

TextOCR: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text

Amanpreet Singh*, Guan Pang*, Mandy Toh*,
Jing Huang, Wojciech Galuba, and Tal Hassner

Facebook AI Research
<https://textvqa.org/textocr>

Abstract

A crucial component for the scene text based reasoning required for TextVQA and TextCaps datasets involve detecting and recognizing text present in the images using an optical character recognition (OCR) system. The current systems are crippled by the unavailability of ground truth text annotations for these datasets as well as lack of scene text detection and recognition datasets on real images disallowing the progress in the field of OCR and evaluation of scene text based reasoning in isolation from OCR systems. In this work, we propose TextOCR, an arbitrary-shaped scene text detection and recognition with 900k annotated words collected on real images from TextVQA dataset. We show that current state-of-the-art text-recognition (OCR) models fail to perform well on TextOCR and that training on TextOCR helps achieve state-of-the-art performance on multiple other OCR datasets as well. We use a TextOCR trained OCR model to create PixelM4C model which can do scene text based reasoning on an image in an end-to-end fashion, allowing us to revisit several design choices to achieve new state-of-the-art performance on TextVQA dataset.

1. Introduction

The computer vision community has recently seen a surge in interest to understand and reason on the text present in the images (scene text) beyond the OCR extraction. In particular, multiple datasets have been introduced that focus on **visual question answering (VQA)** [53, 4, 39] and **image captioning** [51] but in the context of scene text. These tasks involve understanding the objects and text in the image and then reasoning over the spatial and semantic relations between these along with a textual input (e.g. question). Though the

OCR systems have matured, they still don't work well on pictures involving real-life scenarios given the lack of large annotated real scene text OCR datasets. The text extracted by the OCR systems doesn't mean anything in itself until it is used to solve a task which involves using the scene text. Other than VQA and image captioning, the potential use cases include several impactful and interesting tasks such the hate speech and misinformation detection [25].

Although, the field has witnessed success and progress in datasets on downstream OCR applications, the performance of state-of-the-art models on these datasets are nowhere close to human accuracy due to multiple factors which includes the quality of the OCR extracted from existing OCR systems, unavailability of ground-truth text annotations for the real-world images, and no feedback to OCR system to improve detection or extraction based on the errors in the downstream application *i.e.* no end-to-end training.

In this paper, we introduce a new dataset, TextOCR which aims to bridge these gaps by providing (i) high quality and large quantity text annotations on TextVQA images (ii) allowing end-to-end training of downstream application models with OCR systems and thus allowing fine-tuning of OCR pipeline based on the task involved. Prior to TextOCR, many OCR datasets exist [38, 59, 35, 24, 23, 45, 48, 8, 61, 34, 57, 41, 50, 40, 9, 54] that propelled the field's development, but many of these are either relatively small, or focus mostly on outdoor or store-front scenes. As a result, OCR models trained on these datasets usually don't perform well on downstream tasks from other scene types. Moreover, existing datasets usually have a low number of words per image, making them less dense, diverse and ideal to train OCR models for tasks commonly having a high text density. As a solution, we present the TextOCR dataset that contains more than 28k images and 903k words in total, averaging 32 words per image. Jointly with existing TextVQA [53] and TextCaps [51] datasets, it can also serve as an OCR

*Equal Contribution. Correspondence to textvqa@fb.com

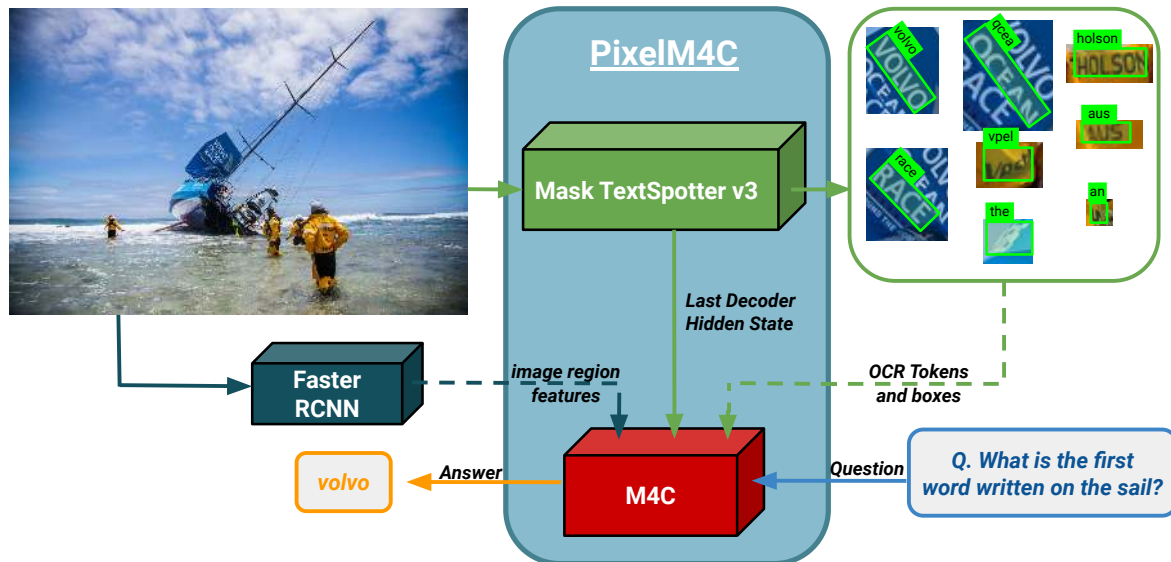


Figure 1: **PixelM4C - An end-to-end TextVQA model.** In this work, we bridge the gap between arbitrary scene-text detection/recognition and scene-text based reasoning in TextVQA [53] and TextCaps [51]. We introduce TextOCR, largest real scene-text detection and recognition dataset with 900k annotated arbitrary-shaped words collected on TextVQA images. Further, we build PixelM4C, an end-to-end TextVQA model which uses TextOCR trained Mask TextSpotter v3 [30] and M4C [17] models to do text-based reasoning directly on the images unlike previous works which rely on pre-extracted OCR text and features. The solid lines in the figure show backpropagable paths.

upper bound for researchers working on them to evaluate their methods’ reasoning capabilities on a fair ground.

