

Image Captioning as an Assistive Technology: Lessons Learned from VizWiz 2020 Challenge

Pierre Dognin

PDOGNIN@US.IBM.COM

Igor Melnyk

IGOR.MELNYK@IBM.COM

Youssef Mroueh

MROUEH@US.IBM.COM

Inkit Padhi

INKPAD@IBM.COM

Mattia Rigotti

MRG@ZURICH.IBM.COM

Jarret Ross

ROSSJA@US.IBM.COM

Yair Schiff

YAIR.SCHIFF@IBM.COM

IBM Research AI, T.J. Watson Research Center, Yorktown Heights, NY, USA.

Richard A. Young

RICHARD.YOUNG2@IBM.COM

IBM Research South Africa, Johannesburg, South Africa.

Brian Belgodere

BMBELGOD@US.IBM.COM

IBM Research, T.J. Watson Research Center, Yorktown Heights, NY, USA.

Abstract

Image captioning has recently demonstrated impressive progress largely owing to the introduction of neural network algorithms trained on curated dataset like MS-COCO. Often work in this field is motivated by the promise of deployment of captioning systems in practical applications. However, the scarcity of data and contexts in many competition datasets renders the utility of systems trained on these datasets limited as an assistive technology in real-world settings, such as helping visually impaired people navigate and accomplish everyday tasks. This gap motivated the introduction of the novel VizWiz dataset, which consists of images taken by the visually impaired and captions that have useful, task-oriented information. In an attempt to help the machine learning computer vision field realize its promise of producing technologies that have positive social impact, the curators of the VizWiz dataset host several competitions, including one for image captioning. This work details the theory and engineering from our winning submission to the 2020 captioning competition. Our work provides a step towards improved assistive image captioning systems.

1. Introduction

Image captioning has witnessed steady progress since 2015, thanks to the introduction of neural caption generators with convolutional and recurrent neural networks (Vinyals et al., 2015; Karpathy & Li, 2015). Such progress, however, has been by and large demonstrated on curated datasets like MS-COCO (T. Lin et al., 2014), whose limited size and scarcity of contexts result in image captioning systems that tend to produce terse and *generic descriptive* captions. This drastically restricts the applicability of image captioning algorithms and prevents them from being widely deployed in practical real-world settings, such as helping visually impaired people navigate and accomplish everyday tasks.



Figure 1: Comparison of images in VizWiz (goal or task oriented captions) and COCO dataset (generic captions). For VizWiz images, along with the caption, we note any quality issues that are present in the image, such as blur, need for rotation, and occlusions.

The realization of these limitations and the drive to produce image captioning systems that constitute mature assistive technologies motivated the introduction of the VizWiz Challenges for captioning images taken by people who are blind in Gurari et al. (2018), Simons et al. (2020), and Gurari et al. (2020). The VizWiz Challenges are the result of a paradigm shift of image captioning towards *goal-oriented captions*, i.e. captions that not only faithfully describe a scene from everyday life, but also answer specific needs, such as those that might help blind users achieve particular tasks. A recent concurrent work to ours, refers to goal-oriented captioning as captioning with purpose (Fisch et al., 2020). Figure 1 contrasts the difference between goal-oriented (VizWiz) and generic image captioning (COCO); for example a goal-oriented caption helps the visually impaired in finding the label of a medicine box as compared to a generic description that does not specify the type of box and its label. The 2020 VizWiz Grand Challenge for image captioning (Gurari et al., 2020) was designed to address challenges in accessibility and to foster advances in designing and building assistive captioning systems from images taken by the visually impaired.

In this paper, we describe our winning entry to the 2020 VizWiz Image Captioning Grand Challenge, and highlight the lessons we learned in terms of system design from both *accessibility* and *utility* perspectives.

Firstly on accessibility, images taken by visually impaired people are captured using phones and may be blurry and have flipped orientations. The paper from Chiu et al. (2020) gives a comprehensive account of the quality challenges encountered in the real world images taken by the blind, such as blur, poor lighting conditions, non-centered framing, occlusion, and rotation issues. Examples of these image quality challenges can be seen in Figure 1. Ensuring that our machine learning pipeline is *robust* to these image quality issues and can compensate for rotation angles of the image framing is a crucial part of our work. In Section 3.1, we show how we address some of these challenges via the use of a ResNeXt network (Mahajan et al., 2018) that was pretrained on billions of Instagram images taken using phones. In Section 3.2.1, we tackle the rotation issue by building a dictionary-guided text recognizer invariant to a finite group of rotations that registers images to their canonical orientation and enables (finite group) rotation invariant text recognition.

Secondly on utility, we augment our system with reading (Sections 3.2.1 and 3.2.2) and semantic scene understanding capabilities (Section 3.2.3). Many of the VizWiz images have text that is crucial to the goal and the task at hand. We equip our pipeline with optical character detection and recognition, OCR (Y. Baek et al., 2019; J. Baek et al., 2019). Then, to ensure the robustness of OCR to rotations as discussed earlier, we perform OCR on four orientations of the image and select the orientation that has a majority of sensible words in a dictionary. In order to improve the semantic understanding of the visual scene, in addition to ResNeXt image features (IMG) and OCR, we augment our pipeline with object detection and recognition pipeline (Tan et al., 2020) (OBJ) that are discussed in Section 3.2.3.

As we describe in Section 4, our model architecture then fuses visual features with detected texts and objects that are embedded using fastText (Bojanowski et al., 2016) with a multimodal transformer (Vaswani et al., 2017). To ensure that vocabulary words coming from OCR and object detection can be used for out-of-vocabulary generation, we incorporate a copy mechanism (Singh et al., 2019; Yao et al., 2017; Gu et al., 2016) in the transformer that allows it to choose between copying an out-of-vocabulary token or predicting an in-vocabulary token. We train our system using Cross-Entropy (CE) pre-training and fine-tune using Self-Critical Sequence Training (SCST) (Rennie et al., n.d.) optimizing CIDEr (Vedantam, Zitnick, & Parikh, 2015) score.

Finally, we give experimental details and extensive ablations studies on the 2020 VizWiz Grand Challenge and the competition results in Section 5. In Section 6, we discuss qualitative evaluation of our winning entry and how it compares to models trained on generic image captions such as MS-COCO (T. Lin et al., 2014) or Google Conceptual Captions (GCC) (Sharma et al., 2018). In Section 7, we describe a real-time demo of our system that uses our captioning pipeline and the Watson Text to Speech API (IBM Research, 2021), showing the potential of image captioning as an assisting technology for the visually impaired.

2. Related Work

Generic image captioning. Recent progress in captioning since the seminal Neural Baby Talk work of Lu et al. (2018) in the pre-Deep Learning era has been accelerated by the wider adoption of deep sequence-to-sequence models. Neural Image Captioning (NIC) on MS-COCO (T. Lin et al., 2014) introduced the influential work of Show-and-Tell using LSTM architectures and cross-entropy training (Karpathy & Li, 2015; Vinyals et al., 2015). NIC won the MS-COCO 2015 challenge (Vinyals et al., 2017). Attention models introduced in Xu, Ba, Kiros, Cho, et al. (2015); Xu, Ba, Kiros, Courville, et al. (2015) further improved performance in image captioning and were refined in bottom-up and top-down attention models (Anderson et al., 2018a). Transformers models (Vaswani et al., 2017) have been adapted to multimodal scenarios, such as image captioning and visual question answering (VQA) in works like Kant et al. (2020) and Luo et al. (2019), which won the conceptual captions challenge on GCC dataset in 2019 (Sharma et al., 2018). Generic image captioning systems were trained on MS-COCO or GCC benchmark using cross-entropy training. Improvements were further obtained using automatic metric training, such as CIDEr and SPICE, combined with Reinforcement Learning techniques, as in Ranzato et al. (2015); Rennie et al. (n.d.); Liu et al. (2017). Generation of more natural sounding captions was also addressed by the use of GAN losses in Dai et al. (2017); Shetty et al. (2017);

Dognin et al. (2019) and better sequence-to-sequence architectures such as Transformers (Vaswani et al., 2017) as in Luo et al. (2019). Copy mechanisms in image captioning and VQA ensure generalization to words not in the vocabulary that are available from other channels, such as object or text detectors (Singh et al., 2019; Yao et al., 2017; Gu et al., 2016).

