

# Visual Madlibs: Fill in the blank Image Generation and Question Answering

Licheng Yu, Eunbyung Park, Alexander C. Berg, and Tamara L. Berg

Department of Computer Science, University of North Carolina, Chapel Hill

{licheng, eunbyung, acberg, tlberg}@cs.unc.edu

## Abstract

In this paper, we introduce a new dataset consisting of 360,001 focused natural language descriptions for 10,738 images. This dataset, the Visual Madlibs dataset, is collected using automatically produced fill-in-the-blank templates designed to gather targeted descriptions about: people and objects, their appearances, activities, and interactions, as well as inferences about the general scene or its broader context. We provide several analyses of the Visual Madlibs dataset and demonstrate its applicability to two new description generation tasks: focused description generation, and multiple-choice question-answering for images. Experiments using joint-embedding and deep learning methods show promising results on these tasks.

## 1. Introduction

Much of everyday language and discourse concerns the visual world around us, making understanding the relationship between the physical world and language describing that world an important challenge problem for AI. Understanding this complex and subtle relationship will have broad applicability toward inferring human-like understanding for images, producing natural human robot interactions, and for tasks like natural language grounding in NLP. In computer vision, along with improvements in deep learning based visual recognition, there has been an explosion of recent interest in methods to automatically generate natural language descriptions for images [5, 9, 15, 32, 16, 20] or videos [31, 8]. However, most of these methods and existing datasets have focused on only one type of description, a generic description for the entire image.

In this paper, we collect a new dataset of focused, targeted, descriptions, the *Visual Madlibs dataset*, as illustrated in Figure 1. To collect this dataset, we introduce automatically produced fill-in-the-blank templates designed to collect a range of different descriptions for visual content in an image. For example, a user might be presented with an



1. This place is a park.
2. When I look at this picture, I feel competitive.
3. The most interesting aspect of this picture is the guys playing shirtless.
4. One or two seconds before this picture was taken, the person caught the frisbee.
5. One or two seconds after this picture was taken, the guy will throw the frisbee.
6. Person A is wearing blue shorts.
7. Person A is in front of person B.
8. Person A is blocking person B.
9. Person B is a young man wearing an orange hat.
10. Person B is on a grassy field.
11. Person B is holding a frisbee.
12. The frisbee is white and round.
13. The frisbee is in the hand of the man with the orange cap.
14. People could throw the frisbee.
15. The people are playing with the frisbee.

Figure 1: An example from the Visual Madlibs Dataset. This dataset collects targeted descriptions for people and objects, denoting their appearances, affordances, activities, and interactions. It also provides descriptions of broader emotional, spatial and temporal context for an image.

image and a fill-in-the-blank template such as “The frisbee is [blank]” and asked to fill in the [blank] with a description of the appearance of frisbee. Alternatively, they could be asked to fill in the [blank] with a description of what the person is doing with the frisbee. Fill-in-the-blank questions can be targeted to collect descriptions about people and objects, their appearances, activities, and interactions, as well as descriptions of the general scene or the broader emotional, spatial, or temporal context of an image. Using these templates, we collect a large collection of 360,001



This place is a(n) road.  
 When I look at this picture, I feel free.  
 The most interesting aspect of this picture is the motorcycles.  
 One or two seconds before this picture was taken, they stopped to chat and decided where to go.  
 One or two seconds after this picture was taken, the bikers ride down the road.



This place is a(n) restaurant.  
 When I look at this picture, I feel like I want donuts.  
 The most interesting aspect of this picture is the box of doughnuts.  
 One or two seconds before this picture was taken, the box was closed.  
 One or two seconds after this picture was taken, a third person out of frame picks up a doughnut.



This place is a(n) waterway.  
 When I look at this picture, I feel concerned.  
 The most interesting aspect of this picture is the men standing on elephants.  
 One or two seconds before this picture was taken, the people were sitting on the elephant.  
 One or two seconds after this picture was taken, the men got off the elephants.



Person A is a girl.  
 Person A is learning how to surf.  
 Person A is kneeling on a surfboard.  
 Person B is a man in blue shorts.  
 Person B is walking on a beach.  
 Person B is at the beach.  
 Person C is a short haired black girl.  
 Person C is practicing surfing.  
 Person C is on a gold surfboard.  
 Person D is a lady in board shorts.  
 Person D is standing around.  
 Person D is next to a blue surfboard.



Person A is wearing a dark suit.  
 Person A is looking at an elephant.  
 Person A is at a zoo.  
 Person B is wearing a grey T-shirt.  
 Person B is talking to two other people.  
 Person B is next to an elephant.  
 Person C is a blonde woman.  
 Person C is getting an elephant.  
 Person C is at a zoo.



Person A is a balding male.  
 Person A is playing a video game.  
 Person A is downstairs.  
 Person B is wearing jeans.  
 Person B is playing wii.  
 Person B is in a basement.  
 Person C is wearing dark clothes.  
 Person C is walking upstairs.  
 Person C is in a building.  
 Person D is mostly out of the photo.  
 Person D is walking away.  
 Person D is in the basement.



Person A is a young man in green.  
 Person A is trying to block a frisbee.  
 Person A is on a field.  
 Person B is wearing purple.  
 Person B is throwing a frisbee.  
 Person B is on a field.



The couches are white.  
 The couches are in the center of the room.  
 People could relax on the couches.  
 The TV is on.  
 The TV is near the wall.  
 People could watch the TV.



The car is white.  
 The car is on a concrete pad.  
 People could ride in the car.  
 The umbrellas are open.  
 The umbrellas are in the people's hands.  
 People could stay dry under the umbrellas.



