.: Mmf: A multimodal framework for vision and language research. <https://github.com/facebookresearch/mmf> (2020) [10](#)

32. Singh, A., Natarajan, V., Jiang, Y., Chen, X., Shah, M., Rohrbach, M., Batra, D., Parikh, D.: Pythia-a platform for vision & language research. In: SysML Workshop, NeurIPS. vol. 2018 (2018) [10](#)
33. Singh, A., Natarajan, V., Shah, M., Jiang, Y., Chen, X., Parikh, D., Rohrbach, M.: Towards vqa models that can read. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8317–8326 (2019) [3](#), [4](#), [5](#), [7](#), [22](#)
34. Smith, R.: An overview of the tesseract ocr engine. In: International Conference on Document Analysis and Recognition (ICDAR 2007). vol. 2, pp. 629–633. IEEE (2007) [1](#), [4](#)
35. Vedantam, R., Lawrence Zitnick, C., Parikh, D.: Cider: Consensus-based image description evaluation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4566–4575 (2015) [10](#)
36. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer networks. In: Advances in Neural Information Processing Systems. pp. 2692–2700 (2015) [9](#)
37. Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., Bowman, S.R.: Glue: A multi-task benchmark and analysis platform for natural language understanding. In: Proceedings of International Conference on Learning Representations (2019) [1](#)
38. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. Transactions of the Association for Computational Linguistics **2**, 67–78 (2014) [3](#)

TextCaps: a Dataset for Image Captioning with Reading Comprehension

(Supplementary Material)

We include the following material in our supplemental material:

- Section A.** Analysis of the influence of GT-OCR on M4C-Captioner model.
- Section B.** Precision and recall of Rosetta OCR tokens on TextCaps.
- Section C.** We show the evaluation on the COCO dataset: Our qualitative (Figure C.1) and quantitative (Table C.1) results show that COCO references captions rarely involve reading comprehension, indicating that COCO captions are not a good dataset for training or evaluating this task.
- Section D.** Qualitative illustration of frequent words in TextCaps images and captions.
- Section E.** Comparison of TextCaps Test and Validation sets.
- Section F.** Data collection User Interface.
- Figure F.1.** Additional examples of M4C-Captioner predictions.

A Analysis of the influence of GT-OCR on M4C-Captioner model

In this section, we provide additional analysis on ground-truth OCRs (GT-OCRs). So far we collected GT-OCR annotations for around 96% on the training, 96% on the validation, and 92% on the test set, we excluded OCR annotation of non-Latin/non-English characters. In Table 1 in the main paper, “M4C-Captioner (w/ GT OCRs)” (line 14 and 18) is evaluated on all TextCaps validation and test set images respectively, where an empty OCR list is used as inputs to the model on those images without GT-OCR annotations. Here, we also specifically compare the methods on the subsets of TextCaps validation and test sets, excluding those images with empty GT-OCR annotations. The results are shown in Table A.1, where ground truth OCR tokens improve the quality of generated predictions significantly.

Given that we would like to have a self-contained captioning model at test time without human in the loop, we also investigate whether ground-truth OCRs can be helpful if they are only available at training time but not at test time. In Table A.2, we compare two M4C-Captioner models trained with automatic (Rosetta) OCRs and GT-OCRs respectively and tested both with automatic (Rosetta) OCRs. It can be seen the results are very close between the two cases, suggesting that direct training on GT-OCRs does not improve the performance out-of-the-box.

The analysis of predictions from M4C-Captioner model with automatically extracted and ground-truth OCR tokens (trained and evaluated) shows that vocabulary size of all tokens used in predictions and OCR tokens in particular does not change significantly (Table A.3). Although, the quality of OCR tokens used increase, as indicated by the precision metric. Precision is calculated as

Table A.1: **Effect of using GT-OCR on performance of M4C-Captioner on TextCaps dataset.** Evaluated on a subsets of images with GT-OCR annotations (96% for the validation set, 92% for the test set).

#	Method	TextCaps validation set metrics				
		B-4	M	R	S	C
1	M4C-Captioner	23.0	21.9	46.1	15.4	88.7
2	M4C-Captioner (evaluated w/ GT OCRs)	24.4	22.8	47.0	16.2	99.4
3	M4C-Captioner (trained and evaluated w/ GT OCRs)	26.3	23.3	48.0	16.4	107.2
#	Method	TextCaps test set metrics				
		B-4	M	R	S	C
4	M4C-Captioner	19.0	19.7	43.2	12.7	80.7
5	M4C-Captioner (evaluated w/ GT OCRs)	20.4	20.8	44.5	13.7	95.1
6	M4C-Captioner (trained and evaluated w/ GT OCRs)	22.2	21.6	45.8	14.0	103.5
7	Human	24.4	26.1	46.9	18.8	125.1

B-4: BLEU-4; M: METEOR; R: ROUGE-L; S: SPICE; C: CIDEr

average ratio of OCR tokens predicted by the model which match OCR tokens used by annotators from total number of OCR predicted in each sentence.