Accessibility and assistive technologies challenges and needed functionalities.

VizWiz Social is a mobile application centered around accessibility whereby visually impaired users can take pictures, ask questions about them and crowd-source real-time answers (Bigham et al., 2010). Brady et al. (2013) analyzed the type of images uploaded and questions asked by the users on this app and grouped them into the following categories: identification, reading, and description (Brady et al., 2013). Kacorri et al. (2017) studied the training of personalized object recognition using VizWiz social. Assistive technologies and machine learning based systems require a set of functionalities to meet the accessibility needs of the visually impaired such as: object, text, and color recognition, as well as counting (Zeng et al., 2020).

VizWiz challenges and datasets. Using the data collected from the VizWiz Social application, several VizWiz datasets and challenges have been proposed on VQA in 2018 (Gurari et al., 2018; Simons et al., 2020) and the image captioning challenge in 2020 (Gurari et al., 2020). The latter is the main focus of our work.

Assistive Technologies. While the space of assistive captioning technology remains an open domain, there are certain solutions available that have begun to address this important need using a variety of mobile applications, wearable devices, and other assistive hardware. For example, ORCAM (OrCam, 2021) is a wearable device solution that combines voice-activated commands with text, object, and facial detection to help users read text and recognize faces and products on-demand. A different approach comes in the form of the NavCog (IBM Research and CMU, 2021) application and smart guiding suitcase prototype, which utilizes radio and Bluetooth signals combined with machine learning optimization to determine a user’s location and assist the visually impaired in navigating to a desired destination. Finally, the solution that most closely aligns to the goal-oriented automated captioning, addressed in this work, is found in the Seeing AI mobile application (Microsoft AI, 2021), which provides text reading, currency, product, and facial recognition, and scene description capabilities.

In addition to these wearable and navigation assistance technologies, there are several solutions that provide automatic enhancements to the screen reader technology, which reads text from computer screens aloud. For a useful taxonomy defining the design space of alternative text, which these screen readers rely on to describe visual media on the computer, see Morris et al. (2018). As demonstrated by the qualitative study conducted in Stangl et al. (2020), the needs of screen reader users often vary by where images are encountered, e.g., on a news site vs. on a social networking platform, and the surrounding context. When alternative text is either missing or provides non-meaningful content, there is a need for automatic solutions that can augment screen readers. For example, Caption Crawler (Guinness et al., 2018) performs reverse image search and crawls the web for more informative captions that can be used by a screen reader. Another technology that is more

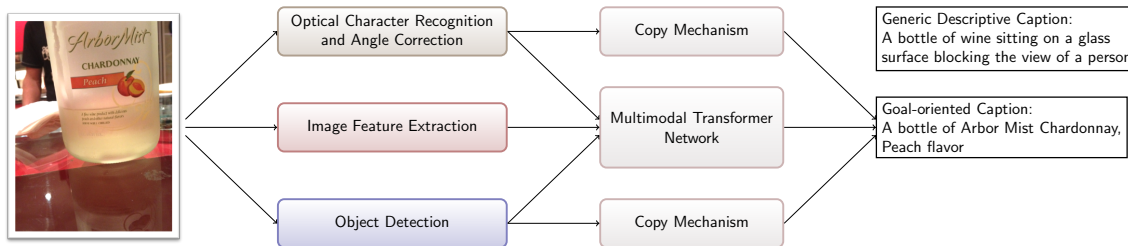


Figure 2: Multimodal Assistive Captioner Overview. An input image is passed through 3 separate modules to perform image feature extraction, OCR and angle correction, and object detection. Their outputs are combined and submitted to a Multimodal Transformer Network. The generated caption is further biased towards the OCR and object detection outputs by means of copy mechanisms to produce more goal-oriented content, rather than a generic descriptive caption.

similar to our captioning work is Twitter A11y (Gleason et al., 2020), which is a browser extension that automatically adds alternative text to images aiming to improve the screen reader experience on social media platforms that increasingly present visual content.

3. Model

In this section, we describe our model and the different data channels we incorporate to address task-oriented image captioning. An overview of our multimodal assistive captioner is given in Figure 2. As seen there, in addition to the image feature extraction discussed in Section 3.1, our system incorporates OCR and object detectors that are detailed in Section 3.2. A technical overview is given in Figure 3.

3.1 Image Feature Extraction with ResNeXt

For image processing, we choose ResNeXt (Mahajan et al., 2018) as our feature extractor since we observed a 10-point CIDEr gain compared to ResNet 101 features (He et al., 2016) (all else equal). We hypothesize that this gain comes from the fact that VizWiz images, which are mainly captured using cell phone cameras, more closely align to the ResNeXt pre-training that uses Instagram phone-taken images. We preserve the size of the image as long as one of the sides is greater than 320 pixels. If this requirement is not met, we increase the offending image axis to 320 pixels while maintaining the aspect ratio. We normalize images using the method from Mahajan et al. (2018). To extract features, we use the 99th layer of the 101 layer ResNeXt (Mahajan et al., 2018) convolutional neural network with cardinality of 32 and bottleneck dimension of 8. These features are finally pushed through a 2-dimensional adaptive max pooling layer forming a $14 \times 14 \times 2048$ feature vector for each image. Before being input into our multimodal Transformer, the image feature vector is reshaped to 196×2048 .