The person is putting food in the bowl.



The people are eating cake at the dining table.  
 The people are serving the cake.

Figure 2: Madlibs description. The first row corresponds to question types 1-5, the second row corresponds to question types 9-11, and the third row is to question types 6-8 and question type 12. All question types are listed in Table 1.

targeted descriptions for 10,738 images. Fig. 2 shows some Madlibs description samples.

With this new dataset, we can develop methods to generate more focused descriptions. Instead of asking an algorithm to “describe the image” we can now ask for more focused descriptions such as “describe the person”, “describe what the person is doing,” or “describe the relationship between the person and the frisbee.” We can also ask questions about aspects of an image that are somewhat beyond the scope of the directly depicted content. For example, “describe what might have happened just before this picture was taken.” or “describe how this image makes you feel.” These types of descriptions reach toward high-level goals of producing human-like visual interpretations for images.

In addition to focused description generation, we also introduce a multiple-choice question-answering task for images. In this task, the computer is provided with an image and a partial description such as “The person is [blank]”. A set of possible answers is also provided, one answer that was written about the image in question, and several ad-

ditional answers written about other images. The computer is evaluated on how well it can select the correct choice. In this way, we can evaluate performance of description generation on a concrete task, making evaluation more straightforward. Varying the difficulty of the negative answers—adjusting how similar they are to the correct answer—provides a nuanced measurement of performance.

For both the generation and question-answering tasks, we study and evaluate a recent state of the art approach for image description generation [32], as well as a simple joint-embedding method learned on deep representations. The evaluation also includes extensive analysis of the Visual Madlibs dataset and comparisons to the existing MS COCO dataset of natural language descriptions for images.

In summary, our contributions are:

- 1) A new description collection strategy, *Visual Madlibs*, for constructing fill-in-the-blank templates to collect targeted natural language descriptions.
- 2) A new Visual Madlibs Dataset consisting of 360,001 targeted descriptions, spanning 12 different types of templates,

for 10,738 images, as well as analysis of the dataset and comparisons to existing MS COCO descriptions.

3) Evaluation of a generation method and a simple joint embedding method for targeted description generation.

4) Definition and evaluation of generation and joint-embedding methods on a new task, *multiple-choice fill-in-the-blank question answering for images*.

The rest of our paper is organized as follows. First, we review related work (Sec 2). Then, we describe our strategy for automatically generating fill-in-the-blank templates and introduce our Visual Madlibs dataset (Sec 3). Next we outline the multiple-choice question answering and targeted generation tasks (Sec 4) and provide several analyses of our dataset (Sec 5). Finally, we provide experiments evaluating description generation and joint-embedding methods on the proposed tasks (Sec 6) and conclude (Sec 7).

## 2. Related work

**Description Generation:** Recently, there has been an explosion of interest in methods for producing natural language descriptions for images or video. Early work in this area generally explored two complementary directions. The first type of approach focused on detecting content elements such as objects, attributes, activities, or spatial relationships and then composing captions for images [18, 33, 26, 10] or videos [17] using linguistically inspired templates. The second type of approach explored methods to make use of existing text either directly associated with an image [11, 1] or retrieved from visually similar images [27, 19, 24].

With the advancement of deep learning for content estimation, there have been many exciting recent attempts to generate image descriptions using neural network based approaches. Some methods first detect words or phrases using Convolutional Neural Network (CNN) features, then generate and re-rank candidate sentences [9, 20]. Other approaches take a more end-to-end approach to generate output descriptions directly from images. Kiros *et al.* [16] learn a joint image-sentence embedding using visual CNNs and Long Short Term Memory (LSTM) networks. Similarly, several other methods have made use of CNN features and LSTM or recurrent neural networks (RNN) for generation with a variety of different architectures [32, 15, 5]. These new methods have shown great promise for image description generation under some measures (e.g. BLEU-1) achieving near-human performance levels. We look at related, but more focused description generation tasks.

**Description Datasets:** Along with the development of image captioning algorithms there have been a number of datasets collected for this task. One of the first datasets collected for this problem was the UIUC Pascal Sentence data set [10] which contains 1,000 images with 5 sentences per image written by workers on Amazon Mechanical Turk. As the description problem gained popularity larger and richer

datasets were collected, including the Flickr8K [28] and Flickr30K [34] datasets, containing 8,000 and 30,000 images respectively. In an alternative approach, the SBU Captioned photo dataset [27] contains 1 million images with existing captions collected from Flickr. This dataset is larger, but the text tends to contain more contextual information since captions were written by the photo owners. Most recently, Microsoft released the MS COCO [21] dataset. MS COCO contains 120,000 images depicting 80 common object classes, with object segmentations and 5 turker written descriptions per image. These datasets have been one of the driving forces in improving methods for description generation, but are currently limited to a single description about the general content of an image. We make use of MS COCO data, extending the types of descriptions associated with images.

**Question-answering** Natural language question-answering has been a long standing goal of NLP, with commercial companies like Ask-Jeeves or Google playing a significant role in developing effective methods. Recently, embedding and deep learning methods have shown great promise for question-answering [30, 3, 4]. Lin *et al.* [22] take an interesting multi-modal approach to question-answering. A multiple-choice text-based question is first constructed from 3 sentences written about an image; 2 of the sentences are used as the question, and 1 is used as the positive answer, mixed with several negative answers from sentences written about other images. The authors develop ranking methods to answer these questions and show that generating abstract images for each potential answer can improve results. Note, here the algorithms are not provided with an image as part of the question. Some recent work has started to look at the problem of question-answering for images. Malinowski *et al.* [23] combine computer vision and NLP in a Bayesian framework, but restrict their method to scene based questions. Geman *et al.* [12] design a visual Turing test to test image understanding using a series of binary questions about image content. We design more general question-answering tasks that allow us to ask a variety of different types of natural language questions about images.