In addition to TextOCR, we present a novel architecture, PixelM4C, that connects an OCR model, Mask TextSpotter (MTS) v3 [30] with downstream TextVQA model, M4C [17], in an end-to-end trainable fashion, as illustrated in Figure 1. Through extensive analysis and ablations possible with end-to-end PixelM4C, we revisit and improve design choices from prior work to achieve new state-of-the-art performance on TextVQA dataset [53] under comparable settings and show TextOCR’s impact on performance on new TextCaps dataset [51]. In summary, our main contributions include:

- A large and diverse OCR dataset with $\sim 1M$ arbitrary-shaped word annotations (3x larger than existing datasets), with high density of ~ 32 words per image.
- Extensive experiments to evaluate TextOCR showing that it is effective both as (i) a training data to push OCR state-of-the-art on multiple datasets and (ii) testing data to offer a new challenge to the community.
- A new end-to-end novel architecture, PixelM4C for TextVQA and TextCaps, which connects Mask TextSpotter (MTS) v3 [30] to M4C [17] allowing extensive analysis and revisiting prior work’s design decisions.
- State-of-the-art on TextVQA [53] using OCR tokens generated from TextOCR trained OCR models and insights from PixelM4C ablations under comparable settings.

2. Related work

2.1. OCR datasets

Recognition. The text recognition datasets which involve recognizing text from cropped words can be categorized as regular or irregular. The regular datasets like IIIT5K-Words (IIIT) [38], Street View Text (SVT) [59], ICDAR2003 (IC03) [35], ICDAR2013 (IC13) [24] have horizontally aligned words while irregular datasets like ICDAR2015 (IC15) [23], SVT Perspective (SVTP) [45], CUTE80 [48], and Total Text (TT) [8] are more challenging as they involve various transformations, such as arbitrary-oriented or curved.

Detection. Compared to older OCR datasets, which only allowed recognition as they came with pre-cropped words, the newer datasets can be used for either detection or end-to-end task as they have full images with labeled instances. The examples IC13[24], IC15[23], and TT[8] can use different word location formats, horizontal box, quadrilateral box, and curved polygon respectively. Additionally, datasets like MSRA-TD500 [61] and CTW1500 [34] with line-level labels are commonly only used for detection task.

Multi-Lingual. In recent years, there has been a surge in large-scale multi-lingual datasets containing (i) upto 7 or 8 different scripts (*e.g.* ICDAR17-MLT [41] and ICDAR19-MLT [41]), (ii) Chinese and English due to large character set (*e.g.* ICDAR17-RCTW [50], ICDAR19-ArT [9] and ICDAR19-LSVT [54]). These usually use test set for a chal-



Figure 2: **TextOCR visualizations.** The figure shows diversity and density in TextOCR images.¹

lenge not releasing the labels and being multi-lingual, the amount of data distributed in each language is smaller.

Synthetic. Synth90k [19] with 9M word-level crops, and SynthText [14] with 800K images (6M words) are most common in OCR research. As these are synthetic and contain a lot of data, they are typically used for model pretraining or joint training with real datasets.

2.2. Downstream OCR applications

In recent years, multiple datasets have been introduced for scene-text applications to study and reason about the text present in an image w.r.t. the objects in it. TextVQA [53] contains 28K images from OpenImages [28] with 45K questions, each with 10 human-annotated answers, which require reading and reasoning over scene-text to answer them. Similarly, ST-VQA [4] contains 32k questions on images from 6 different sources (IC13 [24], IC15 [23], ImageNet [10], VizWiz [3], IIIT Scene Text Retrieval, Visual Genome [27], and COCO-Text [57]). A series of datasets were introduced following these which focusing on specific aspects of text-based VQA including OCR-VQA [39], STE-VQA [60], DocVQA [36], PlotQA [37], and LEAF-QA [7].

TextCaps dataset [51] requires reading comprehension with images and contains 143K captions on TextVQA images [53]. TextCaps requires understanding how OCR words interact with each other and objects to build a caption which

¹All images are licensed under CC BY 2.0. See appendix for full attributions.

is coherent while tackling challenges like parts-of-speech, OCR and fixed vocabulary switching. VizWiz-Captions [15] also contains similar captions on VizWiz images [3] but doesn't explicitly require scene-text reasoning.

2.3. Downstream application models

The state-of-the-art on TextVQA and TextCaps use the pre-extracted text tokens from a standard OCR system as additional input [21]. As the OCR text can be any string, for word embeddings, we use a system that allows out-of-vocabulary words via character-level modeling or piece-based modeling [21]. The other textual input (*e.g.* question) is encoded using a pretrained word-embedding (BERT, GloVe [11, 44]) and fused with image's object features and OCR embeddings. The joint embedding passes through a classifier or decoder to generate the output. The state-of-the-art TextVQA model, M4C [17], uses transformers [55] to model the fusion via self and inter-modality attention to achieve 40% on TextVQA compared to 86% human accuracy. On TextCaps, M4C can be adapted to generate a sentence by taking previously generated words as text inputs at each time step. Multiple models have been introduced recently which ablate various components of M4C for better accuracy [22, 13, 16, 20]. Contrary to M4C and derivative works which treat OCR as a black box, in PixelM4C, we train an end-to-end model and use this capability to apply new design choices in an more informed way. To test our hypothesis, we build and compare PixelM4C with M4C as our base because of its simplicity and modular design.

3. TextOCR dataset

3.1. Annotation Details

For collecting arbitrary shaped scene text, all words within an image are annotated with polygon annotation for detection. For recognition, only Latin words are annotated. All non-Latin and illegible words are then annotated with a “.”. Similar to COCOText, a word is defined as an uninterrupted sequence of letters.

