

BOOSTING IMAGE CAPTIONING WITH ATTRIBUTES

Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, Tao Mei

Microsoft Research Asia

{tiyao, v-yipan, v-yehl, v-zhqi, tmei}@microsoft.com

ABSTRACT

Automatically describing an image with a natural language has been an emerging challenge in both fields of computer vision and natural language processing. In this paper, we present Long Short-Term Memory with Attributes (LSTM-A) - a novel architecture that integrates attributes into the successful Convolutional Neural Networks (CNNs) plus Recurrent Neural Networks (RNNs) image captioning framework, by training them in an end-to-end manner. To incorporate attributes, we construct variants of architectures by feeding image representations and attributes into RNNs in different ways to explore the mutual but also fuzzy relationship between them. Extensive experiments are conducted on COCO image captioning dataset and our framework achieves superior results when compared to state-of-the-art deep models. Most remarkably, we obtain METEOR/CIDEr-D of 25.2%/98.6% on testing data of widely used and publicly available splits in (Karpathy & Fei-Fei, 2015) when extracting image representations by GoogleNet and achieve to date top-1 performance on COCO captioning Leaderboard.

1 INTRODUCTION

Accelerated by tremendous increase in Internet bandwidth and proliferation of sensor-rich mobile devices, image data has been generated, published and spread explosively, becoming an indispensable part of today's big data. This has encouraged the development of advanced techniques for a broad range of image understanding applications. A fundamental issue that underlies the success of these technological advances is the recognition (Szegedy et al., 2015; Simonyan & Zisserman, 2015; He et al., 2016). Recently, researchers have strived to automatically describe the content of an image with a complete and natural sentence, which has a great potential impact for instance on robotic vision or helping visually impaired people. Nevertheless, this problem is very challenging, as description generation model should capture not only the objects or scenes presented in the image, but also be capable of expressing how the objects/scenes relate to each other in a nature sentence.

The main inspiration of recent attempts on this problem (Donahue et al., 2015; Vinyals et al., 2015; Xu et al., 2015; You et al., 2016) are from the advances by using RNNs in machine translation (Sutskever et al., 2014), which is to translate a text from one language (e.g., English) to another (e.g., Chinese). The basic idea is to perform a sequence to sequence learning for translation, where an encoder RNN reads the input sequential sentence, one word at a time till the end of the sentence and then a decoder RNN is exploited to generate the sentence in target language, one word at each time step. Following this philosophy, it is natural to employ a CNN instead of the encoder RNN for image captioning, which is regarded as an image encoder to produce image representations.

While encouraging performances are reported, these CNN plus RNN image captioning methods translate directly from image representations to language, without explicitly taking more high-level semantic information from images into account. Furthermore, attributes are properties observed in images with rich semantic cues and have been proved to be effective in visual recognition (Parikh & Grauman, 2011). A valid question is how to incorporate high-level image attributes into CNN plus RNN image captioning architecture as complementary knowledge in addition to image representations. We investigate particularly in this paper the architectures by exploiting the mutual relationship between image representations and attributes for enhancing image description generation. Specifically, to better demonstrate the impact of simultaneously utilizing the two kinds of representations, we devise variants of architectures by feeding them into RNN in different placements and moments,

e.g., leveraging only attributes, inserting image representations first and then attributes or vice versa, and inputting image representations/attributes once or at each time step.

The main contribution of this work is the proposal of attribute augmented architectures by integrating the attributes into CNN plus RNN image captioning framework, which is a problem not yet fully understood in the literature. By leveraging more knowledge for building richer representations and description models, our work takes a further step forward to enhance image captioning and could have a direct impact of indicating a new direction of vision and language research. More importantly, the utilization of attributes also has a great potential to be an elegant solution of generating open-vocabulary sentences, making image captioning system really practical.

2 RELATED WORK

The research on image captioning has proceeded along three different dimensions: template-based methods (Kulkarni et al., 2013; Yang et al., 2011; Mitchell et al., 2012), search-based approaches (Farhadi et al., 2010; Ordonez et al., 2011; Devlin et al., 2015), and language-based models (Donahue et al., 2015; Kiros et al., 2014; Mao et al., 2014; Vinyals et al., 2015; Xu et al., 2015; Wu et al., 2016; You et al., 2016).