## 3. Designing and collecting Visual Madlibs

The goal of Visual Madlibs is to study targeted natural language descriptions of image content that go beyond describing which objects are in the image, and beyond generic descriptions of the whole image. The experiments in this paper begin with a dataset of images where the presence of some objects have already been labeled<sup>1</sup>. The prompts for the Madlibs-style fill-in-the-blank questions are automatically generated based on image content, in a

<sup>1</sup>More generally, acquiring such labels could be included as part of collecting Madlibs.

Type	Instruction	Prompt	#words
1. image’s scene	Describe the type of scene/place shown in this picture.	The place is a(n) ____ .	4+1.45
2. image’s emotion	Describe the emotional content of this picture.	When I look at this picture, I feel ____ .	8+1.14
3. image’s interesting	Describe the most interesting or unusual aspect of this picture.	The most interesting aspect of this picture is ____ .	8+3.14
4. image’s past	Describe what happened immediately before this picture was taken.	One or two seconds before this picture was taken, ____ .	9+5.45
5. image’s future	Describe what happened immediately after this picture was taken.	One or two seconds after this picture was taken, ____ .	9+5.04
6. object’s attribute	Describe the appearance of the indicated object.	The object(s) is/are ____ .	3.20+1.62
7. object’s affordance	Describe the function of the indicated object.	People could ____ the object(s).	4.20+1.74
8. object’s position	Describe the position of the indicated object.	The object(s) is/are ____ .	3.20+3.35
9. person’s attribute	Describe the appearance of the indicated person/people.	The person/people is/are ____ .	3+2.52
10. person’s activity	Describe the activity of the indicated person/people.	The person/people is/are ____ .	3+2.47
11. person’s location	Describe the location of the indicated person/people.	The person/people is/are ____ .	3.20+3.04
12. pair’s relationship	Describe the relationship between the indicated person and object.	The person/people is/are ____ the object(s).	5.20+1.65

Table 1: All 12 types of Madlibs instructions and prompts. Right-most column shows the average number of words for each description (#words for prompt + #words for answer).

manner designed to elicit more detailed descriptions of the objects, their interactions, and the broader context of the scene shown in each image.

### Visual Madlibs: Image+Instruction+Prompts+Blank

A single fill-in-the-blank question consists of a prompt and a blank, e.g., *Person A is [blank] the car.* The implicit question is, “What goes in the blank?” This is presented to a person along with an image and instructions, e.g., *Describe the relationship between the indicated person and object.* The same image and prompt may be used with different instructions to collect a variety of description types.

### Instantiating Questions

While the general form of the questions for the Visual Madlibs were chosen by hand, see Table 1, most of the questions are instantiated depending on a subset of the objects present in an image. For instance, if an image contained two people and a dog, questions about each person (question types 9-11 in Table 1), the dog (types 6-8), relationships between the two people and the dog (type 12), could be instantiated. For each possible instantiation, the wording of the questions might alter slightly to maintain grammatical consistency. In addition to these types of questions that depend on the objects present in the image, other questions (types 1-5) can be instantiated for an image regardless of the objects present.

Notice in particular the questions about the temporal context – what might have happened before or what might happen after the image was taken. People can make inferences beyond the specific content depicted in an image. Sometimes these inferences will be consistent between people (e.g., when what will happen next is obvious), and other times these descriptions may be less consistent. We can use the variability of returned responses to select images for which these inferences are reliable.

Asking questions about every object and all pairs of objects quickly becomes unwieldy as the number of objects increases. To combat this, we choose a subset of objects

present to use in instantiating questions. Such selection could be driven by a number of factors. The experiments in this paper consider comparisons to existing, general, descriptions of images, so we instantiate questions about the objects mentioned in those existing natural language descriptions. Whether an object is mentioned in an image description can be viewed as an indication of the object’s importance [2].

### 3.1. Data Collection

To collect the Visual Madlibs Dataset we use a subset of 10,738 human-centric images from MS COCO, that make up about a quarter of the validation data [21], and instantiate fill-in-the-blank templates as described above. The MS COCO images are annotated with a list of objects present in the images, segmentations for the locations of those objects, and 5 general natural language descriptions of the image. To select the subset of images for collecting Madlibs, we start with the 19,338 images with a person labeled. We then look at the five descriptions for each and perform a dependency parse [7], only keeping those images where a word referring to a person (woman, man, etc. E.g., in Fig. 3, guys, men) is the head noun for part of the parse. This leaves 14,150 images. We then filter out the images whose descriptions do not include a synonym for any of the 79 non-person object categories labeled in the MS COCO dataset. This leaves 10,738 human-centric images with at least one other object from the MS COCO data set mentioned in the general image descriptions.

Before final instantiation of the fill-in-the blank templates, we need to resolve a potential ambiguity regarding which objects are referred to in the descriptions. There could be several different people or different instances of an object type labeled in an image. It is not immediately obvious which ones are described in the sentences. To address this assignment problem, we estimate the quantity of each described person/object in the sentence by parsing the determinant (two men and a frisbee in Fig. 3), the conjunction (a man and a woman), and the singular/plural form (dog,



dogs). We compare this number with the number of annotated instances for each category, and consider two possible cases: 1) there are fewer annotated instances than the sentences describe, 2) there are more annotated instances than the sentences describe. It is easy to address the first case, just construct templates for all of the labeled instances. For the second case, we sort the area of each segmented instance, and pick the largest ones up to the parsed number for instantiation. Using this procedure, we obtain 26,148 labeled object or person instances in the 10,738 images.

