

Concadia: Towards Image-Based Text Generation with a Purpose

Elisa Kreiss¹

Fei Fang²

Noah D. Goodman^{2,3}

Christopher Potts¹

¹Department of Linguistics ²Department of Computer Science ³Department of Psychology
Stanford University
Stanford, CA 94305 USA
{ekreiss, feifang, ngoodman, cgpotts}@stanford.edu

Abstract

Current deep learning models often achieve excellent results on benchmark image-to-text datasets but fail to generate texts that are useful in practice. We argue that to close this gap, it is vital to distinguish *descriptions* from *captions* based on their distinct communicative roles. *Descriptions* focus on visual features and are meant to replace an image (often to increase accessibility), whereas *captions* appear alongside an image to supply additional information. To motivate this distinction and help people put it into practice, we introduce the publicly available Wikipedia-based dataset **Concadia** consisting of 96,918 images with corresponding English-language descriptions, captions, and surrounding context. Using insights from Concadia, models trained on it, and a preregistered human-subjects experiment with human- and model-generated texts, we characterize the commonalities and differences between descriptions and captions. In addition, we show that, for generating both descriptions and captions, it is useful to augment image-to-text models with representations of the textual context in which the image appeared.

1 Introduction

Image-based natural language generation (NLG) has great potential for productive applications in accessibility, image search, creative writing, and navigational instruction tasks, among other areas. However, current systems fall short of realizing this potential even though they can generally produce fluent, truthful texts in a wide range of scenarios (e.g., Dognin et al., 2022; Guinness et al., 2018; Gurari et al., 2020). In this paper, we identify and address one of the fundamental limitations holding these systems back: insufficient attention to the *communicative purpose* of the text being generated.

We focus on the high-level distinction between *descriptions* and *captions*. Descriptions are created for the purpose of replacing an image in context,

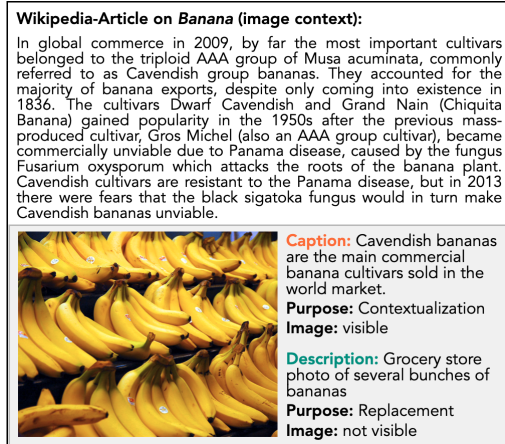


Figure 1: An example from our Wikipedia-based corpus Concadia of an image from the article on bananas with two associated texts: the description as provided in the image’s alt text, and the caption as displayed below the image in the article. For captions, the image content is presupposed whereas descriptions aim to stand in for the image. While grounded in the same image, description and caption convey vastly different information.

most notably used to make an image accessible to users who can’t see them. Captions, in contrast, presuppose access to the image and seek to provide additional information. For example, in Figure 1, the description summarizes the scene and describes the features that are clearly visible in the image, whereas the caption provides almost no visual details and instead highlights the relevance of this image to the article. Both texts are grounded in the same image and have clear purposes, but neither can play the role of the other.

To enable in-depth exploration of the description/caption distinction, we introduce **Concadia**, a Wikipedia-based corpus of 96,918 images with associated English-language captions, alt text descriptions, and accompanying context from the respective Wikipedia article. Concadia provides the opportunity to directly compare descriptions and captions for shared images, and it can be used to

develop image-based NLG systems that serve these distinct communicative purposes as well.

To further substantiate the description/caption distinction, we report on a linguistic analysis of the commonalities and differences between these two kinds of text in Wikipedia (Section 3). We then develop and assess image-based NLG systems using Transformer- and LSTM-based architectures. Two clear lessons emerge from these modeling efforts. First, description generation is easier than caption generation, as one might expect given the tighter relationship between descriptions and images. Second, it is highly beneficial to include representations of the textual context in which an image appeared. This finding aligns with much recent work showing that blind and low-vision (BLV) users value image descriptions that tie the image to its surrounding context (Stangl et al., 2021, 2020; Muehlbradt and Kane, 2022). Similarly, Biten et al. (2019) identify benefits of using context in the caption domain, for instance, for inferring names of people visible in an image. In the example in Figure 1, both the caption and description reflect the context (commercial aspects of bananas), and the caption additionally connects to the type of banana discussed in the article. Most datasets for image-based NLG provide only image–text pairs, with no context, which may fundamentally limit their usefulness. A strength of Concadia is that all the image–text pairs are given within their Wikipedia contexts.

Finally, we provide evidence from a preregistered human-subjects experiment (Section 5). Our results further reinforce the observation that descriptions and captions serve distinct communicative purposes. In addition, the results hold not only for the descriptions and captions from Concadia but extend to our model-generated texts as well. This serves to further validate that systems trained to generate captions are inadequate for the task of generating descriptions, and vice versa, and that model evaluations need to be sensitive to communicative purpose. Together with the modeling results, this speaks to the importance of both communicative purpose and context when it comes to developing effective image-based NLG systems.

2 Captions vs. Descriptions in Prior Work

In AI, the task of generating prose text from images is often referred to as “image captioning”, which glosses over distinctions of communicative purpose.

The current section offers a synopsis of research, within and outside of AI, that is relevant to our proposed description/caption distinction and supports a more nuanced view of image-based NLG tasks.

The conceptual description/caption distinction we propose poses a terminological challenge. We use “caption” to mean “image-supporting information expressed in text”, which we contrast with “(alt) description” to mean “text intended to replace an image”. This is aligned with general popular usage and with the terminology in (photo-)journalism and graph accessibility, fields that conceptually already make a description/caption distinction. However, the term “caption” is now a loaded term in AI, and future work should be careful in considering the conceptual distinctions despite the challenges in terminology.

2.1 Image-to-Text Relations

Most previous work has focused on characterizing the relation between images and the texts that occur alongside them. Marsh and Domas White (2003) summarize cross-disciplinary efforts, and propose a unifying taxonomy based on the different functions an image takes on in a text (e.g., from decorative through explanatory to interpretative).

Work in communication sciences situates image–text relations based on their role in discourse (e.g., Martinec and Salway, 2005), and work in library sciences investigates the role that these relations play in efficient image indexing (e.g., Jaimes and Chang, 1999; Shatford, 1986).

More recently, the relations of images and visually co-occurring texts have been explored in AI, specifically for news (Otto et al., 2020; Oostdijk et al., 2020), advertising (Otto et al., 2020; Zhang et al., 2018), and social media (Kruk et al., 2019; Hodosh et al., 2013). Alikhani et al. (2020) introduce an annotation protocol to investigate coherence relations between images and associated texts. They analyze to what extent texts from a variety of resources address what’s *visible*, *subjective*, representing an *action*, telling a *story*, or containing *meta* information. Our description vs. caption distinction is focused on communicative purpose and can cross-cut these categories. For instance, evidence from BLV users suggests that a description intended to replace an image of politicians in a news article should include visible features that describe the scene, but also meta information that names the politicians and the actions they’re

engaging in (Stangl et al., 2021).

A majority of the previously proposed categorizations are based on analyses of the text content, and most generally assume that images and texts visually co-occur. The description vs. caption distinction is a higher-level categorization that is based on whether the image content is presumed to be visible in the first place. This is motivated by communicative principles: when the image is visible, texts that simply describe the image would be redundant, and instead those texts need to serve distinct communicative purposes (Kruk et al., 2019). Hodosh et al. (2013) follow similar reasoning to explain why crowd-sourced image–text datasets at the time contained high rates of information that can’t be directly extracted from the image. This distinction is particularly pronounced in the area of image accessibility. For instance, recent studies on the accessibility of computational research papers explicitly distinguish figure captions and figure descriptions when assessing a paper’s accessibility (Williams et al., 2022; Chintalapati et al., 2022).