The annotations are performed by a set of annotators familiar with polygon word annotation. We provided the team with annotation guidelines, and a quality control training and performed rigorous checks on their work. The annotators first annotate the bounding box around a word. If the word is near-horizontal, annotators are encouraged to draw a rectangle whenever possible. If the word is curved, then the annotators draw polygon as annotations with multiple points while preserving reading direction from first annotated point to second annotated point. The annotators are also encouraged to annotate the bounding box with as little background space as possible. To ensure accuracy of predictions, our annotation pipelines includes an audit procedure, where expert annotators (authors) provide feedback to individual annotators for re-annotation. Please see appendix for more details on our annotation user interface.

3.2. Statistics and Visualizations

Figure 3 shows that TextOCR is diverse both in terms of words per image (left) as well as the word locations (right). Figure 3 (a) compare and shows high density of word annotations in TextOCR compared with COCOText [57] and ICDAR15 [23]. Figure 3 (b) and (c) compare the density of word bounding boxes in TextOCR and COCOText depicting more uniform, regular and heavy density in TextOCR suggesting that TextOCR is more precisely, uniformly and carefully annotated.

Table 1 shows the statistics of TextOCR compared to other public datasets. TextOCR has more images than most existing public datasets except for ICDAR19-LSVT [54], a bilingual dataset focusing more on street view images in Chinese. TextOCR has much larger number of annotated words than any existing public datasets, with at least 3x more words than ICDAR19-LSVT and 10x more than the rest. ICDAR19-LSVT contains only 44K words in English, while TextOCR predominantly English, thus contains 20x English words. As a result, TextOCR contains on average 32.1 words per images, 3x more than any existing datasets making it a great source for both word-level text recognition task and image-level text spotting task in text heavy scenes.

Table 2 offers more detailed statistics into TextOCR. There are total of 1.32M labeled instances in TextOCR if including empty word annotations where the word box or polygon is labeled but the text is not transcribed (due to

#	Dataset	# Images		# Words		Words per Image
		Train	Test	Train	Test	
1	Synth90k [19] [†]	-	-	8.9M	-	-
2	SynthText [14] [†]	800k	-	5.5M	-	6.9
3	IIIT5K [38]	-	-	2000	3000	-
4	SVT [59]	-	-	257	647	-
5	ICDAR2003 [35]	-	-	1156	1110	-
6	ICDAR2013 [24]	229	233	848	1095	4.2
7	ICDAR2015 [23]	1000	500	4468	2077	4.4
8	SVTP [45]	-	-	-	645	-
9	CUTE80 [48]	-	-	-	288	-
10	Total-Text [8]	1255	300	9276	2215	7.4
11	MSRA-TD500 [61]	300	200	-	-	-
12	CTW-1500 [34]	1000	500	-	-	-
13	COCO-Text [57]	18895	4416	61793	13910	3.2
14	ICDAR17-MLT [41] ^{*‡}	9000	9000	85094	n/a	9.5
15	ICDAR17-RCTW [50] [*]	8034	4229	47107	n/a	5.9
16	ICDAR19-MLT [41] [*]	10000	10000	89407	n/a	8.9
17	ICDAR19-ArT [9] [*]	5603	4563	50042	n/a	8.9
18	ICDAR19-LSVT [54] [*]	30000	20000	243385	n/a	9.1
19	TextOCR (ours)[‡]	24902	3232	822572	80497	32.1

Table 1: **TextOCR vs other datasets.** We only count non-empty words and images with at least one instance). For non end-to-end and test-only datasets, unavailable fields are left blank. * ⇒ multi-lingual datasets with no test labels and small English annotations, † ⇒ synthetic datasets. ‡ ⇒ val set counted in the train set.

Count Type	Train	Val	Test	Total
Images	21749	3153	3232	28134
Labeled instances	1052001	150338	117725	1320064
Empty words	337815	41952	37228	416995
Non-empty words	714186	108386	80497	903069
Non-alphanumeric	102744	15595	11596	129935
Less than 3 chars	197100	28726	24643	250469
Alphanumeric & 3+ chars	414342	64065	44258	522665
Rotated (degree > 20)	118547	18548	13102	150197
Curved (points > 4)	14368	3099	1843	19310

Table 2: **TextOCR statistics.** Details on instance types

illegibility or language). If we remove words that are non-alphanuemic (e.g. symbols) or have less than 3 characters (a standard in some datasets), TextOCR still contains 523k words. Among these, 150k are rotated words (> 20° angle) and 19.3k are curved words (more than 4 points used to draw the polygon), almost twice the total words in Total-Text [8], a dataset focusing on curved text.

4. OCR Experiments

In this section, we evaluate the TextOCR dataset and the challenge it presents, then exhibit its usefulness and empirically show how it can be used for both training su-

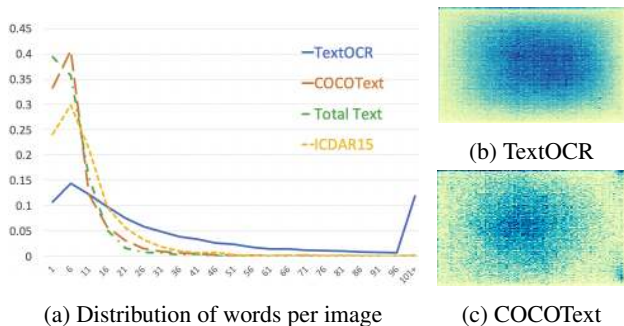


Figure 3: **TextOCR distributions.** (left) Comparison of words per image showing TextOCR’s higher text density, with $> 10\%$ images containing 100+ instances. (right) Word locations’ heatmaps across images with blue indicating higher density. TextOCR is more uniformly annotated/distributed than COCOText [57] which is more uniform than IC13 and IC15.

rior models and surpassing existing baselines on other text-recognition benchmark. We demonstrate this through three types of experiments: (i) cross-dataset empirical analysis, (ii) achieving state-of-the-art on public benchmarks, and (iii) evaluating state-of-the-art methods on TextOCR. Please refer to supplementary material for implementation details.