Each Visual Madlib is answered by 3 workers on Amazon’s Mechanical Turk. To date, we have collected 360,001 answers to Madlib questions. Some example Madlibs answers are shown in Fig. 2,

**Labeled instances:**  
person, person, frisbee, ear, ear, ear

**MS COCO descriptions:**  
 1. Two guys are playing Frisbee in the park.  
 2. Two young men playing a shirtless game of frisbee.  
 3. Two men in a grassy field playing with a frisbee.  
 4. Two men are playing with a frisbee together.  
 5. Two shirtless men playing frisbee in a field.

Figure 3: COCO instance annotation and descriptions for the image of Fig. 1. We show how we map labeled instances to the mentioned person and object in the sentence.

#### 4. Tasks: Multiple-choice question answering and targeted generation

We design two tasks to evaluate targeted natural language description for images. The first task is to automatically generate natural language descriptions of images to fill in the blank for one of the Madlibs questions. This allows for producing targeted descriptions such as: a description specifically focused on the appearance of an object, or a description about the relationship between two objects. The input to this task is an image, instructions, and a Madlibs prompt. As has been discussed at length in the community working on description generation for images, it can be difficult to evaluate free form generation. Our second task tries to address this issue by developing a new targeted multiple-choice question answering task for images. Here the input is again an image, instruction, and a prompt, but instead of a free form text answer, there are a fixed set of multiple-choice answers to fill in the blank. The possible multiple-choice answers are sampled from the Madlibs responses, one that was written for the particular image/instruction/prompt as the correct answer, and distractors chosen from either similar images or random images depending on the level of difficulty desired. This ability to choose distractors to adjust the difficulty of the question as well as the relative ease of evaluating multiple choice answers are attractive aspects of this new task.

In our experiments we randomly select 20% of the 10,738 images to use as our test set for evaluating these tasks. For the multiple-choice questions we form two sets of answers for each, with one set designed to be more difficult than the other. We first establish the easy task distractor answers by randomly choosing three descriptions (of the same question type) from other images [22]. The hard task is designed more delicately. Instead of randomly choosing from the other images, we now only look for those containing the same objects as our question image, and then arbitrarily pick three of their descriptions. Sometimes, the descriptions sampled from “similar” images could also be good answers for our questions (later we experiment with using Turkers to select less ambiguous multiple-choice questions from this set). For the targeted generation task, for question types 1-5, algorithms generate descriptions given the image, instructions, and prompt. For the other question types whose prompts are related to some specific person or object, we additionally provide the algorithm with the location of each person/object mentioned in the prompt. We also experiment with estimating these locations using object detectors.

### 5. Analyzing the Visual Madlibs Dataset

We begin by conducting quantitative analyses of the responses collected in the Visual Madlibs Dataset in Sec. 5.1. A main goal is understanding what additional information is provided by the targeted descriptions in the Visual Madlibs Dataset vs general image descriptions. The MS COCO dataset [21] collects general image descriptions following a similar methodology to previous efforts for collecting general image descriptions, e.g. [28, 34]. So, we provide further analyses comparing the Visual Madlibs to the MS COCO descriptions collected for the same images in Sec. 5.2

#### 5.1. Quantifying Visual Madlibs responses

We analyze the length, structure, and consistency of the Visual Madlibs responses. First, the average length of each type of description is shown in the far right column of Table 1. Note that descriptions of people tend to be longer than descriptions of other objects in the dataset<sup>2</sup>.

Second, we use the phrase chunking [6] to analyze which phrasal structures are commonly used to fill in the blanks for different questions. Fig. 4, top row, shows relative frequencies for the top-5 most frequent templates used for several question types. Object attributes are usually described briefly with a simple adjectival phrase. On the other hand, people use more words and a wider variety of structure to describe possible future events. Except for future and past descriptions, the distribution of structures is generally concentrated on a few likely choices for each question type.

<sup>2</sup>Also note that the length of the prompts varies slightly depending on the object names used to instantiate the Madlib, hence the fractional values in the mean length of the prompts shown in gray.

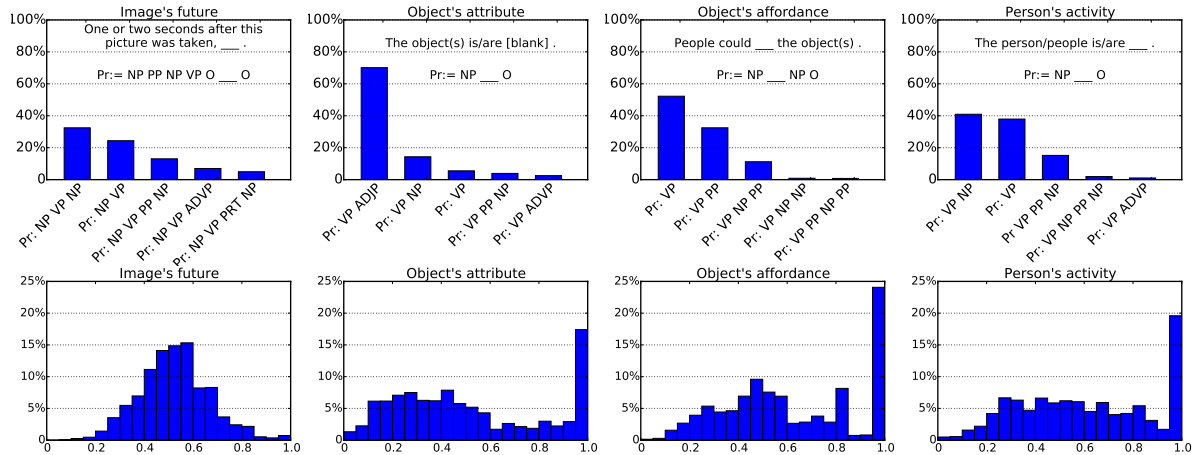


Figure 4: First row shows top-5 most frequent phrase templates for image’s future, object’s attribute, object’s affordance and person’s activity. Second row shows the histograms of similarity between answers.