B Rosetta OCR performance analysis

In our experiments, we use Rosetta [8] to extract OCR tokens from an image. The English-only version of Rosetta is used, which is referred to as **Rosetta-en** in [17]. To measure the performance of the Rosetta OCR system on TextCaps, we evaluated the precision and recall of OCR tokens against the human-annotated text (ground-truth OCRs) over the validation and test set images, following the ICDAR-13 evaluation protocol for end-to-end text recognition [19]. On the validation set images, the Rosetta OCR tokens have a precision of 56.50, a recall of 37.15, and an F-1 score of 44.83. On the test set images, the Rosetta OCR tokens have a precision of 53.60, a recall of 36.92, and an F-1 score of 43.72.

C Automatic evaluation on COCO captioning

Table C.1 shows the automatic evaluation metrics of the M4C-Captioner model on the COCO dataset. Here, the model trained on COCO + TextCaps⁸ has lower evaluation scores than the same model trained only on COCO. We also experiment with different sampling ratios between COCO captions and TextCaps captions, and observe that higher TextCaps ratio (sampling TextCaps captions more frequently) leads to better qualitative results where more OCR tokens are described in the generated captions, but worse CIDEr scores on the COCO

⁸ When training on COCO + TextCaps in this setting, we sample TextCaps captions more frequently than COCO captions to encourage learning text reading.

Table A.2: **Comparison between using automatic (Rosetta) OCRs and GT-OCRs during training.** We experiment with M4C-Captioner trained on our TextCaps dataset, using either automatic (Rosetta) OCRs or GT-OCRs during training. At test time, only automatic OCRs are provided, and no GT-OCRs are used in either case. Line 1 and 3 in this table are the same as line 13 and 17 in Table 1 of the main paper.

		TextCaps validation set metrics				
#	Method	B-4	M	R	S	C
1	M4C-Captioner	23.3	22.0	46.2	15.6	89.6
2	M4C-Captioner (trained w/ GT-OCRs)	23.9	21.9	46.1	15.1	88.6
		TextCaps test set metrics				
#	Method	B-4	M	R	S	C
3	M4C-Captioner	18.9	19.8	43.2	12.8	81.0
4	M4C-Captioner (trained w/ GT-OCRs)	19.7	19.8	43.6	12.4	81.9

B-4: BLEU-4; M: METEOR; R: ROUGE-L; S: SPICE; C: CIDEr

Table A.3: **Statistics of predicted sentences with automatic OCR compared to GT-OCR**

Method	vocab size			OCR	
	Total	OCR	VOCAB	token	precision
M4C-Captioner	3287	2957	545		0.62
M4C-Captioner (trained&evaluated w/ GT OCRs)	3391	3106	491		0.78

validation set. We inspect and find that this is mainly because the human captions in the COCO dataset rarely involve reading comprehension. For example, in Figure C.1, we see that the predicted captions from M4C-Captioner trained on COCO + TextCaps has noticeably lower CIDEr scores, although it learns to read and copy relevant text from the image.

Table C.1: Automatic evaluation metrics of the M4C-Captioner model on the COCO captioning validation set (Karpathy split). Here, training on TextCaps leads to lower metrics on COCO. This is mainly because the human captions in the COCO dataset do not involve reading comprehension in addition to the domain shift between the two datasets. See Sec. C and Figure C.1 for details.

#	Method	Trained on	B-4	M	R	S	C
1	M4C-Captioner	COCO	34.3	27.5	56.2	20.6	112.2
2	M4C-Captioner	COCO+TextCaps	27.1	24.1	51.6	17.4	87.5

B: BLEU-4; M: METEOR; R: ROUGE-L; S: SPICE; C: CIDEr



M4C-Captioner (trained on COCO + TextCaps):
 a large [china] airlines passenger jet is parked on the tarmac (CIDEr: 80.7)
M4C-Captioner (trained on COCO):
 a large passenger jet is sitting on top of an airport tarmac (CIDEr: 175.1)
human captions:
 china airlines plain on the ground at an airport with baggage cars nearby.
 a large passenger jet sitting on top of an airport runway.
 a large commercial plane with a flower on the tail
 a plane parked on the runway with luggage carts parked next to it
 a cargo air plane is parked on the runway



M4C-Captioner (trained on COCO + TextCaps):
 a red light at the intersection of [shoreline] [blvd] (CIDEr: 16.3)
M4C-Captioner (trained on COCO):
 a traffic light with a sky background (CIDEr: 86.1)
human captions:
 a blue sky with puffy white clouds and the top of a stop light.
 a view of the clouds beyond the stoplight.
 a traffic light in front of a cloudy blue sky.
 cloudy sky with a street light set to stop.
 a yellow streetlight beneath a sky full of clouds.