2.2 Image-Based Text Generation Datasets

We can loosely categorize existing datasets according to communicative purpose. *Caption datasets* contain texts that are useful when appearing alongside images, mainly to contextualize them. *Description datasets* contain texts that are intended to replace images. Some datasets are difficult to decisively place due to uncertainty about the nature of the texts or major postprocessing of the data.

2.2.1 Description Datasets

The first group of datasets were created by presenting sighted participants with images and the task to describe them. Presenting the images out of context restricts the text to information that can be extracted from the image itself. We therefore consider these to be *description* datasets. Examples include MS-COCO (Lin et al., 2014; Chen et al., 2015), Flickr8k/30k (Young et al., 2014; Hodosh et al., 2013), and more recently VizWizCaptions (Gurari et al., 2020) and TextCaps (Sidorov et al., 2020). Localized Narratives (Pont-Tuset et al., 2020) and DIDE (van Miltenburg et al., 2018) are similarly constructed, except that the descriptions are spoken and the datasets further contain real-time data aligning regions of the image with the description.¹

¹See van Miltenburg et al. 2018 for a comparison of spoken vs. written image descriptions.

Conceptual Captions (Sharma et al., 2018) and its extension Conceptual 12M (Changpinyo et al., 2021) represent a second group of description datasets. These datasets contain crawled images and their associated alt texts from the Web. A primary purpose of alt texts is to replace the images when they’re absent, and they are what screen readers use for providing non-visual image access. These use-cases align with the assumption that the text stands in for the image, which makes them description datasets.

None of these datasets contain information on the context an image is situated in. In the case of crowd-sourced descriptions, no context was provided, which means that the description content solely relies on cognitive priors and heuristics to determine what is relevant. In the crawled datasets, we can generally assume that the people writing alt texts were aware of the context the images appear in, suggesting that the descriptions are likely to be context-sensitive. However, the context itself is often not preserved in the datasets, meaning that the effect of context can’t be further investigated or learned by a model.

2.2.2 Caption Datasets

Some recent datasets focus on texts expressing information that can’t be extracted from the image alone. The news domain is a prominent example (Biten et al., 2019; Ramisa et al., 2017), since the texts associated with images in newspapers generally presuppose that the user can see the image. Thus, we consider these to be *caption* datasets. These datasets generally do contain contextual information in the form of the articles the images appear in.

2.2.3 Other Datasets

Im2Text is a dataset crawled from Flickr, containing 1 million images and the user-generated text appearing alongside it. In this case, the purpose of the text is difficult to define since this same text is used for image search as well as providing information that might be read alongside the image. Recognizing this challenge, the dataset was post-processed to make the text appear more descriptive. However, subsequent work has suggested that the dataset still contains a large number of non-descriptive texts (Hodosh et al., 2013), making it difficult to decisively place.

While often not summarized in a single static dataset, Twitter data is also used for image-based

split	datapoints	unique articles	avg length (words)	avg word length	vocab size
train	77,534	31,240	caption: 12.79 description: 14.53	caption: 6.07 description: 5.69	caption: 82,745 description: 54,893
dev	9,693	4,952	caption: 12.91 description: 14.57	caption: 6.10 description: 5.70	caption: 24,063 description: 16,800
test	9,691	4,951	caption: 12.67 description: 14.29	caption: 6.10 description: 5.71	caption: 23,650 description: 16,610
all	96,918	41,143	caption: 12.79 description: 14.51	caption: 6.07 description: 5.69	caption: 96,908 description: 63,884

Table 1: Overview of our Wikipedia-based caption and description dataset Concadia. Captions are on average shorter than descriptions but contain longer words. Captions make use of a larger vocabulary.

text generation (e.g., Hessel et al., 2021). The Twitter API offers access to the image as well as the tweet and, if present, a user-written image description to make it non-visually accessible. The tweet assumes that a user has access to the posted image which aligns it with the captioning purpose. The accessibility text seeks to substitute it, aligning with the description purpose.

The WIT dataset contains crawled images with their corresponding captions (which are called *reference descriptions* in WIT), alt texts, and the surrounding paragraph from Wikipedia (Srinivasan et al., 2021).² WIT provides a rich multilingual resource for images and associated texts in context. While it contains captions as well as alt descriptions for each image within a context, the authors note that the descriptions are rare and often contain only the filename. Perhaps for that reason, the subsequent analyses and experiments were not performed on the alt description data. Concadia is specifically created for capturing the alt descriptions, and is controlled and filtered for their quality.

3 The Concadia Dataset

We introduce **Concadia**, **Contextualized captions and alt descriptions** from **Wikipedia**, a corpus extracted from Wikipedia consisting of images with their naturally occurring alt descriptions, captions, and surrounding context. Concadia allows for a direct comparison of descriptions and captions for a specific image in a specific context. In this section, we characterize the distribution of captions and descriptions on all of Wikipedia, describe the Concadia dataset generation process, and provide a qualitative comparison on the similarities and dif-

ferences of the captions and descriptions in Concadia. Appendix A provides implementation details and further illustrations.

3.1 Descriptions and Captions on Wikipedia

Descriptions and captions fulfill distinct communicative purposes that are already apparent in their distribution. Where present, captions are accessible for all readers, often printed below their respective image. Image descriptions are harder to access, since they’re usually defined in the image’s alt tag. Although hidden to most users, alt descriptions are essential for making images accessible to users who can’t see them. The fact that alt tags are hidden by default likely decreases awareness and use of the field (Gleason et al., 2019), resulting in major accessibility challenges specifically for blind and low vision users. The scarcity of alt descriptions online that could make these images nonvisually accessible is widely established, but existing studies have focused on social media (Morris et al., 2016) and most frequently visited websites overall (Bigham et al., 2006; Guinness et al., 2018). We expect to see a similar pattern on Wikipedia.

In 2020, English Wikipedia consisted of over 6 million articles with more than 4 million images.³ Based on 10,000 randomly sampled articles, our approximation suggests that while around 91% of all images on Wikipedia are associated with a caption, only 6% contain alt description tags. Moreover, this small subset includes alt descriptions that simply say *alt text* or *Image*. The number of informative alt descriptions is likely much lower.

²They also extract the image’s *Attribution Description*, which is a context-independent description of the image from Wikimedia, used for image search.

³This number is only an estimate, since confirmed numbers aren’t available.

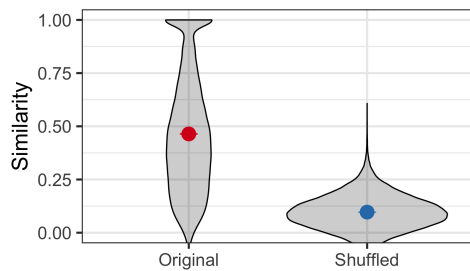


Figure 2: Cosine similarity between SBert embeddings of corresponding descriptions and captions (left, in red) and randomly sampled descriptions and captions (right, in blue) from Concadia. The grey areas visualize the data distribution and the points mark average similarity with 95% bootstrapped confidence intervals.

3.2 Dataset Curation

For Concadia, we extracted all images from Wikipedia that have captions as well as alt descriptions. Images were excluded where the picture wasn't publicly available at Wikimedia Commons, descriptions contained reference to the caption (e.g., *refer to caption*), consisted of fillers (e.g., *alt text*, *text*), only consisted of the file name, or where any associated text was shorter than 3 characters. The extracted context paragraph is the paragraph following the image in the HTML, and is minimally 20 words long.

Table 1 provides basic corpus statistics. The final corpus consists of 96,918 images with descriptions, captions, and surrounding text from 41,143 articles. Additionally, we include a train/dev/test split for the data. Images that occur multiple times (but have potentially different associated texts), and images where caption and description are identical are sorted into the training set to ensure the highest quality in the validation and test sets. All other assignments are random, while ensuring that data from the same article are assigned to the same split. The dataset is publicly available.⁴

3.3 Caption/Description Similarities

Concadia provides the opportunity to directly compare descriptions and captions naturally written by Wikipedia users for a single image. If captions and descriptions show perfect similarity, there is no reason to assume a difference in communicative purpose, and captions could for instance simply replace the rarely present descriptions when they are absent. A complete dissimilarity would indicate a major effect of communicative purpose, and would

suggest few commonalities that could be induced by the shared image and context. Crucially, if the content of captions and descriptions is *partially* related, this would be compatible with the texts fulfilling distinct purposes while still being shaped by their shared image and context.