Third, we analyze how consistent the Mechanical Turk workers’ answers are for each type of question. To compute a measure of similarity between a pair of responses we use the cosine similarity between representations of each response. A response is represented by the mean of the Word2Vec [25] vectors for each word in the response, following [22, 20]. Word2Vec is a 300 dimensional embedding representation for words that encodes the distributional context of words learned over very large word corpora. This measure takes into account the actual words used in a response, as opposed to the previous analyses of parse structure. Each Visual Madlibs question is answered by three workers, providing 3 pairs for which similarity is computed. Fig. 4, bottom row, shows a histogram of all pairwise similarities for several question types. Generally the similarities have a normal-like distribution with an extra peak around 1 indicating the fraction of responses that agree almost perfectly. Once again, descriptions of the future and past are least likely to be (near) identical, while object attributes and affordances are often very consistent.

## 5.2. Visual Madlibs vs general descriptions

We compare the targeted descriptions in the Visual Madlibs Dataset to the general image descriptions in MS COCO. First, we analyze the words used in Visual Madlibs compared to MS COCO descriptions of the same images. For each image, we extract the unique set of words from all descriptions of that image from both datasets, and compute the coverage of each set with respect to the other. We find that on average (across images) 22.45% of the Madlibs’s words are also present in MSCOCO descriptions, while 52.38% of the COCO words are also present in Madlibs.

Second, we compare how Madlibs and MS COCO answers describe the people and objects in images. We ob-

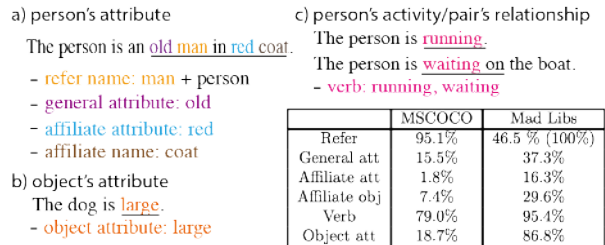


Figure 5: Template used for parsing person’s attributes, activity and interaction with object, and object’s attribute. The percentages below compares Madlibs and MSCOCO on how frequent these templates are used for description.

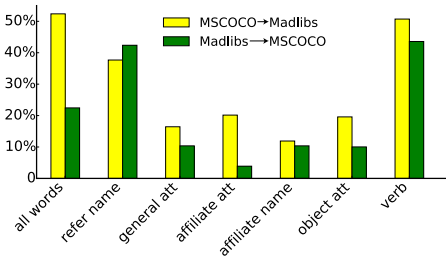


Figure 6: Frequency that a word in a position in the people and object parsing template in one dataset is in the same position for the other dataset.

serve that the Madlibs questions types, Table 1, cover much of the information in MS COCO descriptions [20]. As one way to see this, we run the StanfordNLP parser<sup>3</sup> on both datasets. For attributes of people, we use the parsing template shown in Fig. 5(a) to analyze the structures being used.

<sup>3</sup><http://nlp.stanford.edu/software/lex-parser.shtml>

The *refer name* indicates whether the person was mentioned in the description. Note that the Madlibs descriptions always have one reference to a person in the prompt (The person is [blank]). Therefore, for Madlibs, we report the presence of additional references to the person (e.g., the person is a *man*). The *general attribute* directly describes the appearance of the person or object (e.g., old or small); the *affiliate object* indicates whether additional objects are used to describe the targeted person (e.g. with a bag, coat, or glasses) and the *affiliate attribute* are appearance characteristics of those secondary objects (e.g., red coat). The templates for *object's attribute* and *verbs* are more straightforward as shown in Fig. 5(b)(c). The table in Fig. 5 shows the frequency of each parse component. Overall, more of the potential descriptive elements in these constructions are used in response to the Madlibs prompts than in the general descriptions found in MS COCO.

We also break down the overlap between Visual Madlibs and MS COCO descriptions over different parsing templates for descriptions about people and object (Fig. 6). Yellow bars show how often words for each parse type in MSCOCO descriptions were also found in the same parse type in the Visual Madlibs answers, and green bars measure the reverse direction. Observations indicate that Madlibs provides more coverage in its descriptions than MS COCO for all templates except for person’s refer name. One possible reason is that the prompts already indicates “the person” or “people” explicitly, so workers need not add an additional reference to the person in their descriptions.

**Extrinsic comparison of Visual Madlibs Data and general descriptions:** Here we provide an extrinsic analysis of the information available in the general descriptions compared to Visual Madlibs. We perform this analysis by using either: a) the MS COCO descriptions for an image, or b) Visual Madlibs responses from other Turkers for an image, to select answers for our multiple-choice evaluation task. Specifically, we use one of the human provided descriptions, either from Madlibs or from MS COCO, and select the multiple-choice answer that is most similar to that description. Similarity is measured as cosine similarity between the mean Word2Vec vectors for the words a description compared to the Word2Vec vectors of the multiple-choice answers. In addition to comparing how well the Madlibs or MS COCO descriptions can select the correct multiple-choice answer, we also use the descriptions automatically produced by a recent natural language generation system (CNN+LSTM [32], implementation from [15]) trained on MS COCO dataset. This allows us to make one possible measurement of how close current automatically generated image descriptions are to our Madlibs descriptions. Fig. 7 shows the accuracies resulting from using Madlibs, MSCOCO, or CNN+LSTM [32] to select the correct multiple-choice answer.