Fig. C.1: Predicted and human captions on COCO validation set (Karpathy split), where words copied from OCR tokens are taken in square brackets. As human captions in COCO rarely describe text in the image, generated captions that mention text often have lower CIDEr scores.

D Illustration of most frequent words in TextCaps

We visualize word clouds for the text tokens in TextCaps captions in Fig. C.2. In the left word cloud, it can be seen that OCR tokens copied from the image to the caption with high frequency mainly consist of brand names and other words which can be found on the products and their labels (‘samsung’, ‘nokia’, ‘colgate’, ‘ale’). In the right word cloud, when all the words are taken into account, we can observe great use of words like ‘sign’, ‘says’, ‘written’ which annotators used to incorporate text tokens into their captions.



Fig. C.2: Wordcloud visualizations of most frequent OCR tokens (left) and all words (right) in TextCaps captions.

E Comparison of TextCaps Test and Validation sets

We notice a difference in performance on our validation and test set in Table 1 of the main paper, specifically, the performance on the validation set is always higher. In this section, we discuss the similarities and differences of the two sets to understand the performance difference.

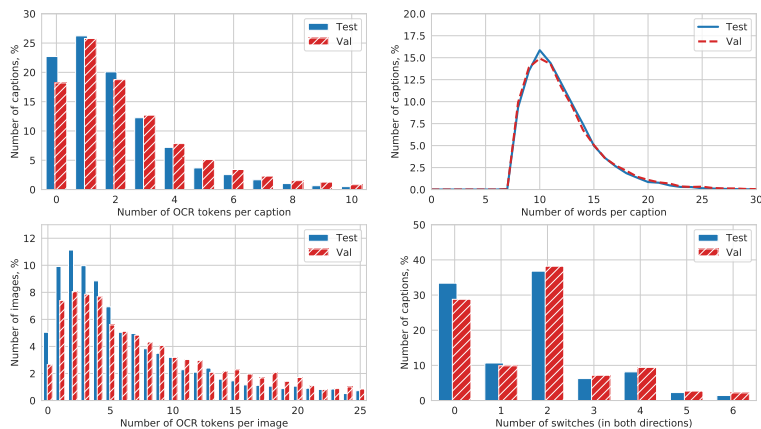


Fig. E.1: Statistics of TextCaps Test and Validation sets.

First, the images for TextCaps ‘training’ and ‘validation’ sets are from the OpenImages [20] ‘training’ set, while TextCaps ‘test’ images are from the OpenImages ‘test’ set. We observe that the image-labels of OpenImages training and test sets have slightly different distributions and categories [21]. As our training and validation set are both from the same image distribution (as we follow TextVQA’s split [33]), it is likely that models trained on the training set better fit the validation set than the test set. This is partially confirmed by evaluating a model trained only on COCO captions on TextCaps validation and test set. Here, the performance difference is smaller, for example for the BUTD model the CIDEr score drops by 4% (relative) from validation to test set when trained on COCO, but 20% when trained on TextCaps.

Second, although the captions for training, validation, and test were collected jointly, the different image distributions might affect the captions. For this, we further compare their statistics. In particular, we observe from Figure E.1 that both images and captions of the validation set have a larger number of OCR tokens on average (note that all these statistics are based on automatically extracted OCR tokens). This also causes a larger number of switches between OCR and vocabulary required in the validation set. On the other hand, in the

test set we observe more captions without any copied OCR tokens, which could suggest more paraphrasing, reasoning, and re-formulation of the OCR tokens in this set. The distribution of captions’ length is almost the same for both sets.

Third, we evaluate and compare the automatic metrics on human-written captions between the test and validation sets. Since there are only 5 human captions (instead of 6) collected on the validation set, we perform a similar leave-one-out evaluation as mentioned in Sec. 4.1 but using only 5 human captions per image (evaluating 1 human caption over the remaining 4 and averaging over the 5 runs). The results are shown in Table E.1, where the BLEU-4, METEOR, ROUGE_L, and CIDEr metrics are higher on the test set than on the validation set. This is a bit surprising, but also indicates, that there is slight domain shift between validation and test set, which humans are not affected by, rather than that the test set is more difficult in itself as the models’ performance drop might suggest.