4.1. Cross-dataset empirical analysis

COCOText [57], one of the largest fully-annotated English OCR dataset, has images from COCO which were originally collected for object detection purpose, resulting in sparse text occurrences. As this setting is different from usual OCR applications and benchmarks, COCOText is not ideal to train or test upon. On the other hand, TextVQA dataset (image source of TextOCR) is designed for visual question answering based on text in the image leading to more prominent text, making it also a great OCR data source.

In our text recognition experiments, shown in Table 3, we use the TPS-ResNet-BiLSTM-Attn model [1], and train it on COCOText (row #1) and TextOCR (row #2) separately from scratch for 100k iterations keeping all other settings same. We evaluate and compare the results TextOCR, COCOText and other common text recognition benchmarks. The model trained on TextOCR is 22.45% better than COCOText on the TextOCR test set, and 10.56% better even on the COCOText’s test set. On the other benchmarks, TextOCR-trained model is consistently better with 10% or more gap. The superior performance can be attributed to the sheer amount of difference in number of words in these datasets compared to TextOCR. Note that training on COCOText alone (w/o synthetic data) only achieves 64.05% word accuracy on ICDAR 2015 [23], signaling it is not a good representative of oriented scene text. Comparatively, training on TextOCR alone

can achieve near state-of-the-art performance of 80.07%. Besides its large scale, we also show TextOCR has good quality compared to previous datasets, by experiments on the same number of instances as ICDAR15 and COCO-Text. Results show TextOCR is 2.5% better than ICDAR15 on average in recognition benchmarks, and 0.3% better than COCO-Text, thanks to its good quality and diversity. Please refer to supplementary experiment details.

Table 4 shows results on end-to-end recognition evaluating TextOCR’s usefulness on the image-level task. We use the latest Mask TextSpotter (MTS) V3 [30]² and train it from scratch (with ResNet50 trunk pretrained on ImageNet) on COCOText (row #1) and TextOCR (row #2) separately. We can see model fine-tuned on TextOCR again has a 0.2% lead over COCOText on its own test set, and 10%+ lead on the TextOCR and Total-Text test sets. This demonstrates the advantage of using TextOCR as a training data as it is more generalizable on other datasets. Since the number of images in TextOCR is comparable to COCO-Text (21749 vs 18895), this result is another evidence of TextOCR’s good quality.

4.2. State-of-the-art on public benchmarks

In this section, using text recognition and end-to-end experiments again, we demonstrate that TextOCR is complementary to existing datasets, and training with it can improve model accuracy significantly on existing public benchmarks, and even outperform state-of-the-art.

For text recognition, we evaluate the state-of-the-art models based upon Baek et al. [1]³, as well as fine-tune them on TextOCR’s train set in Table 3 (rows #3-10). For each method, fine-tuning on TextOCR brings a significant increase on almost all datasets. The irregular datasets (e.g. ICDAR2015 [23], SVT Perspective [45] and CUTE80 [48]), gain most thanks to the rich diversity in TextOCR.

For end-to-end recognition, we fine-tuned the official weights by Mask TextSpotter V3 on TextOCR and Total Text. Table 4 (rows #3-8) again shows that adding TextOCR can further improve the F-measure on Total Text test set by 3.3% and 3.2% with none and weak lexicon respectively. Figure 4(a) shows qualitative examples of the results.

4.3. The challenges of TextOCR

Similar to others [12], we show the challenges of TextOCR, by evaluating pre-trained and TextOCR fine-tuned state-of-the-art methods on TextOCR test set. The end-to-end recognition results on TextOCR are evaluated in the same protocol as described in [29] following ICDAR2015 with support for polygon representation. All experiments were performed with a input short side of 1000 for fair comparison. Note that TextOCR can benefit from higher short sides due to its high resolution.

²<https://github.com/MhLiao/MaskTextSpotterV3>

³<https://github.com/clovaai/deep-text-recognition-benchmark>

#	Method	PW	Train Dataset	Test Dataset (Word accuracy)								
				IIT	SVT	IC03	IC13	IC15	SVTP	CUTE	COCOText	TextOCR
Cross Dataset Experiments												
1	TPS-ResNet-BiLSTM-Attn [1]		COCOText	70.73	73.57	85.58	82.73	64.05	60.31	50.87	53.47	43.20
2	TPS-ResNet-BiLSTM-Attn [1]		TextOCR	80.50	82.84	92.16	91.25	80.07	76.74	70.38	64.03	65.65
Benchmarking state-of-the-art models and fine-tuning on TextOCR												
3	CRNN [49]	✓	S90k+ST	82.63	82.07	92.96	90.55	68.85	71.01	62.37	49.00	43.07
4	CRNN [49] (ours)		S90k+ST+TextOCR	85.97	87.94	92.96	93.70	79.85	78.30	73.87	60.09	58.61
5	Rosetta [5]	✓	S90k+ST	84.00	84.08	92.39	91.13	70.29	74.73	67.60	49.66	43.16
6	Rosetta [5] (ours)		S90k+ST+TextOCR	87.50	89.80	93.77	94.52	81.28	81.40	77.35	62.61	60.85
7	STAR-Net [32]	✓	S90k+ST	86.26	86.09	94.39	91.48	75.81	76.43	72.47	53.56	48.23
8	STAR-Net [32] (ours)		S90k+ST+TextOCR	90.30	92.12	94.69	95.33	86.09	83.88	83.62	67.92	66.84
9	TPS-ResNet-BiLSTM-Attn [1]	✓	S90k+ST	87.37	87.33	95.12	93.00	78.24	80.16	74.22	56.34	50.37
10	TPS-ResNet-BiLSTM-Attn [1] (ours)		S90k+ST+TextOCR	86.70	91.50	94.23	94.63	85.15	84.19	79.44	69.15	69.49

Table 3: **Text recognition experiments on the TextOCR dataset.** PW means the model uses public available weights. S90k and ST refer to the Synth90k [19] and SynthText [14] datasets respectively. Row #1-2 show the cross-dataset comparison between COCOText [57] and TextOCR. Row #3-10 show the experiments on state-of-the-art methods that including TextOCR in training can improve their word accuracy on most public benchmarks, as well as their word accuracy on TextOCR test set.