The first direction, template-based methods, predefine the template for sentence generation which follows some specific rules of language grammar and split sentence into several parts (e.g., subject, verb, and object). With such sentence fragments, many works align each part with image content and then generate the sentence for the image. Obviously, most of them highly depend on the templates of sentence and always generate sentence with syntactical structure. For example, Kulkarni *et al.* employ Conditional Random Field (CRF) model to predict labeling based on the detected objects, attributes, and prepositions, and then generate sentence with a template by filling in slots with the most likely labeling (Kulkarni et al., 2013). Similar in spirit, Yang *et al.* utilize Hidden Markov Model (HMM) to select the best objects, scenes, verbs, and prepositions with the highest log-likelihood ratio for template-based sentence generation in (Yang et al., 2011). Furthermore, the traditional simple template is extended to syntactic trees in (Mitchell et al., 2012) which also starts from detecting attributes from image as description anchors and then connecting ordered objects with a syntactically well-formed tree, followed by adding necessary descriptive information.

Search-based approaches “generate” sentence for an image by selecting the most semantically similar sentences from sentence pool or directly copying sentences from other visually similar images. This direction indeed can achieve human-level descriptions as all sentences are from existing human-generated sentences. The need to collect human-generated sentences, however, makes the sentence pool hard to be scaled up. Moreover, the approaches in this dimension cannot generate novel descriptions. For instance, in (Farhadi et al., 2010), an intermediate meaning space based on the triplet of object, action, and scene is proposed to measure the similarity between image and sentence, where the top sentences are regarded as the generated sentences for the target image. Ordonez *et al.* (Ordonez et al., 2011) search images in a large captioned photo collection by using the combination of object, stuff, people, and scene information and transfer the associated sentences to the query image. Recently, a simple k -nearest neighbor retrieval model is utilized in (Devlin et al., 2015) and the best or consensus caption is selected from the returned candidate captions, which even performs as well as several state-of-the-art language-based models.

Different from template-based and search-based models, language-based models aim to learn the probability distribution in the common space of visual content and textual sentence to generate novel sentences with more flexible syntactical structures. In this direction, recent works explore such probability distribution mainly using neural networks for image captioning. For instance, in (Vinyals et al., 2015), Vinyals *et al.* propose an end-to-end neural networks architecture by utilizing LSTM to generate sentence for an image, which is further incorporated with attention mechanism in (Xu et al., 2015) to automatically focus on salient objects when generating corresponding words. More recently, in (Wu et al., 2016), high-level concepts/attributes are shown to obtain clear improvements on image captioning when injected into existing state-of-the-art RNN-based model and such visual attributes are further utilized as semantic attention in (You et al., 2016) to enhance image captioning.

In short, our work in this paper belongs to the language-based models. Different from most of the aforementioned language-based models which mainly focus on sentence generation by solely

depending on image representations (Donahue et al., 2015; Kiros et al., 2014; Mao et al., 2014; Vinyals et al., 2015; Xu et al., 2015) or high-level attributes (Wu et al., 2016), our work contributes by studying not only jointly exploiting image representations and attributes for image captioning, but also how the architecture can be better devised by exploring mutual relationship in between. It is also worth noting that (You et al., 2016) also additionally involve attributes for image captioning. Ours is fundamentally different in the way that (You et al., 2016) is as a result of utilizing attributes to model semantic attention to the locally previous words, as opposed to holistically employing attributes as a kind of complementary representations in this work.

3 BOOSTING IMAGE CAPTIONING WITH ATTRIBUTES

In this paper, we devise our CNN plus RNN architectures to generate descriptions for images under the umbrella of additionally incorporating the detected high-level attributes. Specifically, we begin this section by presenting the problem formulation and followed by five variants of our image captioning frameworks with attributes.