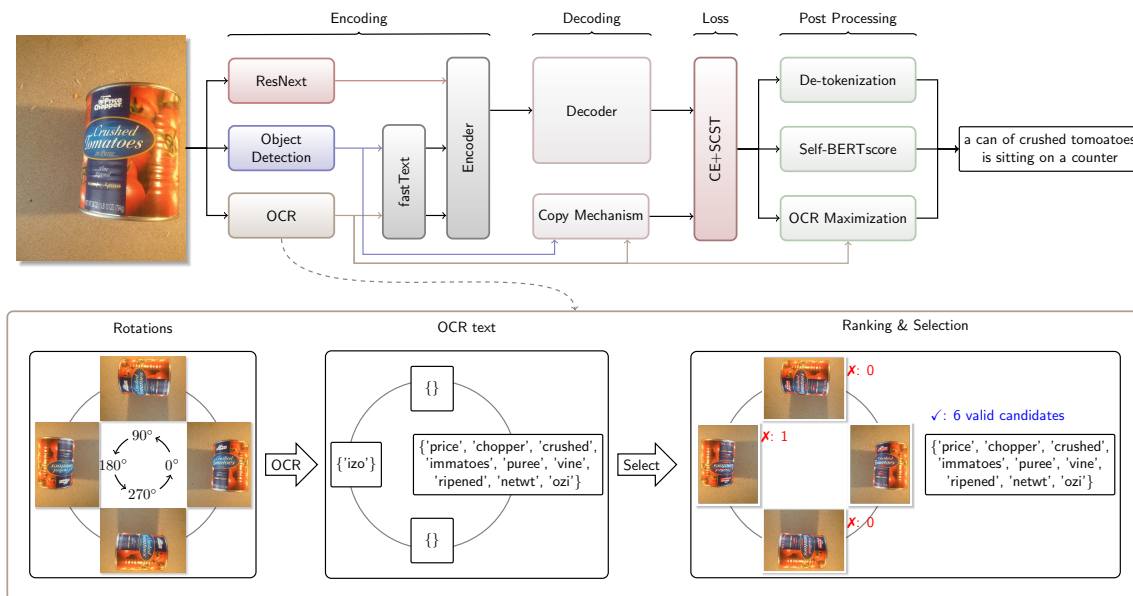


Figure 3: Multimodal Assistive Captioner. Our captioner is a multimodal Transformer with an encoder/decoder architecture composed of 4 distinct stages. First, an Encoding stage provides embeddings for the 3 modality streams of image features, object detection, and OCR results derived from an input image. Image features are extracted by a ResNeXt model, while object detection and OCR modules generate sequences of object categories (10 max) and strings of characters (20 max) respectively, which are then embedded using fastText. Embeddings for these 3 modalities are combined and fed to a Decoding stage where a decoder generates a caption whose content is biased towards OCR and object detection results by a copy mechanism that influences the Loss computation stage. The Post Processing stage finalizes the generated caption by cleaning up the caption, maximizing the use of OCR tokens, and promoting coherent candidate (via self-BERT score). The OCR module is expanded to show the inner workings of our Dictionary-Guided Rotation-Invariant OCR module. The input image rotations lead to 4 sets of OCR results which enable the selection of a proper rotation based on the largest number of intelligible tokens.

3.2 Incorporating OCR and Objects

In this section, we focus on how to incorporate inputs from OCR and object detection.

3.2.1 INCORPORATING OCR

Manual inspection of the VizWiz images and captions reveals that OCR is an integral component of building meaningful captions for images that are provided by visually impaired users. Product labels, nutrition facts, instructions, and other relevant text on objects and screens are often found in the VizWiz dataset images. Therefore, when composing our

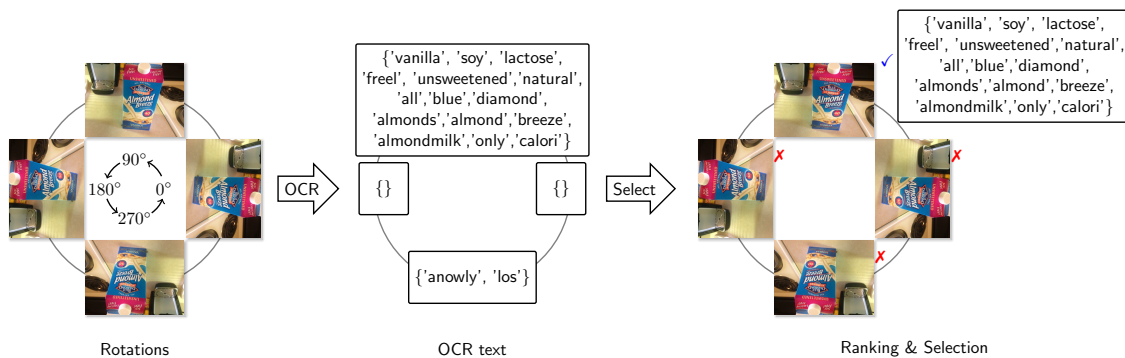


Figure 4: OCR example with two competing rotations. In this case, OCR detects tokens for two rotations, and after dictionary-guided ranking the tokens associated with 90° rotation are selected.

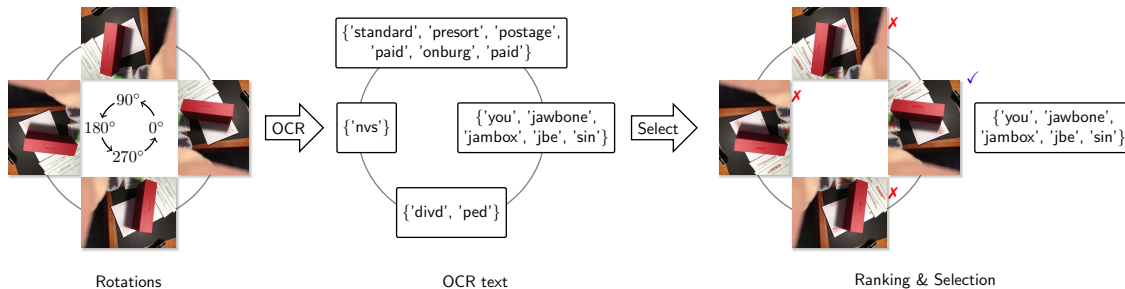


Figure 5: OCR example with four competing rotations. In this case, OCR detects tokens for all four rotations, and after dictionary-guided ranking the tokens associated with 0° rotation are selected.

pipeline to tackle this challenge, it was evident that a state of the art OCR component would be critical to render our captioning system useful.

To that end, we investigated several open source OCR solutions, and eventually decided to incorporate Y. Baek et al. (2019) and J. Baek et al. (2019). Having demonstrated success reading text in various fonts and lighting conditions, this open source solution was a useful addition to our pipeline.

3.2.2 DICTIONARY-GUIDED ROTATION-INVARIANT OCR

Although this open source implementation provides strong OCR capabilities, we noted that it did contain some failure modes, specifically with regards to text orientation. The nature of the VizWiz dataset made this a unique obstacle for this accessibility challenge, as images often have non-horizontal orientations: approximately 28.5% of the training images and 28.3% of validation images have rotation issues, according to the metadata annotations on the competition website. We addressed this shortcoming by designing an OCR module whose output is invariant to a finite group of rotations of the input, inspired by work that related group invariance to robustness and generalization in classification tasks (see e.g. Mroueh et al. (2015)). In practice, we achieved said finite-group rotation invariance by simply augmenting the OCR component of our pipeline by pre-processing images in four possible orientations, i.e. 0, 90, 180, and 270 degree rotations. Specifically, for each rotation,

we pass it through the OCR module and extract the predicted text. The OCR predictions are also passed to the fastText (Bojanowski et al., 2016) tokenizer. The orientation that produces the greatest number of intelligible tokens, according to the fastText vocabulary, is selected. We also observed that in images which contain a lot of text, the detected texts that contribute to the caption the most are usually the ones having the largest bounding box. This inspired us to rank detected OCR texts by bounding area and text confidence in descending order.

In Figures 4 and 5, we present schematic representations of the OCR modality for two different images. For each image, we consider four rotations which are then passed through the OCR module. These examples highlight how the OCR module is sensitive to rotations and, in some cases, yields no legible text. Through the inclusion of dictionary-guided ranking, we select OCR text(s) from a particular view, making the overall modality invariant to the rotation. The final step of the OCR module involves creating features using fastText embeddings of the final text(s).

3.2.3 OBJECT DETECTION

In addition to the OCR component, we also included an object detection module (OBJ) to aid the captioner in composing more detailed and accurate image descriptions. After reviewing several popular off-the-shelf solutions, we decided to use EfficientDet from Tan et al. (2020). The main trade-off in selecting the detector is between the number of recognized objects versus the detection accuracy. Although EfficientDet covers only 80 objects based on MS-COCO categories (T. Lin et al., 2014), its accuracy is high enough to be a reliable input source. This is in contrast to other approaches (e.g, YOLO9000 (Redmon & Farhadi, 2016)), which offer larger object coverage but suffer from less reliable detection accuracy. For the VizWiz challenge, the EfficientDet was used as is, without any modifications or adaptations. Given an image, the OBJ detector returns a list of objects and the corresponding confidence levels. We then sort the objects based on the confidence and select the ones above 0.25. Similar to the OCR module, the detected object strings are encoded using fastText and passed to the captioner as a separate modality.

4. Multimodal Transformer

In this section, we present the architecture choices made to create our multimodal transformer.