To investigate the semantic similarity between captions and descriptions, we computed the cosine similarity of the SBert embeddings (Reimers and Gurevych, 2019) for all matching description/caption pairs (see Figure 2 in red). Descriptions and captions are significantly more similar than would be assumed under a baseline where descriptions and captions are randomly paired (in blue; two-sample Kolmogorov-Smirnov test: $p < 0.0001$), suggesting that they share item-specific content. (We replicate this using Jaccard distance of the raw strings, as presented in Appendix A.3.)

Having established that there is a semantic similarity between captions and descriptions, we turn to how they come apart. A quantitative analysis of the Parts-of-Speech (POS) shows distinct patterns for descriptions and captions. Adjectives and nouns are significantly more frequent in descriptions but captions contain significantly more proper nouns. This aligns with previous work on news captions (Biten et al., 2019), which found that proper nouns contribute 20% of captions from New York Times articles while they're completely absent in datasets such as MS-COCO (a description dataset). This similarly replicates for the adjective and noun distributions (Biten et al., 2019; Ramisa et al., 2017). Intuitively, these patterns are already reflected in the most frequent bigrams on Concadia (see Figure 3). Descriptions are dominated by descriptive attributes such as people's looks (e.g., *white shirt*, *baseball cap*, *dark hair*), and meta information about the image (e.g., *colour photograph*). The most frequent bigrams in captions are dominated by proper noun compounds such as *San Francisco* or *Tour de France*, as well as common places and times (e.g., *national park*, *19th century*).

4 Model Experiments

Our investigation of the similarity between captions and descriptions in Section 3.3 indicates that they differ in principled ways but that they are not unrelated. This suggests that, while captions and descriptions can't stand in for each other, they can potentially inform each other's content when we automatically generate them. We further motivated

⁴<https://github.com/elisakreiss/concadia>

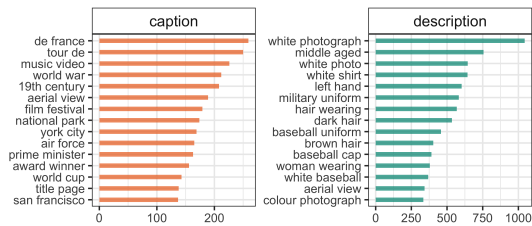


Figure 3: Most frequent bigrams, excluding stopwords (i.e., highly frequent words such as pronouns, prepositions, and forms of *to be*).

that there is evidence that both description and caption generation would benefit from integrating the larger context where the image occurred. We now pursue these hypotheses in the context of large-scale caption and description generation.

4.1 Model Architectures

Previous work in image-based NLG has shown promising results from using both recurrent and Transformer-based models to generate text. To investigate the importance of a description vs. caption distinction and the relevance of context, we compare three model variants. Depending on the condition, the models are trained on predicting description or caption labels and receive as context captions, descriptions, the broader context, or nothing. We provide further details on the models and training in Appendix B.

ResNet-LSTM generates a label token-by-token, using an LSTM decoder (Hochreiter and Schmidhuber, 1997). The input to the LSTM is a concatenated representation of the image using a pretrained ResNet (He et al., 2016), and the context using BERT embeddings (Devlin et al., 2019). We took inspiration from the importance of attention mechanisms (Vaswani et al., 2017) on the input image as used in state-of-the-art “captioning” models (Xu et al., 2015; Lu et al., 2017; Anderson et al., 2018) and extended them to the context input. To ensure the same number of trainable parameters between the with-context vs. no-context models, the no-context models received all-ones vectors in place of context embeddings.

DenseNet-LSTM follows the same structure as the ResNet-LSTM but leverages pretrained DenseNet features instead (Deng et al., 2020; Hosain et al., 2021).

OSCAR(VinVL) is a pretrained Transformer-based vision-language model that achieves state-of-the-art results on many tasks via finetuning (Zhang et al., 2021). For each image, the VinVL pretrained

visual feature extractor provides visual features, i.e., vector embeddings of object-enclosing regions, and object tags in text. To obtain a unified representation of the image–text pair, the OSCAR model concatenates the BERT embeddings (Devlin et al., 2019) of the text and object tags with the visual features (Li et al., 2020). We supply context by appending it to the object tags before the BERT encoding. Since BERT can handle variable-length inputs, the trainable parameters remain constant regardless of whether context is incorporated.

4.2 Evaluation Metrics

Since Concadia is extracted from naturally occurring data, each image is only associated with a single ground-truth caption and description. Currently established evaluations of text generated from images rely on multiple references to reliably estimate performance. As suggested in previous work (Biten et al., 2019), we consider CIDEr (Vedantam et al., 2015) to be most appropriate for our setting. First, CIDEr has been shown to outperform BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) specifically when reference sentences are sparse (Vedantam et al., 2015; Anderson et al., 2016). Second, while METEOR (Denkowski and Lavie, 2014) and SPICE (Anderson et al., 2016) achieve slightly higher alignment with human judgments on some datasets with few references, their use of soft-similarity and lemmatization is not well-defined for proper names (Biten et al., 2019), which make up more than 26% of our caption data and more than 15% of our description data (Section 3.3).

4.3 Experiment 1

In our first experiment, we probe how the distinction between the tasks of description and caption generation affects model performance.

Predictions If descriptions are meant to replace an image while captions are meant to complement it, we expect models to perform better at generating descriptions from image input alone, as compared to captions. In Experiment 1, we test whether this prediction is borne out in the three models. A high performance gap between the two conditions would further highlight the necessity of differentiating those tasks to address their distinct challenges.

Results Table 2 shows the performance on the test set for description and caption generation. Test set evaluation was performed using beam search ($n = 5$) for all models. As predicted, all models achieve higher performance on the description data

Label	Context	RN-LSTM	DN-LSTM	OSC.VinVL
Descr.	None	0.16 (0.00)	0.21 (0.00)	0.21 (0.00)
Caption	None	0.07 (0.00)	0.10 (0.01)	0.10 (0.01)
Descr.	Rand. Par.	0.14	0.19	0.22
Caption	Rand. Par.	0.06	0.09	0.10
Descr.	Paragr.	0.20 (0.00)	0.23 (0.01)	0.32 (0.01)
Caption	Paragr.	0.13 (0.00)	0.15 (0.00)	0.45 (0.00)
Descr.	Caption	0.27 (0.01)	0.28 (0.00)	1.18 (0.01)
Caption	Descr.	0.17 (0.01)	0.19 (0.00)	1.14 (0.01)

Table 2: CIDEr scores for all models: ResNet-LSTM, DenseNet-LSTM and OSCAR(VinVL). Standard deviations (in brackets) are computed over 3 random seeds. Across models, description generation results in higher CIDEr scores and adding contexts boosts performance.

than the caption data. This suggests that the information from the image alone is more helpful to description generation, suggesting a closer connection between the two. Experiments using shuffled versions of the dataset further show that the performance gap can't be explained simply by the smaller vocabulary in descriptions, the higher proportion of proper nouns in captions, or the higher n-gram overlap across descriptions (see Section 4.5).

4.4 Experiment 2

We have seen that descriptions and captions play different roles and pose different modeling challenges. Can these two kinds of text mutually reinforce each other in the context of automatic text generation? We investigate this question by treating one as the context that can inform generation for the other. For a second notion of context, we also test whether the closely occurring paragraph can help the models' performances. A positive effect of context would point to the value of treating image-based NLG as a contextual problem.

Predictions We predict that, while distinct in purpose, descriptions and captions from the same image and context share enough information that they can benefit from one another. Part of this contextual relevance should also be captured when using the closely appearing paragraph. To ensure that not just language in general accounts for potential performance gains, we provide control conditions where the model receives paragraphs from other images. Improvements over these controls can then only be due to semantic relatedness between the context and the generated description/caption.