3.1 PROBLEM FORMULATION

Suppose we have an image I to be described by a textual sentence \mathcal{S} , where $\mathcal{S} = \{w_1, w_2, \dots, w_{N_s}\}$ consisting of N_s words. Let $\mathbf{I} \in \mathbb{R}^{D_v}$ and $\mathbf{w}_t \in \mathbb{R}^{D_s}$ denote the D_v -dimensional image representations of the image I and the D_s -dimensional textual features of the t -th word in sentence \mathcal{S} , respectively. Furthermore, we have feature vector $\mathbf{A} \in \mathbb{R}^{D_a}$ to represent the probability distribution over the high-level attributes for image I . Specifically, we train the attribute detectors by using the weakly-supervised approach of Multiple Instance Learning (MIL) in (Fang et al., 2015) on image captioning benchmarks. For an attribute w_a , one image I is regarded as a positive bag of regions (instances) if w_a exists in image I 's ground-truth sentences, and negative bag otherwise. By inputting all the bags into a noisy-OR MIL model, the probability of the bag b_I which contains attribute w_a is measured on the probabilities of all the regions in the bag as

$$\Pr_I^{w_a} = 1 - \prod_{r_i \in b_I} (1 - p_i^{w_a}), \quad (1)$$

where $p_i^{w_a}$ is the probability of the attribute w_a predicted by region r_i and can be calculated through a sigmoid layer after the last convolutional layer in the fully convolutional network. In particular, the dimension of convolutional activations from the last convolutional layer is $x \times x \times h$ and h represents the representation dimension of each region, resulting in $x \times x$ response map which preserves the spatial dependency of the image. Then, a cross entropy loss is calculated based on the probabilities of all the attributes at the top of the whole architecture to optimize MIL model. With the learnt MIL model on image captioning dataset, we treat the final image-level response probabilities of all the attributes as \mathbf{A} .

Inspired by the recent successes of probabilistic sequence models leveraged in statistical machine translation (Bahdanau et al., 2015; Sutskever et al., 2014), we aim to formulate our image captioning models in an end-to-end fashion based on RNNs which encode the given image and/or its detected attributes into a fixed dimensional vector and then decode it to the target output sentence. Hence, the sentence generation problem we explore here can be formulated by minimizing the following energy loss function as

$$E(\mathbf{I}, \mathbf{A}, \mathcal{S}) = -\log \Pr(\mathcal{S}|\mathbf{I}, \mathbf{A}), \quad (2)$$

which is the negative log probability of the correct textual sentence given the image representations and detected attributes.

Since the model produces one word in the sentence at each time step, it is natural to apply chain rule to model the joint probability over the sequential words. Thus, the log probability of the sentence is given by the sum of the log probabilities over the word and can be expressed as

$$\log \Pr(\mathcal{S}|\mathbf{I}, \mathbf{A}) = \sum_{t=1}^{N_s} \log \Pr(\mathbf{w}_t|\mathbf{I}, \mathbf{A}, \mathbf{w}_0, \dots, \mathbf{w}_{t-1}). \quad (3)$$

By minimizing this loss, the contextual relationship among the words in the sentence can be guaranteed given the image and its detected attributes.

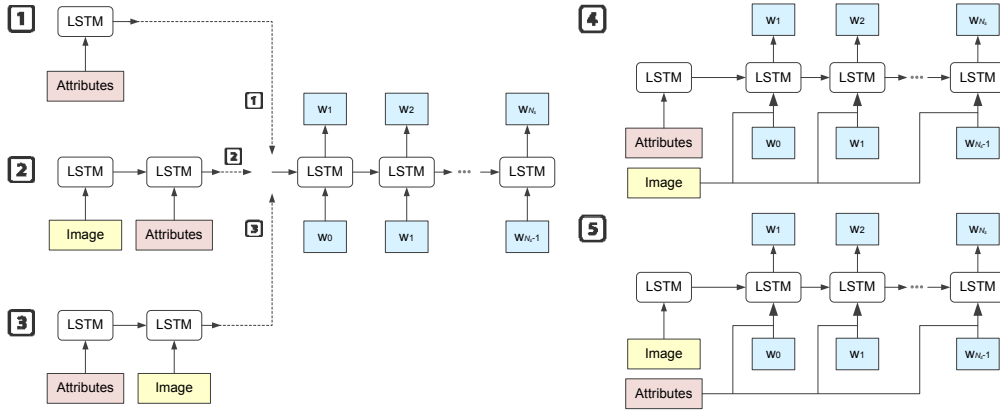


Figure 1: Five variants of our LSTM-A framework (better viewed in color).