4.1 Regular Transformer (without copy mechanism)

Given the image features extracted from ResNeXt that we reshape to a sequence of $n_{\text{pixel}} \times 2048$ (where $n_{\text{pixel}} = 196$), detected objects (clipped to a maximum $n_{\text{obj}} = 10$ objects), and OCR (clipped to maximum $n_{\text{ocr}} = 20$ words) embedded in fastText to 300-dimensional vectors each, we embed those three modalities to a joint space of dimension $d = 512$ using linear projections. Then, we concatenate the three modalities resulting in a sequence of length $n = n_{\text{pixel}} + n_{\text{ocr}} + n_{\text{obj}}$ of dimension d . This simple concatenation is then fed to a regular transformer (Vaswani et al., 2017). Note that the self-attention in this case is a great model for intra and inter correlations between the modalities: the diagonal of the self-attention matrix represents self-attention on each modality alone, and the off-diagonal

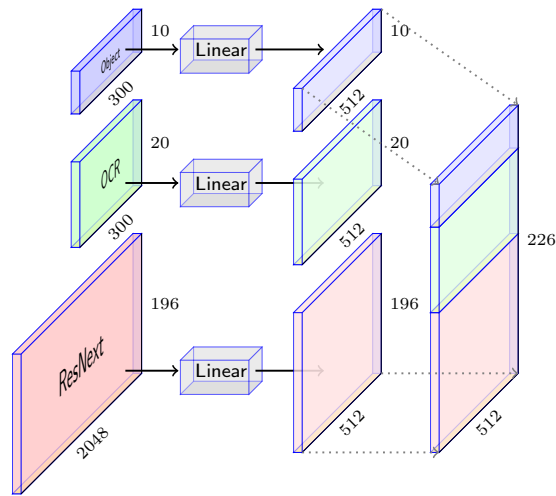


Figure 6: Multimodal features transformation and concatenation.

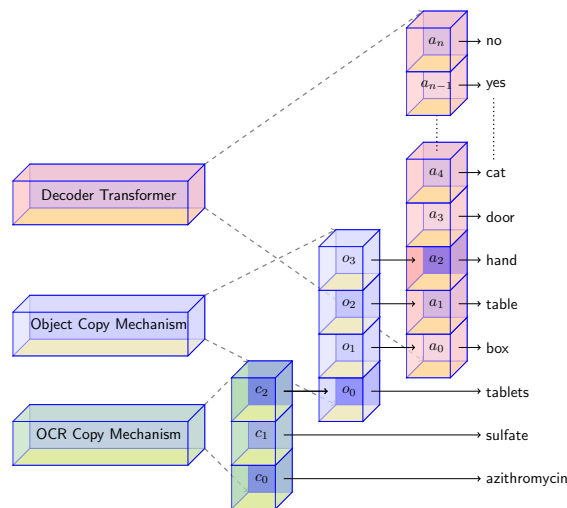


Figure 7: Biasing softmax and vocab expansion with copy mechanism for object detection and OCR.

terms represent cross attention between modalities (IMG, OCR), (IMG, OBJ), and (OBJ, OCR). A summary of this step is given in Figure 6.

4.2 Copy Transformer and Dynamic Vocabulary

OCR and OBJ are important input modalities for building semantically meaningful representations of a scene. In addition, segments from the OCR and OBJ can be directly useful as parts of the output sequence. Specifically, giving the decoder transformer the option to sample from the OCR and OBJ outputs is a straightforward way to extend the vocabulary

of the captions with *out-of-vocabulary tokens* that are potentially relevant to describing the scene. A mechanism that is naturally suited to achieve this is the “copy mechanism”, which was originally proposed for seq2seq NLP tasks in Gu et al. (2016) as a way to accommodate both understanding of the input sequences and rote memorization of parts of it and direct copy in the output sequence. Yao et al. (2017) subsequently incorporated copying in image captioning, and See et al. (2017) use a similar mechanism for summarizing text.

Inspired by this line of work, we decided to implement a copy mechanism that optionally copies tokens generated by the OCR and OBJ modules directly to the output. This is obtained by expanding the vocabulary of the decoder transformer with the outputs of the OCR and OBJ modules. Concretely, we do this by concatenating the caption and OCR and OBJ vocabularies. We then mask out at caption generation the unnormalized logits associated with tokens that were not present in the OCR and OBJ outputs. We compute the probability of sampling each one of the tokens in the resulting expanded vocabulary by applying a softmax non-linearity on the vector of concatenated unmasked unnormalized logits. As a result, the probability that the decoder transformer samples a token y_t at generation step t conditioned on an image and the previously generated tokens can be written as a normalized average of probabilities:

$$p(y_t|\cdot) = \frac{1}{Z} \left(e^{\phi_{\text{IMG}}(y_t)} + e^{\phi_{\text{OCR}}(y_t)} + e^{\phi_{\text{OBJ}}(y_t)} \right),$$

where Z is a normalization term, and $\phi_{\text{IMG}}(\cdot)$, $\phi_{\text{OCR}}(\cdot)$, $\phi_{\text{OBJ}}(\cdot)$ are score functions outputting unnormalized logits over the captions, OCR, and OBJ vocabulary, respectively, with the understanding that $\phi_{\text{OCR}}(\cdot)$ is set to $-\infty$ for tokens that do not appear in the output of the OCR module, and $\phi_{\text{OBJ}}(\cdot)$ is set to $-\infty$ for tokens that do not appear in the output of the OBJ module. In practice, this corresponds to biasing the caption generator to sample from the outputs of the OCR and OBJ modules, as illustrated in Figure 7, and most importantly, it endows our captioning architecture with a “dynamic vocabulary” that comprises out-of-vocabulary tokens that might not be present in the caption vocabulary, but are captured by the OCR and OBJ modules.

5. Experiments

In this section, we describe the VizWiz Captions dataset, and provide details about training our model.

5.1 Dataset

We use the VizWiz Captions dataset for all our experiments. This dataset is split into training, validation, and testing subsets as described in Table 1. Given the original dataset, we prepare the data for training by first removing any special characters, such as `\n` (new line), `\r` (carriage return), and `\t` (tab), from the captions. Then, we remove any punctuation characters from the captions before taking their lowercase versions. Finally, each caption is tokenized using the BERT tokenizer, with any words not in the BERT dictionary replaced by the `<unk>` token. At training time, the captions, together with OCR and OBJ detected tokens bring the total of observed unique tokens to 35,555.

Table 1: VizWiz Captions dataset information.

Subset	Images	Captions
Training	23,431	117,155
Validation	7,750	138,750
Testing	8,000	40,000

5.2 Training Details

In this section, we discuss the details of our model training process. We start with the basic Cross-Entropy (CE) pre-training, then we randomly select one CE model (usually close to the end of the training) to initialize the Self-Critical Sequence Training (SCST) step. We repeat the above CE+SCST steps to build a multitude of models by varying the model hyper-parameter values (such as number of Transformer layers, hidden dimension size, random seed initialization, etc), varying the set of used data modalities (IMG, OCR and OBJ), and turning the copy mechanism on/off. Finally, we combine these model permutations into an ensemble for evaluation and testing.

CE Training The CE training is run for 10 epochs, using a batch size of 80. We employ SGD with ADAM optimizer with $(\beta_1, \beta_2) = (0.9, 0.98)$. We warm the learning rate with a factor of 1 for 2000 iterations (minibatches) and then decay it proportionally to $1/\sqrt{i}$, where i is the iteration step.

Evaluation Metrics To examine the quality of the generated image descriptions, several automatic metrics for caption evaluations have been proposed. We briefly review a few of them that will be used in this paper.

BLEU (BiLingual Evaluation Understudy) (Papineni et al., 2002) is a popular and widely used metric that focuses on precision. Its value is computed as the geometric mean of n-gram (in this work we use 4-gram) matching between the generated caption and multiple references with a brevity penalty to discourage shorter captions.

METEOR (Metric for Evaluation of Translation with Explicit ORdering) (Banerjee & Lavie, 2005) is a metric that is based on the harmonic mean of unigram precision and recall, and is computed via an explicit word-to-word matching (or synonym matching) between generated and one or more reference sentences.