#	Method	Official	Train Dataset				Test Dataset (F-measure)				
			SynthText	Public	COCOText	TextOCR	TT (None)	TT (Full)	COCOText	TextOCR	
Cross Dataset Experiments											
1	Mask TextSpotter v3 [30]				✓			54.2	65.6	52.2	32.5
2	Mask TextSpotter v3 [30]					✓		64.8	74.1	52.4	45.8
Benchmarking state-of-the-art models and fine-tuning on TextOCR											
3	Qin et al. Inc-Res [46]	✓	✓	✓	✓			63.9	–	–	–
4	Mask TextSpotter v2 [29]	✓	✓	✓				65.3	77.4	47.6	–
5	Boundary TextSpotter [58]	✓	✓	✓				65.0	76.1	41.3	–
6	ABCNet [33]	✓	✓	✓	✓			64.2	75.7	–	30.5
7	Mask TextSpotter v3 [30]	✓	✓	✓				71.2	78.4	46.1	34.9
8	Mask TextSpotter v3 [30] (ours)		✓	✓		✓		74.5	81.6	57.9	50.8

Table 4: **End-to-end recognition experiments on the TextOCR dataset.** Official means either using official weights (for testing on TextOCR) or official reported results (other test data). Public refers to the model is trained with public real datasets [24, 23, 8, 41] other than COCOText [57] or TextOCR. TT is short for Total Text [8]. Row #1-2 show the cross-dataset comparison between COCOText and TextOCR. Row #3-7 show results of state-of-the-art methods, where TextOCR tests are obtained with official weights. Row #8 show improvements after fine-tuning with TextOCR train data.

Table 3 rows #3-10 and Table 4 rows #3-8 “TextOCR” column shows performance of state-of-the-art methods on text and end-to-end recognition tasks, respectively. The results demonstrate TextOCR’s challenge; even after fine-tuning with its own large train set of 21k images, the numbers are still much lower than other popular OCR datasets [23, 8], indicating a difficult task with a large room for improvement.

5. TextVQA and TextCaps Experiments

To evaluate the effectiveness and quality of TextOCR for downstream tasks, we calculate various heuristics and conduct experiments on TextVQA and TextCaps dataset using PixelM4C with TextOCR trained OCR module.

5.1. Upper Bounds and Heuristics

First, we set new precedents for the TextVQA dataset in Table 5 by recalculating the OCR-based upper bounds (UB) and heuristics for its val set presented in [53] using Rosetta [6] OCR-en namespace, OCR tokens from TextOCR trained MTS v3 [30, 18], and the annotated text present in TextOCR.

The **human** accuracy (row #1) [53] stays the same at 85.01%. For UB, unlike [53], inspired by M4C’s iterative answer prediction, we calculate the accuracy using multi-word match checking whether the answer can be built using single or multiple token(s) from the source in consideration to cover all possibilities allowing better estimates M4C-like models’ UB. **OCR UB** (row #2) shows the UB achievable by



Figure 4: (a) Examples of Mask TextSpotter V3 [30] improvement on Total Text after fine-tuning on TextOCR compared to official weights; (b) Failure cases by MTS V3 on TextOCR test set; (c) Failure cases by Baek et al. [1] on TextOCR test set

#	Method	TextVQA val accuracy(%)		
		Rosetta	MTS v3	TextOCR
1	Human	85.01	85.01	85.01
2	OCR UB	44.98	53.34	66.90
3	Vocab UB	59.02	59.02	59.02
4	OCR + Vocab UB	79.72	80.64	87.22
5	OCR Biggest	12.06	13.60	16.78
6	OCR Max	9.26	7.50	10.94

Table 5: **TextVQA heuristics.** Val accuracy for various heuristics compared with numbers from [53]. The comparison shows that TextOCR leads to much higher numbers than the original OCR tokens used in the TextVQA.

only using OCR tokens and no vocabulary which is 11% and 22% higher for TextOCR compared to MTS v3 and Rosetta justifying the requirement of a better OCR mechanism while suggesting that training OCR systems on TextOCR would be crucial for TextVQA [53] and TextCaps [51]. **Vocab UB** (row #3) shows the UB achievable by only using a fixed word vocabulary (M4C 5k vocab). **OCR+Vocab UB** (row #4) is UB achievable using both OCR and vocab inter-changeably wherever suitable for prediction. For TextOCR, this surpasses the human accuracy indicating TextOCR’s high quality and the downstream benefits of improved OCR models. **OCR Biggest** and **OCR Max** (row #5 and #6) show the UB obtained by choosing biggest OCR box and the most occurring word in the scene text as an answer respectively advocating TextVQA’s difficulty, TextOCR’s quality and improvement room in current OCR systems.

5.2. Improving the state-of-the art

Given the positive results in Section 5.1, we naturally expect that TextOCR will help with downstream tasks as

well, as we know from literature [13, 22] that OCR is indeed an important aspect. Using TextOCR annotations directly will allow us to evaluate the reasoning capabilities or shortcomings of the TextVQA/TextCaps models in isolation from OCR inconsistencies. Furthermore, this also makes it possible to train an end-to-end model that can take images directly as an input, extract OCR tokens from them and then jointly reason over the object features, OCR and input text with a possibility of backpropagating via the recognition model.

We propose an end-to-end model, PixelM4C shown in Figure 1, that works directly on the images allowing us to test our hypotheses. Specifically, we connect the Mask TextSpotter (MTS) v3 trained on TextOCR with M4C. We extract the OCR tokens and features on-the-fly from MTS v3 and pass them to M4C model allowing more fine-grained control on which features to extract and which specific parts to use based on the downstream task and model. We achieve new state-of-the-art on TextVQA using PixelM4C which allows easy testing of our various hypotheses.