We formulate this task as a variable-length sequence to sequence problem and model the parametric distribution $\Pr(\mathbf{w}_t | \mathbf{I}, \mathbf{A}, \mathbf{w}_0, \dots, \mathbf{w}_{t-1})$ in Eq.(3) with Long Short-Term Memory (LSTM), which is a widely used type of RNN. The vector formulas for a LSTM layer forward pass are summarized as below. For time step t , \mathbf{x}^t and \mathbf{h}^t are the input and output vector respectively, \mathbf{T} are input weights matrices, \mathbf{R} are recurrent weight matrices and \mathbf{b} are bias vectors. Sigmoid σ and hyperbolic tangent ϕ are element-wise non-linear activation functions. The dot product of two vectors is denoted with \odot . Given inputs \mathbf{x}^t , \mathbf{h}^{t-1} and \mathbf{c}^{t-1} , the LSTM unit updates for time step t are:

$$\begin{aligned} \mathbf{g}^t &= \phi(\mathbf{T}_g \mathbf{x}^t + \mathbf{R}_g \mathbf{h}^{t-1} + \mathbf{b}_g), & \mathbf{i}^t &= \sigma(\mathbf{T}_i \mathbf{x}^t + \mathbf{R}_i \mathbf{h}^{t-1} + \mathbf{b}_i), \\ \mathbf{f}^t &= \sigma(\mathbf{T}_f \mathbf{x}^t + \mathbf{R}_f \mathbf{h}^{t-1} + \mathbf{b}_f), & \mathbf{c}^t &= \mathbf{g}^t \odot \mathbf{i}^t + \mathbf{c}^{t-1} \odot \mathbf{f}^t, \\ \mathbf{o}^t &= \sigma(\mathbf{T}_o \mathbf{x}^t + \mathbf{R}_o \mathbf{h}^{t-1} + \mathbf{b}_o), & \mathbf{h}^t &= \phi(\mathbf{c}^t) \odot \mathbf{o}^t, \end{aligned}$$

where \mathbf{g}^t , \mathbf{i}^t , \mathbf{f}^t , \mathbf{c}^t , \mathbf{o}^t , and \mathbf{h}^t are cell input, input gate, forget gate, cell state, output gate, and cell output of the LSTM, respectively.

3.2 LONG SHORT-TERM MEMORY WITH ATTRIBUTES

Unlike the existing image captioning models in (Donahue et al., 2015; Vinyals et al., 2015) which solely encode image representations for sentence generation, our proposed Long Short-Term Memory with Attributes (LSTM-A) model additionally integrates the detected high-level attributes into LSTM. We devise five variants of LSTM-A for involvement of two design purposes. The first purpose is about where to feed attributes into LSTM and three architectures, i.e., LSTM-A₁ (leveraging only attributes), LSTM-A₂ (inserting image representations first) and LSTM-A₃ (feeding attributes first), are derived from this view. The second is about when to input attributes or image representations into LSTM and we design LSTM-A₄ (inputting image representations at each time step) and LSTM-A₅ (inputting attributes at each time step) for this purpose. An overview of LSTM-A architectures is depicted in Figure 1.

3.2.1 LSTM-A₁ (LEVERAGING ONLY ATTRIBUTES)

Given the detected attributes, one natural way is to directly inject the attributes as representations at the initial time to inform the LSTM about the high-level attributes. This kind of architecture with only attributes input is named as LSTM-A₁. It is also worth noting that the attributes-based model in (Wu et al., 2016) is similar to LSTM-A₁ and can be regarded as one special case of our LSTM-A. Given the attribute representations \mathbf{A} and the corresponding sentence $\mathbf{W} \equiv [\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{N_s}]$, the LSTM updating procedure in LSTM-A₁ is as

$$\begin{aligned} \mathbf{x}^{-1} &= \mathbf{T}_a \mathbf{A}, \\ \mathbf{x}^t &= \mathbf{T}_s \mathbf{w}_t, \quad t \in \{0, \dots, N_s - 1\} \quad \text{and} \quad \mathbf{h}^t = f(\mathbf{x}^t), \quad t \in \{0, \dots, N_s - 1\}, \end{aligned}$$

where D_e is the dimensionality of LSTM input, $\mathbf{T}_a \in \mathbb{R}^{D_e \times D_a}$ and $\mathbf{T}_s \in \mathbb{R}^{D_e \times D_s}$ is the transformation matrix for attribute representations and textual features of word, respectively, and f is the