ROUGE (Recall-Oriented Understudy for Gisting Evaluation) (C.-Y. Lin, 2004) is a recall-based metric and thus has a tendency to reward long detailed sentences. The value is usually computed based on n-gram recall between the candidate and the references. Alternatively, longest common sub-sequence-based statistics is also frequently applied, as in this work.

CIDEr (Consensus-based Image Description Evaluation) (Vedantam, Lawrence Zitnick, & Parikh, 2015) is a popular metric that uses TF-IDF scores (to decrease the influence of frequent, less informative words) to weight shorter and longer n-grams across the reference captions, where the higher order n-grams can usually better capture the grammatical properties as well as encode richer semantics.

SPICE (Semantic Propositional Image Caption Evaluation) (Anderson et al., 2016) is another metric which, in contrast to the above examples, does not use n-grams but rather

estimates the caption quality by transforming the candidate and references into a semantic representation called a scene graph, that is eventually converted into a set of word tuples. The final metric is computed as an F-score over these tuples.

SCST Training Recall, that in CE training we maximize the likelihood of the next ground-truth token given the previous ground-truth token, while at test time the next token is predicted given previously self-generated tokens. This creates a train-test mismatch. Moreover, although CE loss is used for training, the evaluation metrics such as BLEU, CIDEr, etc. are used for testing, again creating a mismatch. To address these issues, we employ Self-Critical Sequence Training (SCST) to fine-tune our model after the CE training. SCST is a variant of the REINFORCE algorithm (Williams, 1992) whose goal is to maximize a reward, usually given in the form of a non-differentiable evaluation metric (here we use CIDEr since it is the main ranking metric used in the competition). In SCST, we first sample a caption and compute its CIDEr reward, then we compare it to the reward obtained by a test-time baseline (a greedy max decoding). The sampled captions outperforming the baseline are given a positive weight, while under-performing samples are suppressed.

We perform SCST on CE models from randomly chosen epochs and a variety of initial random number generator seeds. We used a batch size of 80 images/captions as well as the states learned in the ADAM optimizer during CE training. To reduce the learning rate for SCST, we initialize the step number to be $50k$ plus the last step number of CE training. Therefore, there is no warm up in SCST. Finally, we perform SCST for a random number of additional epochs ranging from 15 to 40.

Ensembling The final step in our training pipeline is to form an ensemble from the obtained SCST models. At each step of caption generation, we first average the probabilities across the models for each vocabulary element and then sample the most likely token (greedy decoding).

5.3 Evaluation on EvalAI

The evaluation of VizWiz competition is hosted on *www.eval.ai*, which is an open source platform for evaluating and comparing various ML and AI algorithms at scale. For this challenge, the participating teams submitted generated captions on the full VizWiz test dataset, consisting of 8,000 images. This test dataset was used to support the two main evaluation phases: **test-dev**, consisting of 4,000 test images where teams could submit at most ten times per day, and **test-challenge**, containing all 8,000 test images where teams could submit at most five times over the course of the challenge. Results from the test-challenge determined the challenge winners. For our experimental evaluations, we report the results on the **test-dev** split, while the final competition results are shown on the **test-challenge** set.

5.4 Post Processing of Captions

In order to generate final caption submissions, we incorporate various post-processing steps. Specifically, as shown in Figure 3, we apply three strategies: (a) De-tokenization, (b) self-BERT, and (c) OCR Maximization. While (a) is used to merge the sequences, generated

from the decoder, to their untokenized representations in order to make it look more plausible and human-like, (b) and (c) are strategies that rank candidate captions. For each image, we consider five candidate captions generated from multiple systems. These systems are various ensemble models which achieved high CIDEr on the development set. Note that steps (b) and (c) are used only for competition-purposes, since we observe that these steps help achieve better CIDEr on the validation set. For self-BERT, we follow a similar methodology to that of self-BLEU (Zhu et al., 2018) but we use BERT-score (Zhang et al., 2019) for better semantic similarity. The goal of self-BERT is to select a caption from candidates which is most semantically similar to all the other candidates, and thus, the higher the self-BERT score, the better. Similarly, for the OCR maximization step, we look at the coverage of tokens in generated image caption and tokens detected by the OCR modality. We notice that selecting a caption whose tokens has most overlap with the OCR tokens helps boost the CIDEr score.

5.5 Ablation Studies

In this section, we present our ablation studies to better understand the effect of each modality (IMG, OCR or OBJ) and the benefit of copy mechanism on the generated captions. We build the models following the same CE+SCST pipeline as described above, i.e., the base models are first pre-trained with CE loss followed by SCST fine-tuning. For each selected study, we vary a few model hyper-parameters, such as the number of Transformer layers, hidden dimension or the initialization seed, to finally obtain the ensemble of multiple models, which is then evaluated on the test-dev evaluation dataset. Additionally, we also show the results with and without the post-processing (Section 5.4) to examine the additional improvement that this technique has on the captions.

In Table 2, we study the effect of the data modality, which reveals that while incorporating each of the individual detectors provides gains in automatic scoring metrics above a baseline model that only considers image features, more substantial gains come from the OCR module. This finding is in line with manual inspection of the data that highlighted the prevalence of captions that include a reference to text found in the image, as noted in Section 3.2.1. Also, the results in Table 2 show that the post-processing provides an additional boost in the metrics (e.g., for CIDEr it is usually 2-3 points), which was quite valuable in this competition.

In Table 3 (top 5 rows) we study the benefit of the copy mechanism (see Section 4.2). The setup is the same as before, except that we use all the three data modalities (IMG, OCR and OBJ) and present the post-processed results. As can be seen, for a single model, there is a clear benefit of using the copy mechanism (compare lines 1 and 3 in the table). On the other hand, once the ensemble of such models is used, the advantage of copy-based models disappears (lines 2 and 4). However, combining all the copy and non-copy models together (ensemble of 90 models), which is our winning entry in the VizWiz competition, provides improvement across all metrics, especially CIDEr.

5.6 Competition Results

As discussed in Section 5.5, our best captioner is an ensemble of 90 models, employing all three data modalities (IMG, OCR and OBJ) and a mix of copy and non-copy mechanism

Table 2: Ablation studies. Each ensemble consists of 20 models (without copy mechanism) with varying number of layers, hidden dimension, and random seed. Each cell shows the value of a metric before and after post-processing the caption output. All the results are computed on the test-dev evaluation split.

	BLEU4	METEOR	ROUGE-L	CIDEr	SPICE
IMG only	25.11 / 25.97	21.01 / 21.22	48.22 / 48.55	66.8 / 69.41	15.83 / 16.04
IMG+OCR	25.88 / 26.83	21.49 / 21.73	49.17 / 49.54	72.45 / 75.51	16.23 / 16.47
IMG+OBJ	25.39 / 26.21	21.02 / 21.21	48.33 / 48.63	66.97 / 69.43	15.91 / 16.07

Table 3: Evaluation results of different sets of models that are based on all three modalities (Image, OCR and OBJ), as well as use (and omissions) of the copy mechanism. Each cell shows the value of a metric after post-processing the caption output. The top 5 rows are evaluated on test-dev split, while the bottom row is the evaluation on the test-challenge set.

Evaluation	Copy Mechanism	Model Set	BLEU4	METEOR	ROUGE-L	CIDEr	SPICE
test-dev	Copy	Single	23.66	20.34	47.34	64.48	14.81
		Ensemble (40)	25.50	21.44	49.19	72.57	16.28
	Non Copy	Single	22.10	19.71	46.36	60.90	14.70
		Ensemble (40)	27.43	22.20	50.10	78.80	17.35
	Copy + Non Copy	Ensemble (90)	27.31	22.13	50.02	80.38	17.10
test-challenge	Copy + Non Copy	Ensemble (90)	27.44	22.25	50.20	81.04	17.00

Table 4: CIDEr performance of the competition models based on the presence of text in an image. All values are taken from (Gurari & Zhao, 2020). Our captioner has a clear advantage when text is present confirming the benefits of using the OCR data modality. For images without text, we are slightly behind the SRC-B captioner.