Training. We train PixelM4C and PixelM4C-Captioner (similar to M4C-Captioner) in an end-to-end fashion by extracting OCR tokens from MTS v3 in real time. We use same hyper-parameters and 5k vocabulary as used by M4C [17] but we set batch size to 16 given that model is slow and hard to train on larger batch sizes. We train with Adam [26] optimizer with $1e-4$ learning rate, a step schedule and a linear warmup of 1k iterations. We train PixelM4C and PixelM4C-Captioner for 24k and 12k iterations. We decrease the learning rate to 1/10th at 14k and 19k for PixelM4C and 10k and 11k for PixelM4C-Captioner. We freeze MTS v3 during training as our empirical results suggested that fine-tuning predictor heads hurt TextVQA accuracy. We hypothesize that this happens because MTS v3 is trained using character-level losses while M4C is trained using word-level losses.

#	Method	OCR Source	OCR Feature	TextVQA	TextCaps val metrics				
				val acc (%)	B-4	M	R	S	C
1	Human	–	–	85.01	24.40	26.10	47.00	18.80	125.50
2	M4C / M4C-Captioner	Rosetta (OCR-en)	Object detector fc7	39.40	23.30	22.00	46.20	15.60	89.60
3	M4C w/ STVQA	Rosetta (OCR-en)	Object detector fc7	40.55	–	–	–	–	–
4	PixelM4C / PixelM4C-Captioner	MTS v3 (COCO-Text+TT)	MTS v3 fc7	37.61	23.09	21.08	45.55	14.51	81.44
5	PixelM4C / PixelM4C-Captioner	MTS v3 (COCO-Text+TT)	MTS v3 LH	38.24	23.05	20.88	45.45	14.23	81.55
6	PixelM4C / PixelM4C-Captioner	MTS v3 (TextOCR-en)	MTS v3 fc7	39.69	23.11	21.37	45.57	14.68	84.54
7	PixelM4C / PixelM4C-Captioner	MTS v3 (TextOCR-en)	MTS v3 LH	40.67	23.41	21.45	45.69	14.75	86.87
8	PixelM4C / PixelM4C-Captioner	MTS v3 (TextOCR)	MTS v3 fc7	39.64	23.01	21.27	45.65	14.58	84.99
9	PixelM4C / PixelM4C-Captioner	MTS v3 (TextOCR)	MTS v3 LH	41.23	23.33	21.30	45.71	14.62	85.32
10	PixelM4C w/ STVQA	MTS v3 (TextOCR)	MTS v3 LH	42.12	–	–	–	–	–
11	PixelM4C / PixelM4C-Captioner	TextOCR	MTS v3 fc7 (from 8)	46.28	23.76	21.86	46.38	15.14	91.44
12	PixelM4C / PixelM4C-Captioner	TextOCR	MTS v3 LH (from 9)	46.36	24.10	21.98	46.65	15.08	91.99
13	PixelM4C w/ STVQA	TextOCR	MTS v3 LH (from 9)	48.04	–	–	–	–	–

B-4 = Bleu4 [42], M = METEOR [2], R = ROUGE_L [31], M = METEOR [2], C = CIDEr [56]

Table 6: **PixelM4C experiments on TextVQA/TextCaps.** Val accuracy for ablations compared with M4C [17]. We show that OCR tokens and features from TextOCR trained models and directly help TextVQA and TextCaps models significantly.

Unlike M4C, we conduct ablations on using > 50 tokens given high word density in TextOCR. We train PixelM4C in a distributed fashion on 16 Nvidia Volta V100-SXM2-32GB GPUs using PyTorch based MMF framework [43, 52].

Experiments and Results. We compare PixelM4C with M4C for TextVQA and PixelM4C-Captioner with M4C-Captioner for TextCaps. Table 6 shows results for various experiments and ablations. First, by an extensive sweep (details in appendix), we confirm that batch size of 16 performs better than batch size of 128 used in [17]. We test PixelM4C and PixelM4C-Captioner with four different OCR sources: MTS v3 trained (i) on COCO-Text and Total-Text (row #4 and #5) (ii) on TextOCR but using alphanumeric English only vocabulary (row #6 and #7) (iii) on TextOCR using 240 characters Latin vocabulary (row #8, #9, and #10), (iv) using TextOCR annotations directly as the OCR source (row #11, #12, and #13), extracting features using annotation boxes as the proposals from (iii). Enabled by our end-to-end PixelM4C model, we revisit the choice of OCR feature in [17] and try other features from MTS v3. We found that using last hidden state from prediction decoder (“MTS v3 LH” in Table 6) for $< \text{EOS} >$ token as the OCR representation improves performance. Finally, we add ST-VQA [4] as extra annotation data following M4C [17] (row #10 and #13).

Based on our ablations (see appendix), we use 200 tokens instead of 50 in all experiments. MTS v3’s fc7 features as OCR representation boost accuracy when compared to Visual Genome pretrained FRCNN [47, 27] ones (row #2 vs #8). Further, we achieve state-of-the-art performance using prediction decoder’s last hidden state (row #9 vs #3) when compared to fc7 (row #4 vs #5, #6 vs #7, #8 vs #9, and #11 vs #12) suggesting that MTS v3 representations including decoder’s hidden state contain more relevant information for TextVQA task. Comparing row #5 with #7 and #9, we

observe that TextOCR trained models provide better OCR tokens for TextVQA compared to COCO-Text+TT trained ones. Finally, adding STVQA as additional data boosts the performance to 42.12% setting new state-of-the-art over M4C. In TextCaps, unfortunately, we don’t see significant improvement in metrics except B4 using TextOCR trained model’s OCR tokens testifying TextCaps’s complexity.

Using TextOCR directly as the OCR source gives a significant boost in TextVQA accuracy (6%) and TextCaps metrics (3%) signaling that apart from the gap in reasoning capabilities, there is still a room for improvement in OCR capabilities of the OCR module (MTS v3)⁴.

6. Conclusion

In this work, we introduced the large arbitrary scene text recognition dataset, TextOCR, collected on TextVQA images along with an end-to-end model, PixelM4C, that can perform scene-text reasoning directly on images by incorporating text recognition model as a module. Training on TextOCR, provides better text-recognition models which outperforms state-of-the-art on most text-recognition benchmark. Further, using TextOCR trained text-recognition module in PixelM4C allows us to use different features from it with a possibility of even providing feedback which results PixelM4C surpassing existing state-of-the-art methods on TextVQA. Through TextOCR dataset and PixelM4C model, we take a step towards bridging the communities of OCR and downstream applications based on OCR and hope that research from community will advance both fields at the same time as their is big room for improvement as evident from TextVQA results from training directly on TextOCR.