updating function within LSTM unit. Please note that for the input sentence $\mathbf{W} \equiv [\mathbf{w}_0, \dots, \mathbf{w}_{N_s}]$, we take \mathbf{w}_0 as the start sign word to inform the beginning of sentence and \mathbf{w}_{N_s} as the end sign word which indicates the end of sentence. Both of the special sign words are included in our vocabulary. Most specifically, at the initial time step, the attribute representations are transformed as the input to LSTM, and then in the next steps, word embedding \mathbf{x}^t will be input into the LSTM along with the previous step's hidden state \mathbf{h}^{t-1} . In each time step (except the initial step), we use the LSTM cell output \mathbf{h}^t to predict the next word. Here a softmax layer is applied after the LSTM layer to produce a probability distribution over all the D_s words in the vocabulary as

$$\Pr_{t+1}(w_{t+1}) = \frac{\exp\{\mathbf{T}_h^{(w_{t+1})}\mathbf{h}^t\}}{\sum_{w \in \mathcal{W}} \exp\{\mathbf{T}_h^{(w)}\mathbf{h}^t\}},$$

where \mathcal{W} is the word vocabulary space and $\mathbf{T}_h^{(w)}$ is the parameter matrix in softmax layer.

3.2.2 LSTM-A₂ (INSERTING IMAGE REPRESENTATIONS FIRST)

To further leverage both image representations and high-level attributes in the encoding stage of our LSTM-A, we design the second architecture LSTM-A₂ by treating both of them as atoms in the input sequence to LSTM. Specifically, at the initial step, the image representations \mathbf{I} are firstly transformed into LSTM to inform the LSTM about the image content, followed by the attribute representations \mathbf{A} which are encoded into LSTM at the next time step to inform the high-level attributes. Then, LSTM decodes each output word based on previous word \mathbf{x}^t and previous step's hidden state \mathbf{h}^{t-1} , which is similar to the decoding stage in LSTM-A₁. The LSTM updating procedure in LSTM-A₂ is designed as

$$\begin{aligned} \mathbf{x}^{-2} &= \mathbf{T}_v \mathbf{I} \quad \text{and} \quad \mathbf{x}^{-1} = \mathbf{T}_a \mathbf{A}, \\ \mathbf{x}^t &= \mathbf{T}_s \mathbf{w}_t, \quad t \in \{0, \dots, N_s - 1\} \quad \text{and} \quad \mathbf{h}^t = f(\mathbf{x}^t), \quad t \in \{0, \dots, N_s - 1\}, \end{aligned}$$

where $\mathbf{T}_v \in \mathbb{R}^{D_e \times D_v}$ is the transformation matrix for image representations.

3.2.3 LSTM-A₃ (FEEDING ATTRIBUTES FIRST)

The third design LSTM-A₃ is similar to LSTM-A₂ as both designs utilize image representations and high-level attributes to form the input sequence to LSTM in the encoding stage, except that the orders of encoding are different. In LSTM-A₃, the attribute representations are firstly encoded into LSTM and then the image representations are transformed into LSTM at the second time step. The whole LSTM updating procedure in LSTM-A₃ is as

$$\begin{aligned} \mathbf{x}^{-2} &= \mathbf{T}_a \mathbf{A} \quad \text{and} \quad \mathbf{x}^{-1} = \mathbf{T}_v \mathbf{I}, \\ \mathbf{x}^t &= \mathbf{T}_s \mathbf{w}_t, \quad t \in \{0, \dots, N_s - 1\} \quad \text{and} \quad \mathbf{h}^t = f(\mathbf{x}^t), \quad t \in \{0, \dots, N_s - 1\}. \end{aligned}$$

3.2.4 LSTM-A₄ (INPUTTING IMAGE REPRESENTATIONS AT EACH TIME STEP)

Different from the former three designed architectures which mainly inject high-level attributes and image representations at the encoding stage of LSTM, we next modify the decoding stage in our LSTM-A by additionally incorporating image representations or high-level attributes. More specifically, in LSTM-A₄, the attribute representations are injected once at the initial step to inform the LSTM about the high-level attributes, and then image representations are fed at each time step as an extra input to LSTM to emphasize the image content frequently among memory cells in LSTM. Hence, the LSTM updating procedure in LSTM-A₄ is:

$$\begin{aligned} \mathbf{x}^{-1} &= \mathbf{T}_a \mathbf{A}, \\ \mathbf{x}^t &= \mathbf{T}_s \mathbf{w}_t + \mathbf{T}_v \mathbf{I}, \quad t \in \{0, \dots, N_s - 1\} \quad \text{and} \quad \mathbf{h}^t = f(\mathbf{x}^t), \quad t \in \{0, \dots, N_s - 1\}. \end{aligned}$$