	Ours	SRC-B	aburns	LittlePanda	iim	Baseline
with text	91.60	77.78	67.88	64.76	64.17	64.81
no text	59.80	60.43	54.90	50.37	49.7	49.01

models, followed by post-processing of the generated sentences. Evaluating this model on the test-challenge split gives the result presented at the bottom of Table 3. These test-challenge scores represent the largest values we obtained across most of the metrics, achieving the highest CIDEr score of 81.04.

Tables 4 and 5 (all the evaluation numbers are taken from Gurari and Zhao (2020)) compare our method to a few other challenge participants in terms of CIDEr score performance on the test-challenge evaluation set. In what follows, we first briefly review some of these competitor models and then discuss the table results.

Table 5: CIDEr performance of the competition models based on the image quality. All values are taken from (Gurari & Zhao, 2020). The definition of easy, medium, and hard is given the text. In all three categories, our captioner shows a clear advantage over the competing approaches.

	Ours	SRC-B	aburns	LittlePanda	iim	Baseline
Easy	88.76	77.48	68.15	64.74	64.07	64.52
Medium	66.71	63.26	56.67	51.94	51.43	51.27
Hard	41.31	39.87	35.30	33.83	33.39	31.84

Baseline (Huang et al., 2019): This is the original baseline approach used by the organizers. The captioner is based on Attention on Attention module, which refines the traditional attention mechanism to better determine the relevance between the query and the attended features.

SRC-B (Gan & Lin, 2020) [2nd place]: This is a captioner based on X-Linear Attention Networks (Pan et al., 2020) which uses Faster-RCNN image features pretrained on Visual Genome (Anderson et al., 2018b). The model is trained using the Bayesian SCST-based approach (Bujimalla et al., 2020).

aburns (Burns et al., 2020) [3rd place]: The captioner is an extension of AoANet (Huang et al., 2019) where the visual features from original attention mechanism are fed into another attention module over the OCR semantic features. These OCR tokens are first extracted using GCloud OCR Vision Cloud API and then converted into fastText feature representation.

Table 4 shows the comparison of the models’ performance on images with and without the text. Our captioner is a clear top performer when text is present confirming the benefits of using the OCR data modality. On the other hand, for images without the text, we are slightly behind the **SRC-B** captioner.

Finally, Table 5 compares the models based on the image quality. The organizers defined hard images to be 25% of all images that have a maximum CIDEr score of less than 50 across all the challenge submissions. The easy images are those that have a minimum CIDEr of 100 across all submissions, which account for 6.73% of all the image. The rest of them are deemed medium difficulty. Across all three categories of difficulty, our captioner shows a clear advantage over the competing approaches.

6. Discussion of Goal-Oriented Captioning versus Generic Captioning

In order to provide more nuanced qualitative findings, we look at captions generated by our system for two categorical scenes - indoor (Figure 8) and outdoor (Figure 10). In both figures, we select some of the images and the captions generated by our best ensemble system, V, along with the output from OCR and OBJ modalities. For comparison, and, in the spirit of approach followed by Dognin et al. (2020), we also show the outputs from models trained on a unified dataset combining MS-COCO and GCC (Sharma et al., 2018) data, with only image modality. C is the model selected with highest CIDEr on MS-COCO development dataset





	
<p>V: a person is holding a bottle of cognac C: a bottle of champagne G: a bottle of red wine OCR: cognac, liqueur, marnier, trand Obj: bottle, person</p>	<p>V: a red box of bemiks sitting on a white surface C: a sign for a parking lot G: a sign for a public parking lot OCR: bemiks, film, tablet Obj: book</p>
	
<p>V: a person is holding a box of trail mix C: a hand holding a piece of paper with a red and white background G: a hand holding a piece of cake OCR: actual, bars, net, product, shown, size, nut, bar, trail, mix Obj: book, bottle, person, pizza</p>	<p>V: a bottle of heinz tomato ketchup sitting on a counter C: a bottle of milk G: a jar of brand OCR: 2012, heinz, tomato, ketchup, ozsize Obj: bottle, dining table</p>

Figure 8: Comparison of our model (V) against generic captioners (C, G) on various images from indoor scenes.

whereas, G is the model selected with highest CIDEr on GCC development of the dataset. Note that C and G models are trained on image modality alone and serve as generic image captioning systems, whereas V is trained with all three modalities leading to more meaningful and informative captions, and thus, making it a better task-/goal-oriented system.

For the images provided in Figure 8, it can be noticed that the role of OCR modality is crucial for generating the captions. A particularly interesting case is of the caption containing the token *'bemiks'* (top right image in Figure 8). The token *'bemiks'* does not appear in the training image captions vocabulary and our model manages to use this token, detected by OCR modality, during the decoding step. This underscores the importance of copy mechanism on OCR modality for the scene. The role of OCR modality can further be supported in the ketchup bottle example (bottom right image in Figure 8), where the OCR tokens *'heinz'*, *'tomato'*, *'ketchup'* are captured by the system. Captions in Figure 10 also highlight the positive impact of the OBJ modality in our system. For example, in the bottom left image in Figure 10, our system manages to focus on the tokens generated by the OBJ modality while the C and G models fail to mention anything about the car present in the image. This is also an example which shows the more fine-grained capabilities of the

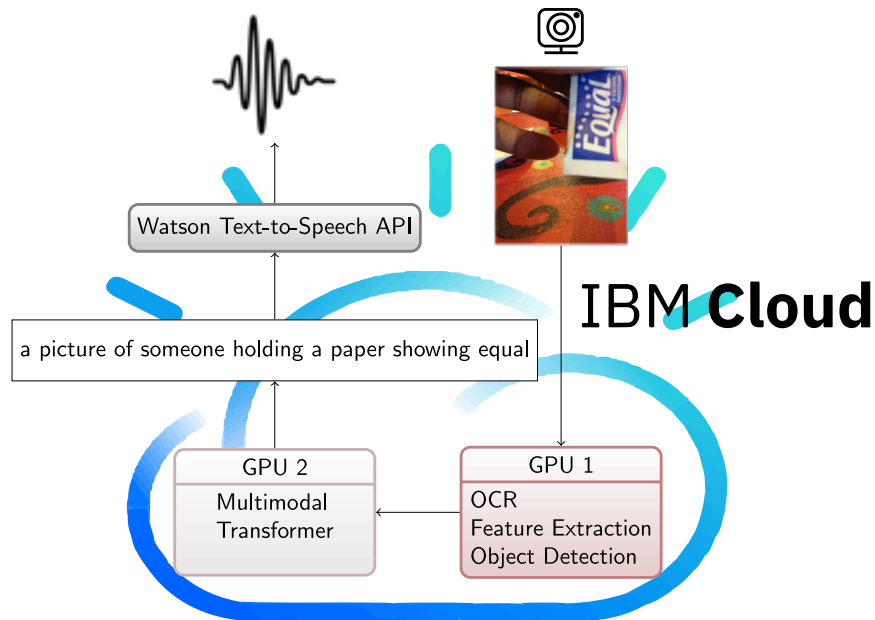


Figure 9: Assistive captioner real time demo in the cloud.

OBJ modality, which detects a 'car' even though it is barely visible in the scene and quite easy to miss.