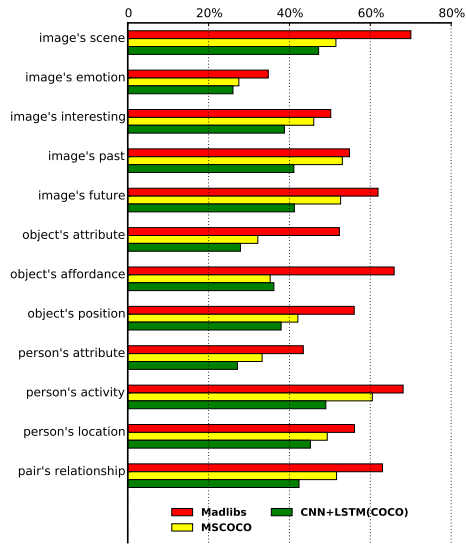


Figure 7: The accuracy of Madlibs, MS COCO and CNN+LSTM [32](trained on MS COCO) used as references to answer the Madlibs hard multiple-choice questions.

Although this approach is quite simple, it allows us we make two interesting observations. First, Madlibs outperforms MS COCO on all types of multiple-choice questions. If Madlibs and MS COCO descriptions provided the same information, we would expect their performance to be comparable. Presumably the performance increase for Madlibs is due to the coverage of targeted descriptions compared to MS COCO’s sentences that describe the overall image content more generally. Second, the automatically generated descriptions from the pre-trained CNN+LSTM perform much worse than the actual MS COCO descriptions, despite doing quite well on general image description generation (The BLEU-1 score of CNN+LSTM, 0.67, is near human agreement 0.69 on MS COCO [32]).

## 6. Experiments

In this section we evaluate a series of methods on the Visual Madlibs Dataset for the targeted natural language generation and multiple-choice question answering tasks, introduced in Sec. 4. As methods, we evaluate simple joint-embedding methods – canonical correlation analysis (CCA) and normalized CCA (nCCA) [14] – as well as a recent deep-learning based method for image description generation – CNN+LSTM [32]. We train these models on 80% of the images in the MadLibs collection and evaluate their performance on the remaining 20%.

In our experiments we extract image features using the VGG Convolutional Neural Network (CNN) [29]. This model has been trained on the ILSVRC-2012 dataset to rec-

Easy Task						
	#Q	CCA	nCCA	nCCA (bbox)	nCCA (all)	CNN+LSTM (madlibs)
1. scene	6277	75.7%	86.8%	—	<b>87.6%</b>	71.1%
2. emotion	5138	41.3%	<b>49.2%</b>	—	42.4%	34.0%
3. past	4903	61.8%	77.5%	—	<b>80.3%</b>	35.8%
4. future	4658	61.2%	78.0%	—	<b>80.2%</b>	40.0%
5. interesting	5095	66.8%	76.5%	—	<b>78.9%</b>	39.8%
6. obj attr	7194	44.1%	47.5%	<b>54.7%</b>	50.9%	45.4%
7. obj aff	7326	59.8%	73.0%	72.2%	<b>76.7%</b>	—
8. obj pos	7290	53.0%	65.9%	58.9%	<b>69.7%</b>	50.9%
9. per attr	6651	40.4%	48.0%	<b>53.1%</b>	44.5%	37.3%
10. per act	6501	70.0%	80.7%	75.6%	<b>82.8%</b>	63.7%
11. per loc	6580	69.8%	82.7%	73.8%	<b>82.7%</b>	59.2%
12. pair rel	7595	54.3%	63.0%	64.2%	<b>67.2%</b>	—

Hard Task						
	#Q	CCA	nCCA	nCCA (bbox)	nCCA (all)	CNN+LSTM (madlibs)
1. scene	6277	63.8%	<b>70.1%</b>	—	68.2%	60.5%
2. emotion	5138	33.9%	<b>37.2%</b>	—	33.2%	32.7%
3. past	4903	47.9%	52.8%	—	<b>54.0%</b>	32.0%
4. future	4658	47.5%	<b>54.3%</b>	—	53.3%	34.3%
5. interesting	5095	51.4%	<b>53.7%</b>	—	<b>55.1%</b>	33.3%
6. obj attr	7194	42.2%	43.6%	<b>49.8%</b>	39.3%	40.3%
7. obj aff	7326	54.5%	63.5%	<b>63.0%</b>	48.5%	—
8. obj pos	7290	49.0%	<b>55.7%</b>	50.7%	53.4%	44.9%
9. per attr	6651	33.9%	38.6%	<b>46.1%</b>	31.6%	36.1%
10. per act	6501	59.7%	65.4%	65.1%	<b>66.6%</b>	53.6%
11. per loc	6580	56.8%	<b>63.3%</b>	57.8%	62.6%	49.3%
12. pair rel	7595	49.4%	54.3%	<b>56.5%</b>	52.0%	—