⁴We don’t claim this as state-of-the-art because it would be non-ideal for community to train directly on TextOCR except for understanding reasoning capabilities in isolation from OCR systems.

References

- [1] Jeonghun Baek, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoon Yun, Seong Joon Oh, and Hwal-suk Lee. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *Proc. Int. Conf. Comput. Vision*, 2019. 5, 6, 7
- [2] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. 8
- [3] 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, 2010. 3
- [4] Ali Furkan Biten, Ruben Tito, Andres Maffla, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4291–4301, 2019. 1, 3, 8
- [5] Fedor Borisjuk, Albert Gordo, and Viswanath Sivakumar. Rosetta: Large scale system for text detection and recognition in images. In *KDD*, 2018. 6
- [6] Fedor Borisjuk, Albert Gordo, and Viswanath Sivakumar. Rosetta: Large scale system for text detection and recognition in images. In *KDD*, 2018. 6
- [7] Ritwick Chaudhry, Sumit Shekhar, Utkarsh Gupta, Pranav Maneriker, Prann Bansal, and Ajay Joshi. Leaf-qa: Locate, encode & attend for figure question answering. In *Winter Conf. on App. of Comput. Vision*, pages 3512–3521, 2020. 3
- [8] Chee Kheng Ch'ng, Chee Seng Chan, and Cheng-Lin Liu. Total-text: Towards orientation robustness in scene text detection. *Int. J. on Document Analysis and Recognition*, 2019. 1, 2, 4, 6
- [9] Chee-Kheng Chng, Yuliang Liu, Yipeng Sun, Chun Chet Ng, Canjie Luo, Zihan Ni, ChuanMing Fang, Shuaitao Zhang, Junyu Han, Errui Ding, Jingtuo Liu, Dimosthenis Karatzas, Chee Seng Chan, and Lianwen Jin. Icdar2019 robust reading challenge on arbitrary-shaped text (rrc-art). In *Proc. Int. Conf. on Document Analysis and Recognition*, 2019. 1, 2, 4
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *Proc. Conf. Comput. Vision Pattern Recognition*, 2009. 3
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 3
- [12] Eran Eidinger, Roei Enbar, and Tal Hassner. Age and gender estimation of unfiltered faces. *Trans. on Inform. Forensics and Security*, 9(12):2170–2179, 2014. 5
- [13] Chenyu Gao, Qi Zhu, Peng Wang, Hui Li, Y. Liu, A. V. D. Hengel, and Qi Wu. Structured multimodal attentions for textvqa. *ArXiv*, abs/2006.00753, 2020. 3, 7
- [14] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *Proc. Conf. Comput. Vision Pattern Recognition*, 2016. 3, 4, 6
- [15] Danna Gurari, Yinan Zhao, Meng Zhang, and Nilavra Bhattacharya. Captioning images taken by people who are blind. *arXiv preprint arXiv:2002.08565*, 2020. 3
- [16] Wei Han, Hantao Huang, and T. Han. Finding the evidence: Localization-aware answer prediction for text visual question answering. *ArXiv*, abs/2010.02582, 2020. 3
- [17] Ronghang Hu, Amanpreet Singh, Trevor Darrell, and Marcus Rohrbach. Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9992–10002, 2020. 2, 3, 7, 8
- [18] Jing Huang, Guan Pang, Rama Kovvuri, Mandy Toh, Kevin J Liang, Praveen Krishnan, Xi Yin, and Tal Hassner. A multiplexed network for end-to-end, multilingual ocr. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 6
- [19] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. In *Neural Inform. Process. Syst.*, 2014. 3, 4, 6
- [20] Zan-Xia Jin, Heran Wu, C. Yang, Fang Zhou, Jingyan Qin, Lei Xiao, and XuCheng Yin. Ruart: A novel text-centered solution for text-based visual question answering. *ArXiv*, abs/2010.12917, 2020. 3
- [21] Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Herve Jégou, and Tomas Mikolov. Fasttext. zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*, 2016. 3
- [22] Yash Kant, Dhruv Batra, Peter Anderson, A. Schwing, D. Parikh, Jiasen Lu, and Harsh Agrawal. Spatially aware multimodal transformers for textvqa. In *European Conf. Comput. Vision*, 2020. 3, 7
- [23] Dimosthenis Karatzas, Lluís 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 *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pages 1156–1160. IEEE, 2015. 1, 2, 3, 4, 5, 6
- [24] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere De Las Heras. Icdar 2013 robust reading competition. In *2013 12th International Conference on Document Analysis and Recognition*, pages 1484–1493. IEEE, 2013. 1, 2, 3, 4, 6
- [25] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. The hateful memes challenge: Detecting hate speech in multimodal memes. *arXiv preprint arXiv:2005.04790*, 2020. 1
- [26] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015. 7