3.2.5 LSTM-A₅ (INPUTTING ATTRIBUTES AT EACH TIME STEP)

The last design LSTM-A₅ is similar to LSTM-A₄ except that it firstly encodes image representations and then feeds attribute representations as an additional input to LSTM at each step in decoding stage to emphasize the high-level attributes frequently. Accordingly, the LSTM updating procedure in LSTM-A₅ is as

$$\begin{aligned} \mathbf{x}^{-1} &= \mathbf{T}_v \mathbf{I}, \\ \mathbf{x}^t &= \mathbf{T}_s \mathbf{w}_t + \mathbf{T}_a \mathbf{A}, \quad t \in \{0, \dots, N_s - 1\} \quad \text{and} \quad \mathbf{h}^t = f(\mathbf{x}^t), \quad t \in \{0, \dots, N_s - 1\}. \end{aligned}$$

4 EXPERIMENTS

We conducted our experiments on COCO captioning dataset (COCO) (Lin et al., 2014) and evaluated our approaches for image captioning.

4.1 DATASET

The dataset, COCO, is the most popular benchmark for image captioning, which contains 82,783 training images and 40,504 validation images. There are 5 human-annotated descriptions per image. As the annotations of the official testing set are not publicly available, we follow the widely used settings in prior works (You et al., 2016; Zhou et al., 2016) and take 82,783 images for training, 5,000 for validation and 5,000 for testing.

4.2 EXPERIMENTAL SETTINGS

Data Preprocessing. Following (Karpathy & Fei-Fei, 2015), we convert all the descriptions in training set to lower case and discard rare words which occur less than 5 times, resulting in the final vocabulary with 8,791 unique words in COCO dataset.

Features and Parameter Settings. Each word in the sentence is represented as “one-hot” vector (binary index vector in a vocabulary). For image representations, we take the output of 1,024-way $pool5/7 \times 7_{s1}$ layer from GoogleNet (Szegedy et al., 2015) pre-trained on Imagenet ILSVRC12 dataset (Russakovsky et al., 2015). For attribute representations, we select 1,000 most common words on COCO as the high-level attributes and train the attribute detectors with MIL model (Fang et al., 2015) purely on the training data of COCO, resulting in the final 1,000-way vector of probabilities of attributes. The dimension of the input and hidden layers in LSTM are both set to 1,024.

Implementation Details. We mainly implement our image captioning models based on Caffe (Jia et al., 2014), which is one of widely adopted deep learning frameworks. Specifically, with an initial learning rate 0.01 and mini-batch size set 1,024, the objective value can decrease to 25% of the initial loss and reach a reasonable result after 50,000 iterations (about 123 epochs).

Testing Strategies. For sentence generation in testing stage, there are two common strategies. One is to choose the word with maximum probability at each time step and set it as LSTM input for next time step until the end sign word is emitted or the maximum length of sentence is reached. The other strategy is beam search which selects the top- k best sentences at each time step and considers them as the candidates to generate new top- k best sentences at the next time step. We adopt the second strategy and the beam size k is empirically set to 3.

Moreover, to avoid model-level overfitting, we utilize ensembling strategy to fuse the prediction results of 5 identical models as previous works (Vinyals et al., 2015; You et al., 2016). Please note that all the 5 identical models are trained with different initializations separately.

Evaluation Metrics. For the evaluation of our proposed models, we adopt five metrics: BLEU@ N (Papineni et al., 2002), METEOR (Banerjee & Lavie, 2005), ROUGE-L (Lin, 2004), CIDEr-D (Vedantam et al., 2015) and SPICE (Anderson et al., 2016). All the metrics are computed by using the codes¹ released by COCO Evaluation Server (Chen et al., 2015).

4.3 COMPARED APPROACHES

To verify the merit of our LSTM-A models, we compared the following state-of-the-art methods.