Results Table 2 shows that, where models are provided contextual representations, performance

improves for both description and caption generation compared to the image-only input baselines. As predicted, this improvement is most pronounced when providing the respective caption or description. Even the closely appearing paragraph leads to increased CIDEr scores compared to the no-context baselines. Furthermore, this improvement goes beyond what models achieved when training on randomly assigned paragraphs, which suggests that they make use of item-specific information and points to the more general benefit of integrating the broader context an image appears in.

4.5 Discussion

Across models, we find that they achieve higher CIDEr ratings on the description than the caption generation task (see Table 2). In part this might be due to the smaller set of tokens present in descriptions (Table 1), the higher proportion of proper nouns in captions (Figure 7), and the higher n-gram overlap across descriptions (Figure 3). An additional potential explanation for the gap is that descriptions are more tightly connected to their image than captions, and are therefore easier to learn.

The first set of factors are image-independent, i.e., they are simply properties of the language labels. If these components solely drive the higher performance on the description data, this gap should be reflected in a version of the dataset where the image-label mappings are shuffled. We conducted this experiment using the ResNet-LSTM model instead of the OSCAR(VinVL) model since the OSCAR(VinVL) model already carries biases from its pretraining. Although the shuffled description model achieves higher CIDEr scores than the shuffled caption model (0.017 vs. 0.016, respectively), it doesn't compare to the performance gain on the non-shuffled data (0.09). The tighter relationship of descriptions and their images predicts this gap and falls in line with the general description/caption purpose distinction.

Further support for the purpose of the description/caption directly affecting their predictability comes from the experiment where we provided the closely appearing paragraph as the context for the images (Table 2). In the best-performing model, OSCAR(VinVL), caption generation performance surpasses description performance when the paragraph is provided. Furthermore, while caption performance doesn't surpass description generation performance in the LSTM models, providing the

paragraph as context to the caption models results in a higher performance increase than for the description models. This is intuitive given the predicted connection of the descriptions/captions to the images and context. If the primary purpose of the caption is to connect the image to its context, the paragraph should be crucial for generating it. For the description, on the other hand, the image takes on a primary role, and adding the context should result in lesser gains. Taken together, the quantitative model results are in line with the description/caption purpose distinction, suggesting a closer alignment of the description with the image and the caption with the context.

In sum, we find that providing linguistic context has the potential to improve model-generated descriptions as well as captions. These results also have implications for applications. Our results suggest that models trained to generate descriptions, which are generally sparse, might benefit from receiving the more available caption information as additional input, providing a new potential path for addressing image accessibility challenges.

5 Human Evaluation

In Section 3.3 and Section 4, we provided initial evidence that descriptions and captions fulfill separate communicative purposes. We now present a human-subject experiment in which we directly test whether captions are more suitable than descriptions for providing information that can't be obtained from the image alone, and whether descriptions are more useful than captions for replacing an image. Furthermore, we investigate whether models might capture these differences. Our experiment focuses on the Resnet-LSTM model, the weakest of our models, which should then provide a conservative assessment of how closely automatic generations align with human texts.

Our quantitative and qualitative hypotheses, along with our exclusion criteria and the analyses, were preregistered.⁵ Details on the recruitment and data exclusions are provided in Appendix C.

Materials We sampled 300 images from the validation set in Concadia. In the experiment, these could be paired with their original descriptions from Concadia, their original captions, or the model-generated descriptions or captions, resulting in four text conditions for each image.

⁵OSF preregistration link: <https://osf.io/qw9nr>.

Design In each trial, an image was displayed next to one of the descriptions/captions sampled from the four text conditions, followed by two questions (Figure 4). Participants assessed how useful the text would be to help someone imagine the picture (rated from *not useful* to *very useful*), and how much they learned from the text that they couldn't have learned from the image (rated from *nothing* to *a lot*). Each question was associated with a continuous scale initialized at zero and ranging to one. Participants could further opt out of responding to the questions by selecting a checkbox to indicate that image and text seemed to be unrelated.

Procedure Each participant saw 32 unique images, and the corresponding texts were uniformly sampled from the four text conditions. Image selection and trial order were randomized and question order was randomized between participants.

Participants We recruited 421 participants over Amazon's Mechanical Turk, who were paid \$2.60 for participation with an average completion time of 12 minutes (\$13/hr). After preregistered exclusions, the analysis is based on 277 participants.

5.1 Predictions

Overall, we predicted that descriptions and captions fulfill different communicative purposes. Concretely, we made the following predictions for the original human-generated descriptions and captions in Concadia:

- H1** The original descriptions are judged more useful for imagining the picture than the original captions.
- H2** A reader receives more extra-image information from the original captions than from the original descriptions.

We further predicted that a model should be capable of reflecting these differences when trained on descriptions and captions separately. Predictions (H1) and (H2) should therefore also hold for model-generated descriptions and captions.

Finally, we qualitatively predicted that models would be strong enough to show the following: generated descriptions should be more similar to original descriptions than the original captions would be, and vice versa for generated captions.

5.2 Results

To test our quantitative predictions, we performed a Bayesian mixed effects regression analysis, where

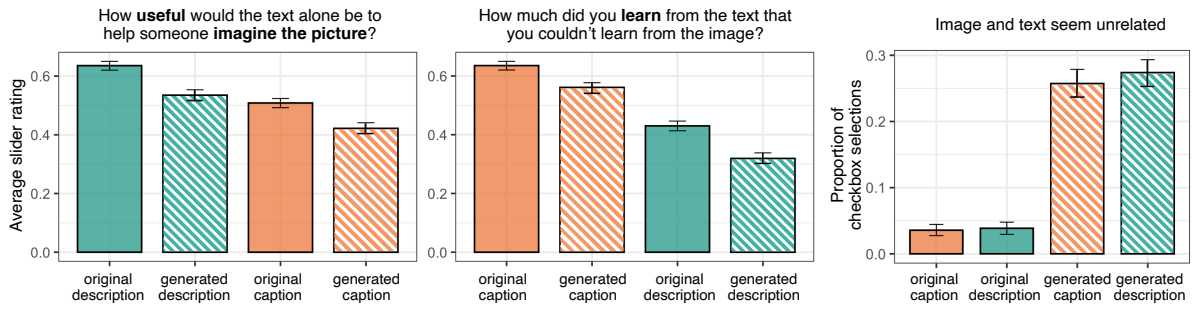


Figure 4: Results from the human subject experiments showing that descriptions and captions fulfill separate purposes and that these differences are picked up by models trained on them. Error bars indicate 95% bootstrapped confidence intervals.

we predicted continuous slider ratings from the centered categorical text condition variable with random intercepts and slopes for each participant and image. We find strong evidence for all of our quantitative predictions and we also see the predicted qualitative patterns, as shown in Figure 4. The results lead us to two main conclusions. Firstly, there are multiple dimensions in which an image-text pair can be successful. Our results show that we can reverse which texts are rated as “better” depending on what goal we prioritize: the ability of to reconstruct the image from the text in its absence, or the amount of information learned from the text in its presence. Secondly, models trained on caption and description data separately capture core aspects of these differences. Even with its limitations, the generated descriptions and captions provide more useful alternatives for their original counterparts than original captions and descriptions, respectively.

Unsurprisingly, the generated text was more often rejected as being unrelated to the image than the original descriptions and captions were, which were only rejected in less than 5% of all trials. These results underline the quality of Concadia and shortcomings of the ResNet-LSTM model.

Overall, our results highlight the importance of specifying the purpose a text is intended to fulfill in the image-text generation domain – both for dataset creation and model evaluation. The questions posed for evaluation have to be assessed against the purpose that they’re intended to fulfill.

6 Conclusion

We argue for a distinction between *descriptions* (text intended to replace an image) and *captions* (text intended to contextualize an image), and we provide evidence for their similarities and dif-

ferences through linguistic analyses and human-subject evaluations. We introduce a corpus, Concadia, consisting of images and their corresponding alt descriptions, captions, and context, and we show that Concadia is a valuable resource for training NLG systems for both description and caption generation. The structure of Concadia additionally allows us to show that both generation tasks benefit from contextual information. Beyond the conceptual relevance for advancing image-based NLG systems, this has practical implications specifically for image accessibility. Captions which often co-occur with images that lack an accessibility description could be a valuable resource for building more powerful description generation models.