Filtered questions from Hard						
	#Q	CCA	nCCA	nCCA (bbox)	nCCA (all)	CNN+LSTM (madlibs)
1. scene	4938	70.4%	<b>77.6%</b>	—	76.3%	66.3%
2. emotion	1936	43.7%	<b>49.4%</b>	—	44.2%	34.5%
3. future	3628	52.0%	<b>60.2%</b>	—	59.4%	33.4%
4. past	3811	51.8%	58.0%	—	<b>60.1%</b>	31.1%
5. interesting	4061	56.5%	60.1%	—	<b>61.7%</b>	35.5%
6. obj attr	5313	45.3%	47.1%	<b>54.5%</b>	43.0%	43.4%
7. obj aff	3829	62.7%	<b>72.3%</b>	72.0%	59.2%	—
8. obj pos	5240	53.8%	<b>61.2%</b>	55.5%	58.6%	47.6%
9. per attr	4887	36.5%	42.4%	<b>52.2%</b>	34.4%	37.1%
10. per act	5707	62.1%	68.6%	68.1%	<b>69.9%</b>	54.9%
11. per loc	4992	63.2%	70.2%	63.0%	<b>70.3%</b>	51.5%
12. pair rel	5976	52.2%	57.6%	<b>60.0%</b>	56.5%	—

Table 2: Accuracies computed for different approaches on the easy and hard multiple-choice answering task, and the filtered hard question set. CCA, nCCA, and CNN+LSTM are trained on the whole image representation for each type of question. nCCA(bbox) is trained and evaluated on ground-truth bounding-boxes from COCO segmentations. nCCA(all) trains a single embedding using all question types.

	#Q	Easy Task			Hard Task		
		nCCA	nCCA (bbox)	nCCA (dbox)	nCCA	nCCA (bbox)	nCCA (dbox)
6. obj attr	2021	47.6%	53.6%	51.4%	43.9%	47.9%	45.2%
9. per attr	4206	50.2%	55.4%	51.2%	40.0%	47.0%	43.3%

Table 3: Multiple-choice answering using automatic detection for 42 object/person categories. “bbox” denotes ground-truth bounding box and “dbox” denotes detected bounding box.

ognize images depicting 1000 object classes, and generates a 4,096 dimensional image representation. On the sentence

side, we average the Word2Vec of all words in a sentence to obtain a 300 dimensional representation.

CCA is an approach for finding a joint embedding between two multi-dimensional variables, in our case image and text vector representations. In an attempt to increase the flexibility of the feature selection and for improving computational efficiency, Gong *et al.* [14] proposed a scalable approximation scheme of explicit kernel mapping followed by dimension reduction and linear CCA. In the projected latent space, the similarity is measured by the eigenvalue-weighted normalized correlation. This method, nCCA, provides high-quality retrieval results, improving over the original CCA performance significantly [14].

We train CCA and nCCA models for each question type separately using the training portion of the Visual Madlibs Dataset. These models allow us to map from an image representation, to the joint-embedding space, to vectors in the Word2Vec space, and vice versa. For targeted generation, we map an image to the joint-embedding space and then choose the answer from the training set text that is closest to this embedded point. In order to answer a multiple-choice question we embed each multiple choice answer, and then select the answer whose embedding is closest to image.

Following the recent “Show and Tell” description generation technique [32] (using an implementation from [15]), we train a CNN+LSTM model for each question type on the Visual Madlibs training set. This approach has demonstrated state of the art performance on generating general natural language descriptions for images. These models directly learn a mapping from an image to a sequence of words which we can use to evaluate the targeted generation task. Note that we input the words from the prompt, e.g., The chair is, and then let the CNN+LSTM system generate the remaining words of the description<sup>4</sup>. For the multiple choice task, we compute cosine similarity between Word2Vec representations of the generated description and each question answer and select the most similar answer.

## 6.1. Discussion of results

Table 2 shows accuracies of each algorithm on the easy and hard versions of the multiple-choice task. Fig. 8, shows example correct and wrong answer choices. There are several interesting observations we can make. First, training nCCA on all types of question together, labeled as nCCA(all), is helpful for the easy variant of the task, however it is less useful on the “fine-grained” hard version of the task. Second, extracting visual features from the bounding box of the relevant person/object yields higher accuracy for predicting attributes, but not for other questions. Based on this finding, we try answering the attribute question using automatic detection methods. The detectors are trained on

<sup>4</sup>The missing entries for questions 7 and 12 are due to this priming failing for a fraction of the questions.





Figure 8: Some example question-answering results from nCCA. First row shows correct choices. Second row shows incorrect choices.

	BLEU-1			BLEU-2		
	nCCA	nCCA(bbox)	CNN+LSTM	nCCA	nCCA(box)	CNN+LSTM
1. scene	0.52	—	0.62	0.17	—	0.19
2. emotion	0.17	—	0.39	0	—	0
3. future	0.38	—	0.32	0.12	—	0.08
4. past	0.39	—	0.42	0.12	—	0.11
5. interesting	0.49	—	0.51	0.14	—	0.15
6. obj attr	0.28	0.36	0.45	0.02	0.02	0.01
7. obj aff	0.56	0.60	—	0.10	0.11	—
8. obj pos	0.53	0.55	0.71	0.24	0.25	0.50
9. per attr	0.26	0.29	0.55	0.06	0.07	0.25
10. per act	0.47	0.41	0.52	0.14	0.11	0.22
11. per loc	0.52	0.46	0.64	0.22	0.19	0.39
12. pair rel	0.46	0.48	—	0.07	0.08	—

Table 4: BLEU-1 and BLEU-2 computed on Madlibs testing dataset for different approaches.

ImageNet using R-CNN [13], covering 42 MS COCO categories. We observe similar performance between ground-truth and detected bounding boxes in Table 3.