7. Image Captioning As an Assistive Technology : Real-Time Demo

Our real-time assistive image captioning demo summarized in Figure 9 is designed as a progressive web application (PWA) with the intent to be used in a cross-platform as well as cross-device manner. Using a PWA allows the visually impaired to take as input an image captured via either a cellphone camera, webcam or uploaded locally from a variety of systems. After the image has been obtained, it is sent to the first GPU in the pipeline. Three different transforms of the image are performed. One transformation consists of feature extraction via the ResNeXt network. The next transformation is an extraction of objects defined by the object detector. Lastly, text found in the image is read from the image and extracted by the OCR module. These separate features are combined into single multimodal feature and then sent to the second GPU. This multimodal feature is then pushed into the Transformer and the captions are generated in the form of a string of text using greedy max decoding. Once completed, the generated caption string, the detected objects and texts from OCR are sent through the Watson Text-to-Speech API to be transformed into an mp3 file on a cloud server. After the mp3 file is returned, we play the audio file for the user. The audio file consists of the caption, listing of detected texts and objects. We also visualize the image, caption, objects and the words detected by the OCR to the screen (see the video of the real time demo on https://github.com/IBM/IBM_VizWiz). The realization of this





	
<p>V: a cloudy sky with a blue sky and cars in the C: a view of the house G: a storm over a city OCR: - Obj: car</p>	<p>V: a porch with a white and green flowers in front of a house C: a red flower pot on a brick wall G: a beautiful red flower in the corner of a brick wall OCR: - Obj: -</p>
	
<p>V: a street with a house and a car in front of a house C: the house is on the side of the road G: the road to the property OCR: - Obj: car</p>	<p>V: a backyard with a green grass and trees in the on the C: a picnic table in a garden G: a picnic table in a garden OCR: - Obj: bench</p>

Figure 10: Comparison of our model (V) against generic captioners (C, G) on various images from outdoor scenes.

pipeline into a real time demo was made possible using flask (Grinberg, 2018), a Python web package.

8. Conclusion

In this paper, we presented our contributions to realizing the potential of image captioning as an assistive technology for the visually impaired, focusing on our winning submission to the VizWiz 2020 Image Captioning Challenge. We described the multiple set of functionalities needed by a machine learning system to achieve meaningful task-oriented captioning, including object and text detection and recognition. Dictionary-guided rotation-invariant text recognition is a crucial component of our system, as images captured by the visually impaired may have severe quality issues in terms of orientation, blur, and occlusion. Our winning entry to the 2020 VizWiz Grand Challenge and its real time demo is a step towards translating the success of deep multimodal models on curated datasets, such as MS-COCO, to useful assistive technology. A more interactive system, in the form of visual dialog, would be an important next step, as it would allow for rounds of feedback between the

visually impaired users and the system in order to facilitate users in reaching their goals and completing tasks at hand. This type of dialog will help with occlusions and bad image conditions that the system would eventually detect, e.g. by asking for additional images from the user, and would refine the system response in case a user was not satisfied with a provided caption. Applications of computer vision in accessibility and science for social good will require this feedback loop as well as the goal-oriented components we highlighted throughout this work. Note that other modalities and capabilities can further augment and improve the system such as improved color detection, counting and reasoning capability via incorporation of a scene graph, or location and navigation capability via addition of GPS coordinates. Beyond vision, other wearable sensors such as LIDAR or sonar technologies could further improve the assistive technology.

Acknowledgments

All authors contributed to the writing of the paper. Authors Pierre Dognin, Igor Melnyk, Youssef Mroueh, Inkit Padhi, Mattia Rigotti, Jarret Ross, and Yair Schiff contributed *equally* in the 2020 VizWiz Grand Challenge and are listed in alphabetical order. Authors Richard A. Young, and Brian Belgodere contributed to the real-time demo and the baseline generic captioner system.

The authors would like to thank the support of the science for social good program in IBM Research AI, and the MIT-IBM Watson AI lab for providing cloud machines for the real-time demo. Authors are also thankful to the 2020 VizWiz Grand Challenge competition organizers, in particular Danna Gurari and Yinan Zhao.

References

- Anderson, P., Fernando, B., Johnson, M., & Gould, S. (2016). Spice: Semantic propositional image caption evaluation. In *European conference on computer vision* (pp. 382–398).
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., & Zhang, L. (2018a). Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*.
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., & Zhang, L. (2018b). 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).
- Baek, J., Kim, G., Lee, J., Park, S., Han, D., Yun, S., ... Lee, H. (2019). What is wrong with scene text recognition model comparisons? dataset and model analysis. In *International conference on computer vision (ICCV)*.
- Baek, Y., Lee, B., Han, D., Yun, S., & Lee, H. (2019). Character region awareness for text detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 9365–9374).
- Banerjee, S., & Lavie, A. (2005). 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* (Vol. 29, pp. 65–72).
- Bigham, J. P., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R. C., . . . Yeh, T. (2010). Vizwiz: Nearly real-time answers to visual questions. In *Proceedings of the 23rd annual acm symposium on user interface software and technology* (p. 333–342). New York, NY, USA: Association for Computing Machinery.
- Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2016). Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606*.
- Brady, E., Morris, M. R., Zhong, Y., White, S., & Bigham, J. P. (2013, April). Visual challenges in the everyday lives of blind people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2117–2126). Association for Computing Machinery.
- Bujimalla, S., Subedar, M., & Tickoo, O. (2020). B-SCST: Bayesian Self-Critical Sequence Training for Image Captioning. *arXiv preprint arXiv:2004.02435*.
- Burns, A., Saenko, K., & Bryan, P. (2020). *Feature Refinement for Common Sense Captioning*. Boston University, "https://ivc.ischool.utexas.edu/~yz9244/VizWiz_workshop/videos/aburns-oral.mp4".
- Chiu, T.-Y., Zhao, Y., & Gurari, D. (2020, June). Assessing image quality issues for real-world problems. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (cvpr)*.
- Dai, B., Lin, D., Urtasun, R., & Fidler, S. (2017). Towards diverse and natural image descriptions via a conditional GAN. *International Conference on Computer Vision (ICCV)*.
- Dognin, P., Melnyk, I., Mroueh, Y., Padhi, I., Rigotti, M., Ross, J., & Schiff, Y. (2020). Alleviating noisy data in image captioning with cooperative distillation. In *Vizwiz grand challenge workshop, cvpr*. (https://ivc.ischool.utexas.edu/~yz9244/VizWiz_workshop/poster_pdf/vizwiz_workshop-co-distill.pdf)
- Dognin, P., Melnyk, I., Mroueh, Y., Ross, J., & Sercu, T. (2019). Adversarial semantic alignment for improved image captions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 10463–10471).
- Fisch, A., Lee, K., Chang, M.-W., Clark, J. H., & Barzilay, R. (2020). CapWAP: Captioning with a Purpose. *ACL 2020*. Retrieved 2020-12-01, from <http://arxiv.org/abs/2011.04264>
- Gan, Z., & Lin, K. (2020). *Box Features Masking and Bayesian Self-Critical Sequence Training for VizWiz Captions Challenge*. Samsung Research China-Beijing, "https://ivc.ischool.utexas.edu/~yz9244/VizWiz_workshop/videos/SRC-B_VCLab.mp4".
- Gleason, C., Pavel, A., McCamey, E., Low, C., Carrington, P., Kitani, K. M., & Bigham, J. P. (2020). Twitter ally: A browser extension to make twitter images accessible. In *Proceedings of the 2020 chi conference on human factors in computing systems* (p. 1–12). Association for Computing Machinery.
- Grinberg, M. (2018). *Flask web development: developing web applications with python*. "O'Reilly Media, Inc."
- Gu, J., Lu, Z., Li, H., & Li, V. O. (2016, August). Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 1631–1640).