Limitations and Ethics

The proposed dataset Concadia is crawled from Wikimedia Commons. Our understanding is that the images and associated texts are compliant with the policies of Wikimedia and Wikipedia.

Concadia is a subset of Wikipedia, filtered for images that had associated alt texts. Therefore, data might disproportionately be extracted from articles historically containing many images, such as the Tour de France or Award Shows, frequently visited articles where lacking image accessibility is more likely to be noticed, and from authors that are aware of the alt texts’ uses, for instance, from accessibility-related Wikipedia articles. Additionally, Concadia inherits biases generally present on Wikipedia (extraction date: December 2020). For instance as of February 2020, only 19% of biographical articles were about women (Tripodi, 2021), and it’s likely that gender as well as ethnicity biases are perpetuated in Concadia. Concadia is further constrained to English-language entries since this is the language we chose for our case

study. We provide a corresponding datasheet in our Github repository (Gebru et al., 2021).

The image-based natural language generation models were trained to allow for comparisons between description and caption generation, and to investigate the role of context in both cases. The models are not intended for deployment of any kind since they were not tested for crucial dimensions such as their robustness, truthfulness, or output of potentially discriminatory language. The model outputs further need to be tested to determine the extent to which they actually capture the intended use cases. For instance, a description might be intended to serve an image accessibility purpose, but rigorous testing is required to determine the extent to which the data and associated models fulfill this goal for blind and low vision users.

The human subject experiment was conducted under an IRB protocol. Participants were paid \$2.60 with an average completion time of 12 minutes (\$13/hr).

Acknowledgements

This work is supported in part by a grant from Google through the Stanford Institute for Human-Centered AI. We thank our experiment participants for their invaluable input. We are further grateful to Desmond Elliott, and Dan Jurafsky for their thorough feedback on previous drafts. We thank Rachit Dubey, all members of Stanford’s CoCoLab, and Jessica Mankewitz for their insightful comments on this work, and to Mike Wu for sharing useful materials.

References

- Malihe Alikhani, Piyush Sharma, Shengjie Li, Radu Soricut, and Matthew Stone. 2020. [Cross-modal coherence modeling for caption generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6525–6535, Online. Association for Computational Linguistics.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. SPICE: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, pages 382–398. Springer.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. 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*, pages 6077–6086.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). *ICLR*.
- Jeffrey P. Bigham, Ryan S. Kaminsky, Richard E. Ladner, Oscar M. Danielsson, and Gordon L. Hempton. 2006. [WebInSight: Making web images accessible](#). In *Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility - Assets '06*, page 181, Portland, Oregon, USA. ACM Press.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Ali Furkan Biten, Lluís Gomez, Marçal Rusinol, and Dimosthenis Karatzas. 2019. [Good News, Everyone! Context driven entity-aware captioning for news images](#). In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12458–12467, Long Beach, CA, USA. IEEE.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. 2021. Conceptual 12M: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3557–3567. IEEE.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. 2015. [Microsoft COCO captions: Data collection and evaluation server](#). arXiv:1504.00325.
- Sanjana Chintalapati, Jonathan Bragg, and Lucy Lu Wang. 2022. [A dataset of alt texts from HCI publications: Analyses and uses towards producing more descriptive alt texts of data visualizations in scientific papers](#). arXiv:2209.13718.
- Zhenrong Deng, Zhouqin Jiang, Rushi Lan, Wenming Huang, and Xiaonan Luo. 2020. Image captioning using densenet network and adaptive attention. *Signal Processing: Image Communication*, 85:115836.
- Michael Denkowski and Alon Lavie. 2014. Meteor Universal: Language Specific Translation Evaluation for Any Target Language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 376–380.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Pierre Dognin, Igor Melnyk, Youssef Mroueh, Inkit Padhi, Mattia Rigotti, Jarret Ross, Yair Schiff, Richard A Young, and Brian Belgodere. 2022. Image captioning as an assistive technology: Lessons learned from VizWiz 2020 Challenge. *Journal of Artificial Intelligence Research*, 73:437–459.
- Pierre Dognin, Igor Melnyk, Youssef Mroueh, Jarret Ross, and Tom Sercu. 2019. Adversarial semantic alignment for improved image captions. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10455–10463. IEEE.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92.
- Cole Gleason, Patrick Carrington, Cameron Cassidy, Meredith Ringel Morris, Kris M. Kitani, and Jeffrey P. Bigham. 2019. “It’s almost like they’re trying to hide it”: How user-provided image descriptions have failed to make Twitter accessible. In *The World Wide Web Conference on - WWW ’19*, pages 549–559, San Francisco, CA, USA. ACM Press.
- Darren Guinness, Edward Cutrell, and Meredith Ringel Morris. 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 - CHI ’18*, pages 1–11, Montreal QC, Canada. ACM Press.
- Danna Gurari, Yanan Zhao, Meng Zhang, and Nilavra Bhattacharya. 2020. Captioning images taken by people who are blind. In *European Conference on Computer Vision*, pages 417–434. Springer.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. CLIPScore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899.
- Md Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, Hamid Laga, and Mohammed Bennamoun. 2021. Text to image synthesis for improved image captioning. *IEEE Access*, 9:64918–64928.
- Alejandro Jaimes and Shih-Fu Chang. 1999. Conceptual framework for indexing visual information at multiple levels. In *Electronic Imaging*, pages 2–15, San Jose, CA.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *ICLR*.
- Julia Kruk, Jonah Lubin, Karan Sikka, Xiao Lin, Dan Jurafsky, and Ajay Divakaran. 2019. Integrating text and image: Determining multimodal document intent in Instagram posts. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4622–4632, Hong Kong, China. Association for Computational Linguistics.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. 2020. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pages 121–137. Springer.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision*, volume 8693, pages 740–755, Cham. Springer International Publishing.
- Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. 2017. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 375–383.
- Emily E. Marsh and Marilyn Domas White. 2003. A taxonomy of relationships between images and text. *Journal of Documentation*, 59(6):647–672.
- Radan Martinec and Andrew Salway. 2005. A system for image–text relations in new (and old) media. *Visual communication*, 4(3):337–371.
- Meredith Ringel Morris, Annuska Zolyomi, Catherine Yao, Sina Bahram, Jeffrey P. Bigham, and Shaun K. Kane. 2016. “With most of it being pictures now, I rarely use it”: Understanding Twitter’s evolving accessibility to blind users. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5506–5516.
- Annika Muehlbradt and Shaun K. Kane. 2022. What’s in an ALT tag? Exploring caption content priorities through collaborative captioning. *ACM Transactions on Accessible Computing*, 15(1):6:1–6:32.
- Nelleke Oostdijk, Hans van Halteren, Erkan Başar, and Martha Larson. 2020. The connection between the text and images of news articles: New insights for