As an additional experiment we ask humans to answer the multiple choice task, with 5 Turkers answering each question. We use their results to filter out a subset of the hard multiple-choice questions where at least 3 Turkers choose the correct answer. Results of the methods on this subset are shown in Table 2 bottom set of rows. These results show the same pattern as on the unfiltered set, with slightly higher accuracy.

Table 4 shows BLEU-1 and BLEU-2 scores for targeted generation. Although the CNN+LSTM models we trained on Madlibs were not quite as accurate as nCCA for selecting the correct multiple-choice answer, they did result in better, sometimes much better, accuracy (as measured by BLEU scores) for targeted generation.

## 7. Conclusions

We have introduced a new fill-in-the-blank strategy for targeted natural language descriptions and used this to collect a Visual Madlibs dataset. Our analyses show that these descriptions are usually more detailed than generic whole image descriptions. We also introduce a targeted natural language description generation task, and a multiple-choice question answering task, then train and evaluate joint-embedding and generation models. Data produced by this paper will be publicly released upon acceptance.

## Acknowledgement

We thank the vision and language community for feedback regarding this dataset, especially Julia Hockenmaier, Kate Saenko, and Jason Corso. This research is supported by NSF Awards #1417991, 1405822, 144234, and 1452851, and Microsoft Research.

## References

- [1] A. Aker and R. Gaizauskas. Generating image descriptions using dependency relational patterns. In *ACL*, 2010.
- [2] A. C. Berg, T. L. Berg, H. D. III, J. Dodge, A. Goyal, X. Han, A. Mensch, M. Mitchell, A. Sood, K. Stratos, and K. Yamaguchi. Understanding and predicting importance in images. In *CVPR*, 2012.
- [3] A. Bordes, S. Chopra, and J. Weston. Question answering with subgraph embeddings. *arXiv preprint arXiv:1406.3676*, 2014.
- [4] A. Bordes, J. Weston, and N. Usunier. Open question answering with weakly supervised embedding models. In *Machine Learning and Knowledge Discovery in Databases*. Springer, 2014.
- [5] X. Chen and C. L. Zitnick. Learning a recurrent visual representation for image caption generation. *arXiv preprint arXiv:1411.5654*, 2014.
- [6] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuska. Natural language processing (almost) from scratch. *JMLR*, 2011.
- [7] M.-C. De Marneffe, B. MacCartney, C. D. Manning, et al. Generating typed dependency parses from phrase structure parses. In *Proceedings of LREC*, 2006.
- [8] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. *arXiv preprint arXiv:1411.4389*, 2014.
- [9] H. Fang, S. Gupta, F. Iandola, R. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. Platt, et al. From captions to visual concepts and back. *arXiv preprint arXiv:1411.4952*, 2014.
- [10] A. Farhadi, M. Hejrati, M. A. Sadeghi, P. Young, C. Rashtchian, J. Hockenmaier, and D. Forsyth. Every picture tells a story: Generating sentences from images. In *ECCV*. 2010.
- [11] Y. Feng and M. Lapata. Topic models for image annotation and text illustration. In *ACL*, 2010.
- [12] D. Geman, S. Geman, N. Hallonquist, and L. Younes. Visual Turing test for computer vision systems. *Proceedings of the National Academy of Sciences*, 2015.
- [13] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [14] Y. Gong, Q. Ke, M. Isard, and S. Lazebnik. A multi-view embedding space for modeling internet images, tags, and their semantics. *IJCV*, 2014.
- [15] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. *arXiv preprint arXiv:1412.2306*, 2014.
- [16] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [17] N. Krishnamoorthy, G. Malkarnenkar, R. J. Mooney, K. Saenko, and S. Guadarrama. Generating natural-language video descriptions using text-mined knowledge. pages 10–19, July 2013.
- [18] G. Kulkarni, V. Premraj, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg. Baby talk: Understanding and generating image descriptions. In *CVPR*, 2011.
- [19] P. Kuznetsova, V. Ordonez, A. C. Berg, T. L. Berg, and Y. Choi. Collective generation of natural image descriptions. In *ACL*, 2012.
- [20] R. Lebrecht, P. O. Pinheiro, and R. Collobert. Phrase-based image captioning. *CoRR*, abs/1502.03671, 2015.
- [21] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014.
- [22] X. Lin and D. Parikh. Don’t just listen, use your imagination: Leveraging visual common sense for non-visual tasks. *arXiv preprint arXiv:1502.06108*, 2015.
- [23] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *NIPS*, 2014.
- [24] R. Mason. Domain-independent captioning of domain-specific images. In *HLT-NAACL*, 2013.
- [25] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [26] M. Mitchell, X. Han, J. Dodge, A. Mensch, A. Goyal, A. Berg, K. Yamaguchi, T. Berg, K. Stratos, and H. Daumé III. Midge: Generating image descriptions from computer vision detections. In *EACL*, 2012.
- [27] V. Ordonez, G. Kulkarni, and T. L. Berg. Im2text: Describing images using 1 million captioned photographs. In *NIPS*, 2011.
- [28] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier. Collecting image annotations using amazon’s mechanical Turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 139–147. Association for Computational Linguistics, 2010.
- [29] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [30] S. Sukhbaatar, A. Szlam, J. Weston, and R. Fergus. Weakly supervised memory networks. *arXiv preprint arXiv:1503.08895*, 2015.
- [31] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. *arXiv preprint arXiv:1412.4729*, 2014.
- [32] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. *arXiv preprint arXiv:1411.4555*, 2014.
- [33] Y. Yang, C. L. Teo, H. Daumé III, and Y. Aloimonos. Corpus-guided sentence generation of natural images. In *EMNLP*, 2011.
- [34] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *TACL*, 2014.