Table E.1: **Comparison of automatic metrics on human captions between TextCaps test and validation set**, using 5 human captions per image (evaluating 1 human caption over the remaining 4 and averaging over the 5 runs).

#	Method	B-4	M	R	S	C
1	Human captions on the TextCaps validation set	22.1	24.8	44.6	20.3	118.0
2	Human captions on the TextCaps test set	22.6	25.4	45.5	20.3	127.9

B-4: BLEU-4; M: METEOR; R: ROUGE_L; S: SPICE; C: CIDEr

F Data collection User Interface

TextCaps was collected using the interfaces presented in Figures F.2-F.5. Before starting, users were presented with a list of detailed instructions (Fig. F.2 and Fig. F.4 for annotation and evaluation, respectively). The main interface window includes a panel with the same list of instructions on the left, a panel with an image in the center, and short instruction followed by the answer field on the right (Fig. F.3, F.5). It is worth noting that in the case of small or hardly-readable text users had options either to open the image in a full size in a new window or to use interactive magnifier lens with 3x zoom. Users annotated images in mini-batches of 5, and evaluated captions in mini-batches of 10. There was no restriction by time (except of very extreme limit of 15 minutes per mini-batch).



(a) a framed picture with the year [2012] on it



(b) a pair of [merrell] brand products are on a table




(c) a bottle of [deluse] sits on a table next to a small small small plastic bottle



(d) a woman is looking at a screen that says [£20] on it



(e) a trash can with a sticker on it that says [seniors] [only]



(f) a sign that says ' [cash] [custor] [pay] [firs] ' on it



(g) a plane with the number [202] on the side of it



(h) a red telephone booth with a red telephone booth sign that says [bar] [bar]



(i) a restaurant with a red sign that says [bar] [bar]



(j) a man is holding a box that says ' i ' m ' on it



(k) a black box with the word [bizhub] on it



(l) a bottle of wine with the word [ch-enet] on the label

Fig. F.1: Additional examples of M4C-Captioner predictions on the test set. Square brackets denote tokens which model selected from OCR tokens while others are from vocabulary.

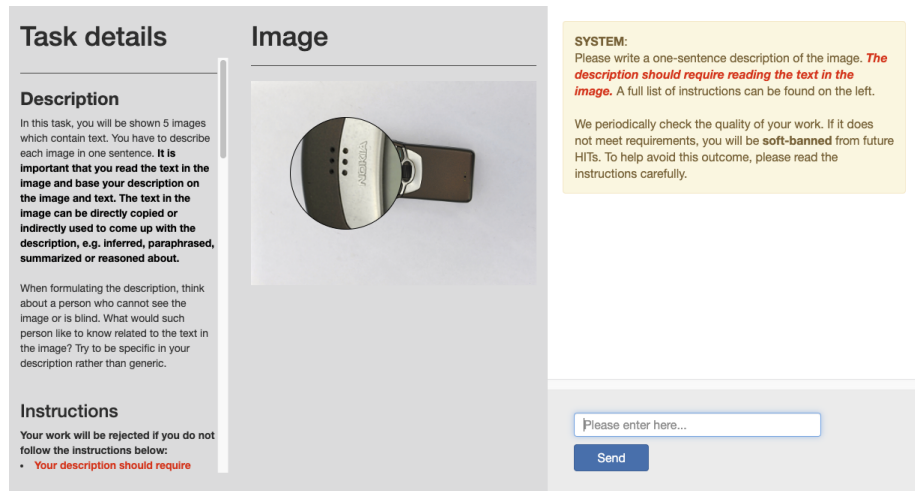


Fig. F.2: The main interface window for the annotation stage of the data collection. Detailed instructions are shown in the next Figure. The circle is an interactive zoom tool which users can move with the mouse cursor like a magnifier lens.