- multimedia analysis. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4343–4351.
- Christian Otto, Matthias Springstein, Avishek Anand, and Ralph Ewerth. 2020. [Characterization and classification of semantic image-text relations](#). *International Journal of Multimedia Information Retrieval*, 9(1):31–45.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, and Luca Antiga. 2019. PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32:8026–8037.
- Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. 2020. [Connecting vision and language with localized narratives](#). In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, volume 12350, pages 647–664. Springer International Publishing, Cham.
- Arnau Ramisa, Fei Yan, Francesc Moreno-Noguer, and Krystian Mikolajczyk. 2017. BreakingNews: Article annotation by image and text processing. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1072–1085.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [Conceptual Captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, Melbourne, Australia. Association for Computational Linguistics.
- Sara Shatford. 1986. [Analyzing the subject of a picture: A theoretical approach](#). *Cataloging & Classification Quarterly*, 6(3):39–62.
- Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. 2020. [TextCaps: A dataset for image captioning with reading comprehension](#). *arXiv:2003.12462 [cs]*.
- Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021. WIT: Wikipedia-based image text dataset for multimodal multilingual machine learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2443–2449.
- Abigale Stangl, Meredith Ringel Morris, and Danna Gurari. 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*, pages 1–13, Honolulu HI USA. ACM.
- Abigale Stangl, Nitin Verma, Kenneth R Fleischmann, Meredith Ringel Morris, and Danna Gurari. 2021. Going beyond one-size-fits-all image descriptions to satisfy the information wants of people who are blind or have low vision. In *The 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–15.
- Francesca Tripodi. 2021. Ms. Categorized: Gender, notability, and inequality on Wikipedia. *New Media & Society*.
- Emiel van Miltenburg, Ákos Kádár, Ruud Koolen, and Emiel Kraemer. 2018. [DIDEC: The Dutch image description and eye-tracking corpus](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3658–3669, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Emiel van Miltenburg, Ruud Koolen, and Emiel Kraemer. 2018. Varying image description tasks: Spoken versus written descriptions. In *Proceedings of the Fifth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2018)*, pages 88–100, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. CIDEr: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4566–4575.
- Candace Williams, Lilian de Greef, Ed Harris III, Leah Findlater, Amy Pavel, and Cynthia Bennett. 2022. Toward supporting quality alt text in computing publications. In *Proceedings of the 19th International Web for All Conference*, pages 1–12.

Kelvin Xu, Jimmy Lei Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*, pages 2048–2057.

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.

Mingda Zhang, Rebecca Hwa, and Adriana Kovashka. 2018. Equal but not the same: Understanding the implicit relationship between persuasive images and text. In *British Machine Vision Conference (BMVC)*.

Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. VinVL: Revisiting visual representations in vision-language models. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5575–5584. IEEE.

Appendix

A Additional Details on Concadia

A.1 Datasheet

A datasheet (Geburu et al., 2021) for Concadia is included in our Github repository.

A.2 Occurrence Frequency Estimates

To estimate the sparsity of alt descriptions and captions for images on Wikipedia, we randomly sampled 10,000 articles, extracted the number of images, and counted how many of those contained captions and how many contained descriptions. We

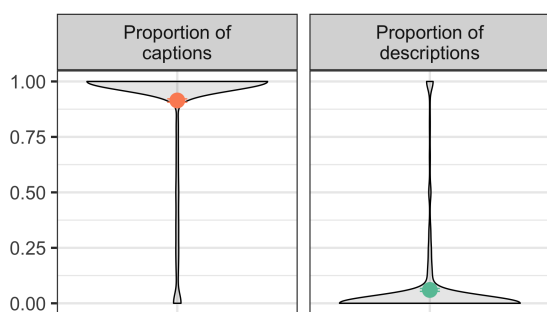


Figure 5: Proportion of images with captions vs. descriptions on Wikipedia, as of December 2020. The grey areas visualize the data distribution and the points mark average proportion of captions/descriptions for an image with 95% bootstrapped confidence intervals. Each datapoint is the average proportion for each of the 10,000 randomly sampled articles.

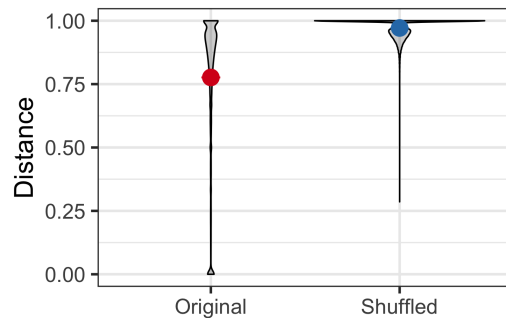


Figure 6: Jaccard distance between the originally matched captions and descriptions and shuffled caption-description pairs.

then took the average proportion of captions and descriptions per image for each article, which yields the data distribution displayed in grey in Figure 5.

A.3 Similarity of Captions and Descriptions

In Section 3.3, we demonstrate that captions and descriptions for the same image in Concadia are semantically similar to each other using SBert embeddings. These results replicate when calculating similarity using Jaccard distance of the raw strings, a more form-based and superficial perspective on the caption/description comparison. Figure 6 summarizes this analysis. The random pairings are significantly more distant from each other than their ordered counterparts (two-sample Kolmogorov-Smirnov test: $p < 0.0001$). The results still support the observation of distinctness between captions and descriptions since the Jaccard distance in the ordered set is still high (0.78).

In Section 3.3, we further present differences in occurrence frequencies of the syntactic categories

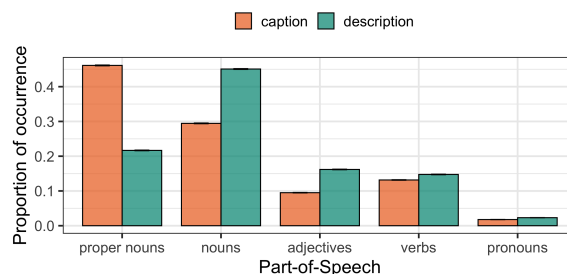


Figure 7: Aggregated proportions of Part-of-Speech frequencies for descriptions and captions with 95% confidence intervals. The following POS tags constitute the categories on the x axis: proper nouns (NNP, NNPS), nouns (NN, NNS), adjectives (JJ, JJR, JJS), verbs (VB, VBD, VBG, VBN, VBP, VBZ), pronouns (PRP, PRPS).

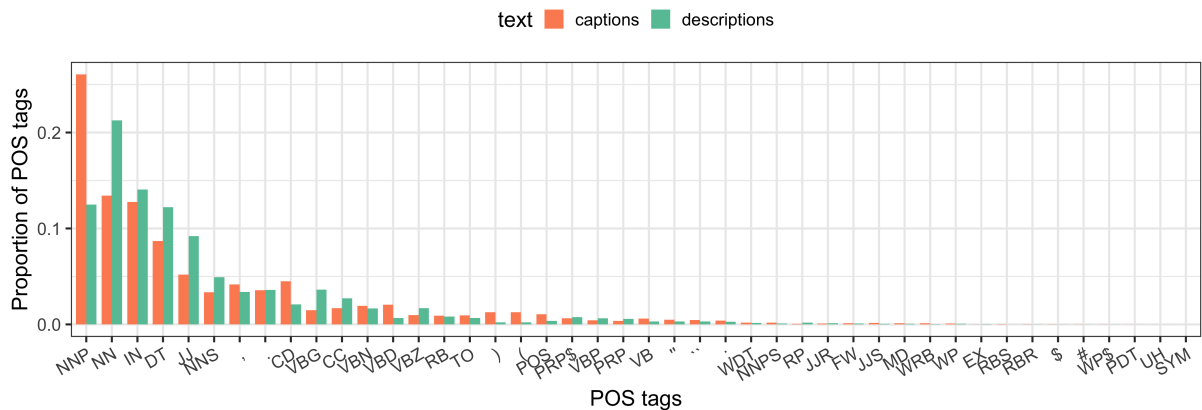


Figure 8: Proportion of part-of-speech tokens in descriptions vs. captions from Concadia.