- Berlin, Germany: Association for Computational Linguistics. Retrieved from <https://www.aclweb.org/anthology/P16-1154> doi: 10.18653/v1/P16-1154
- Guinness, D., Cutrell, E., & Morris, M. R. (2018). Caption crawler: Enabling reusable alternative text descriptions using reverse image search. In *Proceedings of the 2018 chi conference on human factors in computing systems* (p. 1–11). Association for Computing Machinery.
- Gurari, D., Li, Q., Stangl, A. J., Guo, A., Lin, C., Grauman, K., ... Bigham, J. P. (2018, June). Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition (cvpr)*.
- Gurari, D., & Zhao, Y. (2020). *Vizwiz Grand Challenge:Caption Images Taken by People Who Are Blind*. "https://ivc.ischool.utexas.edu/~yz9244/VizWiz_workshop/slides/zhao-talk.pdf".
- Gurari, D., Zhao, Y., Zhang, M., & Bhattacharya, N. (2020). Captioning images taken by people who are blind. In A. Vedaldi, H. Bischof, T. Brox, & J.-M. Frahm (Eds.), *Computer vision – eccv 2020* (pp. 417–434). Springer International Publishing.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Huang, L., Wang, W., Chen, J., & Wei, X.-Y. (2019). Attention on Attention for Image Captioning. *arXiv preprint arXiv:1908.06954*.
- IBM Research. (2021). *Watson Text to Speech*. "https://www.ibm.com/cloud/watson-text-to-speech". (Accessed: 2021-06-21)
- IBM Research and CMU. (2021). *NavCog AI Suite case*. "https://www.cs.cmu.edu/~NavCog/navcog.html". (Accessed: 2021-06-21)
- Kacorri, H., Kitani, K. M., Bigham, J. P., & Asakawa, C. (2017). People with visual impairment training personal object recognizers: Feasibility and challenges. In *Proceedings of the 2017 chi conference on human factors in computing systems* (pp. 5839–5849).
- Kant, Y., Batra, D., Anderson, P., Schwing, A., Parikh, D., Lu, J., & Agrawal, H. (2020). Spatially Aware Multimodal Transformers for TextVQA. *arXiv preprint arXiv:2007.12146*.
- Karpathy, A., & Li, F.-F. (2015). Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Lin, C.-Y. (2004, July). ROUGE: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the acl-04 workshop* (pp. 74–81). Association for Computational Linguistics.
- Lin, T., Maire, M., Belongie, S. J., Bourdev, L. D., Girshick, R. B., Hays, J., ... Zitnick, C. L. (2014). Microsoft COCO: Common Objects in Context. *EECV*.
- Liu, S., Zhu, Z., Ye, N., Guadarrama, S., & Murphy, K. (2017). Improved image captioning via policy gradient optimization of spider. In *International conference on computer vision (iccv)*.
- Lu, J., Yang, J., Batra, D., & Parikh, D. (2018). Neural baby talk. In *Proceedings of the IEEE conference on computer vision and pattern recognition (cvpr)* (p. 7219–7228). IEEE Computer Society.
- Luo, R., Vered, G., Bracha, L., Chechik, G., & Shakhnarovich, G. (2019). *Winning Google Conceptual captions challenge*. <https://ttic.uchicago.edu/~rluo/files/>

- ConceptualWorkshopSlides.pdf.
- Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., ... van der Maaten, L. (2018). *Exploring the limits of weakly supervised pretraining*. Retrieved from <http://arxiv.org/abs/1805.00932>
- Microsoft AI. (2021). *Seeing AI Mobile App*. "https://www.microsoft.com/en-us/ai/seeing-ai". (Accessed: 2021-06-21)
- Morris, M. R., Johnson, J., Bennett, C. L., & Cutrell, E. (2018). Rich representations of visual content for screen reader users. In *Proceedings of the 2018 chi conference on human factors in computing systems* (p. 1–11). Association for Computing Machinery.
- Mroueh, Y., Voinea, S., & Poggio, T. A. (2015). Learning with group invariant features: A kernel perspective. In *Advances in neural information processing systems 28*.
- OrCam. (2021). *OrCam Wearable Device*. "https://www.orcam.com/en/about/". (Accessed: 2021-06-21)
- Pan, Y., Yao, T., Li, Y., & Mei, T. (2020). X-Linear Attention Networks for Image Captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10971–10980).
- Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the association for computational linguistics* (pp. 311–318).
- Ranzato, M., Chopra, S., Auli, M., & Zaremba, W. (2015). Sequence level training with recurrent neural networks. *ICLR*.
- Redmon, J., & Farhadi, A. (2016). *YOLO9000: Better, Faster, Stronger*.
- Rennie, S. J., Marcheret, E., Mroueh, Y., Ross, J., & Goel, V. (n.d.). Self-critical sequence training for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), year = 2017,*
- See, A., Liu, P. J., & Manning, C. D. (2017). Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th annual meeting of the association for computational linguistics (volume 1: Long papers)* (pp. 1073–1083).
- Sharma, P., Ding, N., Goodman, S., & Soricut, R. (2018). Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of acl*.
- Shetty, R., Rohrbach, M., Hendricks, L. A., Fritz, M., & Schiele, B. (2017). Speaking the same language: Matching machine to human captions by adversarial training. *arXiv preprint arXiv:1703.10476*.
- Simons, R. N., Gurari, D., & Fleischmann, K. R. (2020, October). "I Hope This Is Helpful": Understanding Crowdworkers' Challenges and Motivations for an Image Description Task. *Proc. ACM Hum.-Comput. Interact.*, 4(CSCW2).
- Singh, A., Natarajan, V., Shah, M., Jiang, Y., Chen, X., Batra, D., ... Rohrbach, M. (2019, June). Towards VQA Models That Can Read. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Stangl, A., Morris, M. R., & Gurari, D. (2020). "person, shoes, tree. is the person naked?" what people with vision impairments want in image descriptions. In *Proceedings of the 2020 chi conference on human factors in computing systems* (p. 1–13). New York, NY, USA: Association for Computing Machinery.
- Tan, M., Pang, R., & Le, Q. V. (2020). *Efficientdet: Scalable and efficient object detection*.

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., . . . Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems 30* (pp. 5998–6008).
- Vedantam, R., Lawrence Zitnick, C., & Parikh, D. (2015). Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4566–4575).
- Vedantam, R., Zitnick, C. L., & Parikh, D. (2015). Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Retrieved from <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2015.html#VedantamZP15>
- Vinyals, O., Toshev, A., Bengio, S., & Erhan, D. (2015). Show and Tell: A Neural Image Caption Generator. *Pattern Analysis and Machine Intelligence, IEEE Transactions on.*
- Vinyals, O., Toshev, A., Bengio, S., & Erhan, D. (2017, Apr). Show and Tell: Lessons Learned from the 2015 MSCOCO Image Captioning Challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 39*(4), 652–663. Retrieved from <http://dx.doi.org/10.1109/TPAMI.2016.2587640> doi: 10.1109/tpami.2016.2587640
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning, 8*(3-4), 229–256.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A. C., Salakhutdinov, R., . . . Bengio, Y. (2015). Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *Icml*.
- Xu, K., Ba, J., Kiros, R., Courville, A., Salakhutdinov, R., Zemel, R., & Bengio, Y. (2015). Show, attend and tell: Neural image caption generation with visual attention. *arXiv preprint arXiv:1502.03044*.
- Yao, T., Pan, Y., Li, Y., & Mei, T. (2017). Incorporating copying mechanism in image captioning for learning novel objects. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6580–6588).
- Zeng, X., Wang, Y., Chiu, T.-Y., Bhattacharya, N., & Gurari, D. (2020, October). Vision skills needed to answer visual questions. *Proc. ACM Hum.-Comput. Interact., 4*(CSCW2). Retrieved from <https://doi.org/10.1145/3415220> doi: 10.1145/3415220
- Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., & Artzi, Y. (2019). Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J., & Yu, Y. (2018). Texus: A benchmarking platform for text generation models. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (pp. 1097–1100).