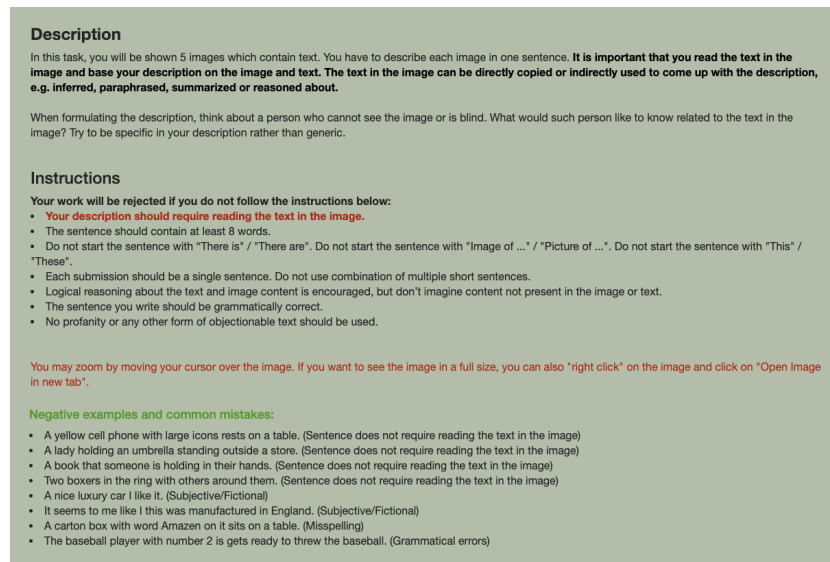


Fig. F.3: Instructions for the annotation stage of data collection. First time users saw the instructions before starting the task, after which they could find it on the left panel of our main task interface.

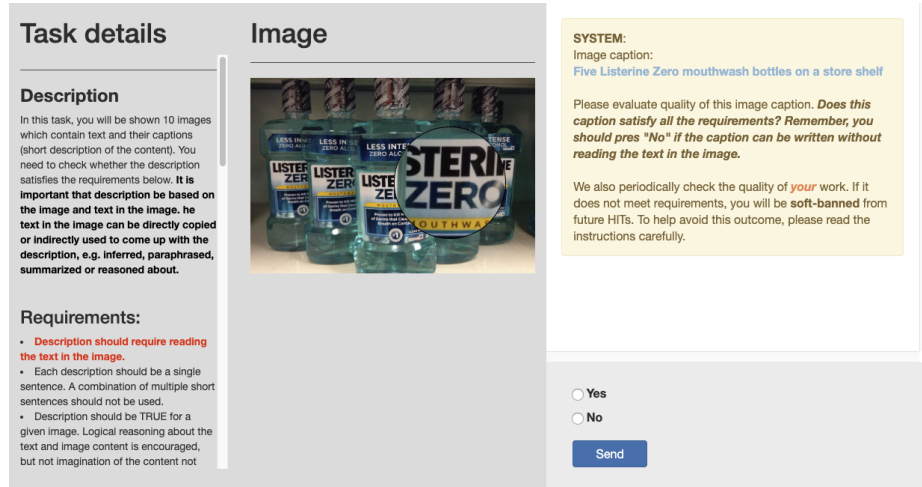


Fig. F.4: The main interface window for the evaluation stage of the data collection. Detailed instructions are shown in the next Figure. The circle is an interactive zoom tool which users can move with the mouse cursor like a magnifier lens.

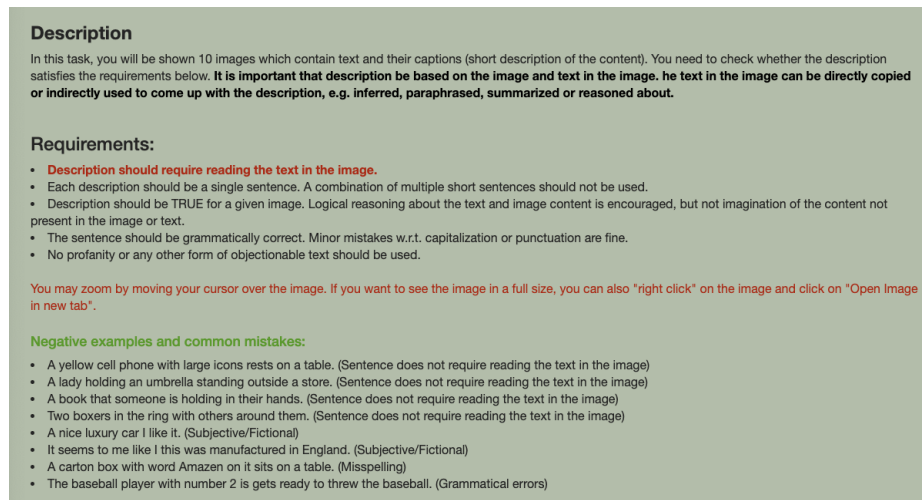


Fig. F.5: Instructions for the evaluation stage of data collection. First time users saw the instructions before starting the task, after which they could find it on the left panel of our main task interface.