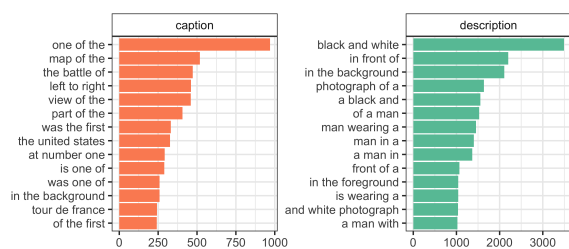


Figure 9: Most frequent trigrams in Concadia for captions and descriptions.

proper nouns, nouns, and adjectives for descriptions and captions, as illustrated in Figure 7. However, the Part-of-Speech (POS) tag analysis using the Python library NLTK (Bird and Loper, 2004) allows for much more finegrained distinctions. Figure 8 contains all analyzed POS tags and shows the proportion at which they occur in captions and descriptions of the Concadia corpus, revealing informative sub-patterns. For example, while there is a clear over-representation of adjectives in descriptions, this pattern qualitatively flips for superlatives. We speculate that superlatives are potentially more evaluative and therefore considered less appropriate in descriptions. There is also a clear over-representation of determiners in descriptions, which we attribute to the higher frequency of common nouns over proper names, since in English *Jesse* doesn't need a determiner to make an utterance grammatical whereas *person* most often does. Proportions are computed by dividing the number of occurrences of a POS tag by the number of all tokens in that specific text form. This allows for a direct comparison between descriptions and captions, even though captions contain more tokens overall. Table 3 contains an overview of all POS

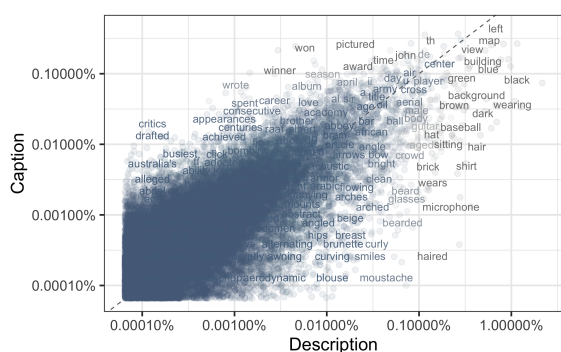


Figure 10: Correlation of word frequency between captions and descriptions.

tags and their respective definition, as provided by NLTK (Bird and Loper, 2004).

Section 3.3 also addresses qualitative semantic differences between captions and descriptions. Due to space constraints, we only use the most frequent bigrams (without stopwords) to demonstrate these differences in the paper. Figure 9 illustrates the most frequent trigram frequencies when including stopwords. Similarly to the bigrams in Figure 3, the most frequent trigrams in captions have instances of event descriptions and proper names (e.g., *the battle of, Tour de France*). For descriptions, there are again primarily descriptive attributes for people (e.g., *man wearing a, is wearing a*), and meta information about the image (e.g., *(black) and white photograph, in the foreground*).

A similar pattern arises when directly comparing which words occur more frequently in captions and descriptions, as shown in Figure 10. Words that lie on the dashed line occur equally often in both text forms. Words below that line occur more frequently in descriptions and words above it occur more often in captions. Words that occur more

frequently in descriptions focus on descriptive adjectives and nouns such as *hat*, *bearded*, or *blouse*. Words that occur more frequently in captions rather refer to an event, e.g., *award* or *winner*, and use more evaluative language, e.g., *achieved*.

Finally, a simple binary classifier based on the SBert embeddings of the descriptions and captions achieves an accuracy of 79.3% (train set: 79.5%, dev set: 79.5%) providing further quantitative evidence suggesting that descriptions and captions carry distinguishable features.

These additional perspectives are in further support of the quantitative and qualitative analyses of the principled ways in which the content of captions and descriptions differs (see Section 3.3).

A.4 Illustrations

In Figure 11, we provide examples from Conca-dia, illustrating a range of topics and image types: photographs, drawings and paintings.

B Models

B.1 ResNet-LSTM

B.1.1 Architecture

Encoders Two pretrained encoders create image and context representations which are input to an LSTM with attention. As an image encoder, we use ResNet-101 (He et al., 2016), pretrained on ImageNet.⁶ For the context representation, we use pretrained BERT (Devlin et al., 2019).⁷ When predicting labels in absence of context, we induce an uninformative context representation by setting all values to 1. This way, we ensure that potential model performance improvements are not due to an increase of model complexity when adding context. This is also one of the reasons why neither image nor context encoders were finetuned. To establish the importance of input components, e.g., context, we needed to fix the number of trainable parameters to avoid conflation from changes in model complexity.

Decoder with Attention The initial input to the LSTM hidden states of the decoder is a vector concatenating an image and a context representation. The LSTM decoder generates each token in a predicted label based on a previous token, the previous

⁶We used the pretrained weights available with Pytorch’s torchvision module version 0.8.2.

⁷We used the pretrained weights available with the Python module pytorch-pretrained-bert version 0.6.2.

POS tags	Explanation
CC	Coordinating Conjunction
CD	Cardinal Digit
DT/PDT	Determiner/Predeterminer
EX	Existential There
FW	Foreign Word
IN	Preposition/Subordin. Conjunct.
JJ/JJR/JJS	Adj./Comparative/Superlative
LS	List Marker 1
MD	Modal
NN/NNS	Noun Singular/Plural
NNP/NNPS	Proper Noun Singular/Plural
POS	Possessive Ending
PRP/PRP\$	Personal/Possessive Pronoun
RB/RBR/RBS	Adverb/Comparative/Superlative
RP	Particle
TO	to
UH	Interjection
VB	Verb, Base Form
VBD/VBN	V, Past Tense/Past Participle
VBG	V, Gerund or Present Participle
VBP/VPZ	V, Sg. Present, non-3rd/3rd P.
WDT/WRB	wh-determiner/wh-abverb
WP/WP\$	wh-pronoun/possessive

Table 3: Part-of-speech tokens from Figure 8.

hidden state, and an attention vector. The attention vector is obtained by using soft attention (Xu et al., 2015; Bahdanau et al., 2015) separately parameterized on the image and the context and then concatenating the resulting representations.

B.1.2 Implementation

The ResNet-LSTM models were implemented in PyTorch (Paszke et al., 2019) using a codebase that has successfully replicated results from Xu et al. (2015).⁸ Additional details on model training and optimization will be available in our code release.


B.1.3 Hyperparameters and Training Details

We used the hyperparameters suggested in this codebase. The only exception is that we used a reduced batch size (32 instead of 80) since our models that include context have more parameters.

Overview of hyperparameters:

- dimension of image encoder output: 2048 (predetermined by the pretrained ResNet-101 He et al. (2016) output size)

⁸<https://github.com/sgrvinod/a-PyTorch-Tutorial-to-Image-Captioning>



Caption: Taapaca rises above Putre
Description: A town within a valley, with two snow-covered summits rising above
Paragraph: Taapaca reaches a height of 5,860 metres (19,230ft) above sea level and is a volcanic complex elongated from west to east; Larancagua lies just east of Taapaca. It consists primarily of many overlapping lava domes, with shapes ranging from almost elliptical to circular. Lava flows are uncommon. The volcano has generated an apron of black-and-ash flow deposits [...]



Caption: Scene from the play, published 1856
Description: black and white drawing of stage scene: an older man, holding a pot-plant, confronts a younger one
Paragraph: The Italian Straw Hat (Un chapeau de paille d'Italie) is a five-act comedy by Eugne Labiche and Marc-Michel. It premiered at the Thtre du Palais-Royal in Paris on 14 August 1851. It has been adapted for the cinema in French, English, German, Czech and Russian, and as a musical play in English and Italian versions. [...]



Caption: Hathor, in bovine form, emerges from a hill representing the Theban necropolis, in a copy of the Book of the Dead from the 13th century BC
Description: Painting of a cow whose head protrudes from a hill, in front of which stand papyrus stalks and a pyramidal chapel
Paragraph: Hathor was one of several goddesses believed to assist deceased souls in the afterlife. One of these was Imentet, the goddess of the west, who personified the necropolises, or clusters of tombs, [...]



Caption: A US Marine during the fighting for the Vegas outpost, March 1953.
Description: A Caucasian soldier wearing a camouflage helmet and thick winter clothing lies on his side in a trench facing the right of the photo, holding a canteen in his left hand and using an entrenching tool with his right hand. The soldier is surrounded by rolls of communication line, jerries and boxes of equipment. Behind him, the skyline can be seen above the top of the earth works.
Paragraph: In the western sector, American outposts such as Detroit, Berlin, East Berlin, Reno, Carson, Elko and Vegas had been bitterly contested over the previous months by the US 1st Marine Division [...]

Figure 11: Sample images with the corresponding descriptions, captions and excerpts of the context paragraphs from Concadia.