(1) NIC & LSTM (Vinyals et al., 2015): NIC attempts to directly translate from image pixels to natural language with a single deep neural network. The image representations are only injected into LSTM at the initial time step. We directly extract the results reported in (You et al., 2016) and name this run as NIC. Furthermore, for fair comparison, we also include one run LSTM which is our implementation of NIC.

¹<https://github.com/tylin/coco-caption>

Table 1: Performance of our proposed models and other state-of-the-art methods on COCO, where $B@N$, M, R, C and S are short for $BLEU@N$, METEOR, ROUGE-L, CIDEr-D and SPICE scores. All values are reported as percentage (%).

Model	B@1	B@2	B@3	B@4	M	R	C	S
NIC (Vinyals et al., 2015)	66.6	45.1	30.4	20.3	-	-	-	-
LRCN (Donahue et al., 2015)	69.7	51.9	38.0	27.8	22.9	50.8	83.7	15.8
HA (Xu et al., 2015)	71.8	50.4	35.7	25	23	-	-	-
SA (Xu et al., 2015)	70.7	49.2	34.4	24.3	23.9	-	-	-
ATT (You et al., 2016)	70.9	53.7	40.2	30.4	24.3	-	-	-
SC (Zhou et al., 2016)	72	54.6	40.4	29.8	24.5	-	95.9	-
LSTM (Vinyals et al., 2015)	68.4	51.2	38	28.4	23.1	50.7	84.3	16
LSTM-A ₁	72.3	55.8	42	31.7	24.9	53.3	96	17.8
LSTM-A ₂	72.8	56.4	42.7	32.2	25	53.5	97.5	18
LSTM-A ₃	73.1	56.4	42.6	32.1	25.2	53.7	98.4	18.2
LSTM-A ₄	71.1	54.5	40.9	30.6	24	52.5	90.6	16.8
LSTM-A ₅	73	56.5	42.9	32.5	25.1	53.8	98.6	18.2
LSTM-A*	95.7	82.5	68.5	55.9	34.1	67.3	150.5	26.8

(2) LRCN (Donahue et al., 2015): LRCN inputs both image representations and previous word into LSTM at each time step for sentence generation.

(3) Hard-Attention (HA) & Soft-Attention (SA) (Xu et al., 2015): Spatial attention on convolutional features of an image is incorporated into the encoder-decoder framework through two kinds of mechanisms: 1) “hard” stochastic attention mechanism equivalently by reinforce learning and 2) “soft” deterministic attention mechanism with standard back-propagation.

(4) ATT (You et al., 2016): ATT utilizes attributes as semantic attention to combine image representations and attributes in RNN for image captioning.

(5) Sentence-Condition (SC) (Zhou et al., 2016): Sentence-condition is proposed most recently and exploits text-conditional semantic attention to generate semantic guidance for sentence generation by conditioning image features on current text content.

(6) MSR Captivator (Devlin et al., 2015): MSR Captivator employs both Multimodal Recurrent Neural Network (MRNN) and Maximum Entropy Language Model (MELM) (Fang et al., 2015) for sentence generation. Deep Multimodal Similarity Model (DMSM) (Fang et al., 2015) is further exploited for sentence re-ranking.

(7) CaptionBot (Tran et al., 2016): CaptionBot is a publicly image captioning system² which is mainly built on vision models by using Deep residual networks (ResNets) (He et al., 2016) to detect visual concepts, MELM (Fang et al., 2015) language model for sentence generation and DMSM (Fang et al., 2015) for caption ranking. Entity recognition model for celebrities and landmarks is further incorporated to enrich captions and the confidence scoring model is finally utilized to select the output caption.

(8) LSTM-A: LSTM-A₁, LSTM-A₂, LSTM-A₃, LSTM-A₄, and LSTM-A₅ are five variants derived from our proposed LSTM-A framework. In addition, LSTM-A* is an oracle run that inputs ground-truth attributes in the LSTM-A₃ architecture.