- [27] 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. *International journal of computer vision*, 123(1):32–73, 2017. 3, 8
- [28] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*, 2018. 3
- [29] Minghui Liao, Pengyuan Lyu, Minghang He, Cong Yao, Wenhao Wu, and Xiang Bai. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 5, 6
- [30] Minghui Liao, Guan Pang, Jing Huang, Tal Hassner, and Xiang Bai. Mask textspotter v3: Segmentation proposal network for robust scene text spotting. In *European Conf. Comput. Vision*, 2020. 2, 5, 6, 7
- [31] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 8
- [32] Wei Liu, Chaofeng Chen, Kwan-Yee K. Wong, Zhizhong Su, and Junyu Han. Star-net: A spatial attention residue network for scene text recognition. In *Proc. British Mach. Vision Conf.*, 2016. 6
- [33] Yuliang Liu*, Hao Chen*, Chunhua Shen, Tong He, Lianwen Jin, and Liangwei Wang. Abcnet: Real-time scene text spotting with adaptive bezier-curve network. In *Proc. Conf. Comput. Vision Pattern Recognition*, 2020. 6
- [34] Yuliang Liu, Lianwen Jin, Shuaitao Zhang, and Sheng Zhang. Detecting curve text in the wild: New dataset and new solution. *arXiv preprint arXiv:1712.02170*, 2017. 1, 2, 4
- [35] Simon M Lucas, Alex Panaretos, Luis Sosa, Anthony Tang, Shirley Wong, and Robert Young. Icdar 2003 robust reading competitions. In *Proc. Int. Conf. on Document Analysis and Recognition*, pages 682–687, 2003. 1, 2, 4
- [36] Minesh Mathew, Dimosthenis Karatzas, R. Manmatha, and C. V. Jawahar. DocVQA: A dataset for vqa on document images, 2020. 3
- [37] Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. Plotqa: Reasoning over scientific plots. In *Winter Conf. on App. of Comput. Vision*, pages 1527–1536, 2020. 3
- [38] Anand Mishra, Karteek Alahari, and C.V. Jawahar. Scene text recognition using higher order language priors. In *Proc. British Mach. Vision Conf.*, pages 1–11, 2012. 1, 2, 4
- [39] Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. Ocr-vqa: Visual question answering by reading text in images. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 947–952. IEEE, 2019. 1, 3
- [40] Nibal Nayef, Yash Patel, Michal Busta, Pinaki Nath Chowdhury, Dimosthenis Karatzas, Wafa Khlif, Jiri Matas, Umapada Pal, Jean-Christophe Burie, Cheng lin Liu, and Jean-Marc Ogier. Icdar2019 robust reading challenge on multi-lingual scene text detection and recognition – rrc-mlt-2019. In *Proc. Int. Conf. on Document Analysis and Recognition*, 2019. 1
- [41] N. Nayef, F. Yin, I. Bizid, H. Choi, Y. Feng, D. Karatzas, Z. Luo, U. Pal, C. Rigaud, J. Chazalon, W. Khlif, M. M. Luqman, J.-C. Burie, C.-L. Liu, and J.-M. Ogier. Icdar2017 robust reading challenge on multi-lingual scene text detection and script identification - rrc-mlt. In *Proc. Int. Conf. on Document Analysis and Recognition*, 2017. 1, 2, 4, 6
- [42] Kishore Papineni, Salim Roukos, Todd Ward, and Wei Jing Zhu. Bleu: a method for automatic evaluation of machine translation. 10 2002. 8
- [43] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Proceedings of NeurIPS*, pages 8024–8035. Curran Associates, Inc., 2019. 8
- [44] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. pages 1532–1543, 2014. 3
- [45] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *Proc. Int. Conf. Comput. Vision*, pages 569–576, 2013. 1, 2, 4, 5
- [46] Siyang Qin, Alessandro Bissacco, Michalis Raptis, Yasuhisa Fujii, and Ying Xiao. Towards unconstrained end-to-end text spotting. In *Proc. Int. Conf. Comput. Vision*, 2019. 6
- [47] Shaoqing Ren, Kaiming He, Ross B. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Trans. Pattern Anal. Mach. Intell.*, 39:1137–1149, 2015. 8
- [48] Anhar Risnumawan, Palaiahankote Shivakumara, Chee Seng Chan, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. *Expert Systems with Applications*, 41(18):8027–8048, 2014. 1, 2, 4, 5
- [49] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *Trans. Pattern Anal. Mach. Intell.*, 39:2298–2304, 2017. 6
- [50] Baoguang Shi, Cong Yao, Minghui Liao, Mingkun Yang, Pei Xu, Linyan Cui, Serge Belongie, Shijian Lu, and Xiang Bai. Icdar2017 competition on reading chinese text in the wild (rctw-17). In *Proc. Int. Conf. on Document Analysis and Recognition*, 2017. 1, 2, 4
- [51] Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. Textcaps: a dataset for image captioning with reading comprehension. *arXiv preprint arXiv:2003.12462*, 2020. 1, 2, 3, 7
- [52] Amanpreet Singh, Vedanuj Goswami, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Mmf: A multimodal framework for vision and language research. <https://github.com/facebookresearch/mmf>, 2020. 8

- [53] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8317–8326, 2019. [1](#), [2](#), [3](#), [6](#), [7](#)
- [54] Yipeng Sun, Zihan Ni, Chee-Kheng Chng, Yuliang Liu, Canjie Luo, Chun Chet Ng, Junyu Han, Errui Ding, Jingtuo Liu, Dimosthenis Karatzas, Chee Seng Chan, and Lianwen Jin. Icdar 2019 competition on large-scale street view text with partial labeling – rrc-lsvt. In *Proc. Int. Conf. on Document Analysis and Recognition*, 2019. [1](#), [2](#), [4](#)
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [3](#)
- [56] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proc. Conf. Comput. Vision Pattern Recognition*, pages 4566–4575, 2015. [8](#)
- [57] 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. [1](#), [3](#), [4](#), [5](#), [6](#)
- [58] Hao Wang*, Pu Lu*, Hui Zhang*, Mingkun Yang, Xiang Bai, Yongchao Xu, Mengchao He, Yongpan Wang, and Wenyu Liu. All you need is boundary: Toward arbitrary-shaped text spotting. In *AAAI Conf. on Artificial Intelligence*, 2020. [6](#)
- [59] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *Proc. Int. Conf. Comput. Vision*, pages 1457–1464, 2011. [1](#), [2](#), [4](#)
- [60] Xinyu Wang, Yuliang Liu, Chunhua Shen, Chun Chet Ng, Canjie Luo, Lianwen Jin, Chee Seng Chan, Anton van den Hengel, and Liangwei Wang. On the general value of evidence, and bilingual scene-text visual question answering. In *Proc. Conf. Comput. Vision Pattern Recognition*, June 2020. [3](#)
- [61] Cong Yao, Xiang Bai, Wenyu Liu, Yi Ma, and Zhuowen Tu. Detecting texts of arbitrary orientations in natural images. In *Proc. Conf. Comput. Vision Pattern Recognition*, 2012. [1](#), [2](#), [4](#)