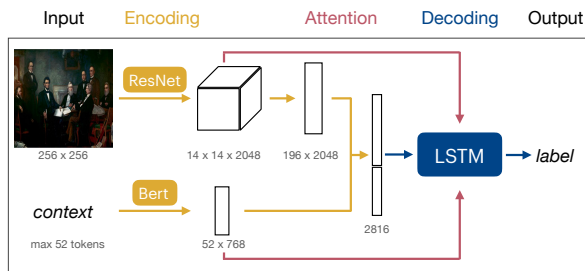


Figure 12: ResNet-LSTM model architecture. The LSTM decoder takes image as well as text input with attention on both representations. If no context is provided, the context representation contains ones only. In our experiments, the model receives the caption as context to produce descriptions, and vice versa.

- dimension of context encoder output: 768 (predetermined by the pretrained Bert Dognin et al. (2019) output size)
- dimension of word embeddings: 512
- dimension of attention linear layers: 512
- dimension of decoder RNN: 512
- dropout: 0.5
- loss function: cross-entropy loss
- optimizer: Adam (Kingma and Ba, 2015)
- batch size: 32
- learning rate for decoder: 4e-4
- clip gradients at an absolute value of 5
- regularization parameter for “doubly stochastic attention”: 1

The model contains 396M parameters, of which 43M are in the image encoder, 110M are in the context encoder, and the rest are in the decoder. The model has 244M trainable parameters, all of which are in the decoder, as the encoders are frozen.

B.1.4 Compute

All models were trained and evaluated on AWS EC2. Description-generation models were trained either on an g4dn.xlarge instance or an g5.xlarge instance. All caption-generation models were trained on an g5.xlarge instance. When training on an g4dn.xlarge instance, each model takes about 1.5 days. When training on an g5.xlarge instance, each model takes about 1 day. This includes evaluation on the validation set after each epoch, where the model generates a sequence with a beam size of 1 for each example.

For each model, evaluation on the test set takes about 2 hours on a p2.xlarge instance. Evaluation on the test set uses a beam size of 5.

B.2 DenseNet-LSTM

The DenseNet-LSTM follows the same structure as the ResNet-LSTM except that the ResNet-101 image encoder is replaced by a DenseNet-161. Consequently, the models differ in the number of parameters. The DenseNet-LSTM contains 382M parameters, of which 26M are in the image encoder, 110M are in the context encoder, and the rest are in the decoder. The model has 245M trainable parameters, all of which are in the decoder, as the encoders are frozen. For all other aspects,

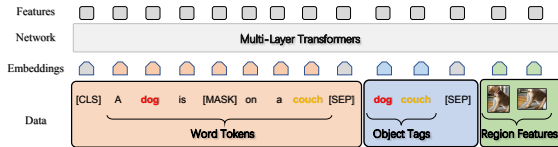


Figure 13: Illustration of the OSCAR architecture (crop taken from Li et al. (2020)). For each image–text pair, the image is first preprocessed by a pretrained visual feature extractor, which outputs object tags in text and region features, i.e., vector embeddings of object-enclosing regions. The OSCAR model represents the image–text pair by concatenating the text and the object tags, embedding it through BERT (Devlin et al., 2019), then concatenating the BERT embedding with the visual features (Li et al., 2020). We use VinVL as the visual feature extractor (Zhang et al., 2021).

consult the ResNet-LSTM specifications.

B.3 OSCAR(VinVL)

B.3.1 Architecture

OSCAR is a Transformer-based visual-language model that is pretrained on a large-scale corpus and can be finetuned for a variety of downstream vision-language tasks (Li et al., 2020), including image-to-text generation. See Figure 13 for a diagram of the architecture, and refer to Section 4 of Li et al. 2020 for the image-to-text finetuning procedure.

There are potentially many ways in which one could integrate context into the model. We took the approach of appending the context to the object tags that OSCAR receives.

B.3.2 Implementation

We applied the VinVL feature extractor to images in Concadia using the VinVL codebase available at https://github.com/microsoft/scene_graph_benchmark. We adopted the OSCAR model implementation and training script from the OSCAR codebase available at <https://github.com/microsoft/Oscar>.

B.3.3 Hyperparameters and Training Details

We used the pretrained BERT_{Base}-based OSCAR(VinVL) checkpoint as our starting point and finetuned it on Concadia. We used the same hyperparameters as the original OSCAR implementation, with two exceptions: due to compute constraints, we fine-tuned OSCAR for 20 epochs (compared to 40) and used a batch size of 32 (instead of 256).

Aside from optimizing for cross-entropy loss, the original implementation BERT further finetuned the model with CIDEr optimization as the objective.

Label	Context	RN-LSTM	DN-LSTM	OSC.VinVL
Descr.	None	0.15 (0.01)	0.19 (0.00)	0.24 (0.00)
Caption	None	0.07 (0.00)	0.10 (0.01)	0.12 (0.00)
Descr.	Rand. Par.	0.14	0.17	0.23 (0.00)
Caption	Rand. Par.	0.06	0.08	0.12 (0.00)
Descr.	Paragr.	0.17 (0.01)	0.20 (0.00)	0.31 (0.00)
Caption	Paragr.	0.11 (0.00)	0.13 (0.00)	0.44 (0.00)
Descr.	Caption	0.22 (0.01)	0.24 (0.00)	0.98 (0.01)
Caption	Descr.	0.15 (0.00)	0.17 (0.00)	1.00 (0.00)

Table 4: CIDEr scores on validation set for all models: ResNet-LSTM, DenseNet-LSTM and OSCAR(VinVL). Standard deviations (in brackets) are computed over 3 random seeds.

Since it wasn’t central to our investigation, we omitted the CIDEr optimization step. By the end of training, the differences between conditions had already persisted over several epochs and validation set performance had converged.

The model contains ≈ 110 M parameters, 88M of which are trainable.

B.3.4 Compute

All models were trained and evaluated on AWS EC2 g4dn.xlarge instances. Each model takes about 2 days to train. This includes evaluation on the validation set after each epoch, where the model generates a sequence with a beam size of 1 for each example.

For each model, evaluation on the test set takes about 8 hours on a g4dn.xlarge instance. Evaluation on the test set uses a beam size of 5.

B.4 Validation Set Performance

The test set performances in Table 2 are similarly reflected on Concadia’s validation set. Table 4 shows the validation set results. Note that evaluation on the validation set uses a beam size of 1 (different from 5 on the test set).

C Human Subject Experiment

C.1 Participants

We recruited 421 participants over Amazon’s Mechanical Turk, restricted to participation within the US and a previous average approval rate above 98%. Participants were paid \$2.60 for participation with an average completion time of 12 minutes (\$13/hr). Participants were only allowed to complete the experiment once.

New Orleans lithograph from 1852

Q1: How much did you learn from the text that you couldn't learn from the image?

Nothing A lot

Q2: How useful would the text alone be to help someone imagine this picture (e.g. a visually impaired person)?

Not useful Very useful

Can't say because image and text seem to be unrelated.

Figure 14: A critical trial as it appeared to the participants. Question order was randomized between participants but consistent throughout the experiment.

C.2 Details on Materials & Procedure

Throughout the experiment, participants never saw the same picture twice. Each participant saw at least seven samples of each text condition, resulting in 28 trials. The last two trial conditions were randomly sampled. An example trial as it appeared to participants is provided in Figure 14. Image selection and trial order was completely randomized and question order was randomized between participants. The experiments themselves are released as part of the Github repository.

C.3 Data Exclusions

Participants were excluded based on their self-reported native language, their self-assessment of task performance, two attention check questions, and their overall response. Firstly, we only included participants who indicated that English is among their native languages (12 exclusions), and who reported that they thought they did the experiment correctly (74 exclusions). We constructed attention check questions where the text clearly described something different than what was displayed in the image and excluded participants who didn't select the checkbox "Can't say because image and text seem to be unrelated" (85 exclusions). Furthermore, we excluded participants who selected the checkbox in more than 80% of all trials (6 exclusions). Overall, we excluded 177 partici-

pants (42%). Furthermore, we excluded trials from the quantitative analysis where participants indicated that text and image didn't relate. If more than 80% of participants agreed that image and text didn't relate, we excluded that data sample from our quantitative analysis as well (13,942 out of 15,616 datapoints (89.3%) remain).