4.4 PERFORMANCE COMPARISON

Performance on COCO Table 1 shows the performances of different models on COCO image captioning dataset. It is worth noting that the performances of different approaches here are based on different image representations. Specifically, VGG architecture (Simonyan & Zisserman, 2015) is utilized as image feature extractor in the methods of Hard-Attention & Soft-Attention and Sentence-Condition, while GoogleNet (Szegedy et al., 2015) is exploited in NIC, LRCN, ATT, LSTM and our LSTM-A. In view that the GoogleNet and VGG features are comparable, we compare directly with results. Overall, the results across eight evaluation metrics consistently indicate that our proposed LSTM-A exhibits better performance than all the state-of-the-art techniques including non-attention models (NIC, LSTM, LRCN) and attention-based methods (Hard-Attention, Soft-Attention, ATT, Sentence-Condition). In particular, the CIDEr-D can achieve 98.6%, which is to date the high-

²<https://www.captionbot.ai>

Table 2: Leaderboard of the published state-of-the-art image captioning models on the online COCO testing server (<http://mscoco.org/dataset/#captions-leaderboard>), where B@N, M, R, and C are short for BLEU@N, METEOR, ROUGE-L, and CIDEr-D scores. All values are reported as percentage (%).

Model	B@1		B@2		B@3		B@4		M		R		C	
	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40
MSM@MSRA (LSTM-A₃)	75.1	92.6	58.8	85.1	44.9	75.1	34.3	64.6	26.6	36.1	55.2	70.9	104.9	105.3
ATT (You et al., 2016)	73.1	90	56.5	81.5	42.4	70.9	31.6	59.9	25	33.5	53.5	68.2	94.3	95.8
Google (Vinyals et al., 2015)	71.3	89.5	54.2	80.2	40.7	69.4	30.9	58.7	25.4	34.6	53	68.2	94.3	94.6
MSR Captivator (Devlin et al., 2015)	71.5	90.7	54.3	81.9	40.7	71	30.8	60.1	24.8	33.9	52.6	68	93.1	93.7

Table 3: The user study using two criteria: M1 - percentage of captions generated by different methods that are evaluated as better or equal to human caption and M2 - percentage of captions that pass the Turing Test.

	Human	LSTM-A ₃	CaptionBot	LSTM	LRCN
M1	-	62.8	58.2	49.2	43.9
M2	90.1	72.2	66.3	57.3	55.9

est performance reported on COCO dataset when extracting image representations by GoogleNet. LSTM-A₁ inputting only high-level attributes as representations makes the relative improvement over LSTM which feeds into image representations instead by 11.6%, 7.8%, 5.1%, 13.9% and 11.25% in BLEU@4, METEOR, ROUGE-L, CIDEr-D and SPICE, respectively. The results basically indicate the advantage of exploiting high-level attributes than image representations for image captioning. Furthermore, by additionally incorporating attributes to LSTM model, LSTM-A₂, LSTM-A₃ and LSTM-A₅ lead to a performance boost, indicating that image representations and attributes are complementary and thus have mutual reinforcement for image captioning. Similar in spirit, LSTM-A₄ improves LRCN by further taking attributes into account. There is a significant performance gap between ATT and LSTM-A₅. Though both runs involve the utilization of image representations and attributes, they are fundamentally different in the way that the performance of ATT is as a result of modulating the strength of attention on attributes to the previous words, and LSTM-A₅ is by employing attributes as auxiliary knowledge to complement image representations. This somewhat reveals the weakness of semantic attention model, where the prediction errors will accumulate along the generated sequence.

Compared to LSTM-A₁, LSTM-A₂ which is augmented by integrating image representations performs better, but the performances are lower than LSTM-A₃. The results indicate that LSTM-A₃, in comparison, benefits from the mechanism of first feeding high-level attributes into LSTM instead of starting from inserting image representations in LSTM-A₂. The chance that a good start point can be attained and lead to performance gain is better. LSTM-A₄ feeding the image representations at each time step yields inferior performances to LSTM-A₃, which only inputs image representations once. We speculate that this may be because the noise in the image can be explicitly accumulated and thus the network overfits more easily. In contrast, the performances of LSTM-A₅ which feeds attributes at each time step show the improvements on LSTM-A₃. The results demonstrate that the high-level attributes are more accurate and easily translated into human understandable sentence. Among the five proposed LSTM-A architectures, LSTM-A₃ achieves the best performances in terms of BLEU@1 and METEOR, while LSTM-A₅ performs the best in other six evaluation metrics. The performances of the oracle run LSTM-A* could be regarded as the upper bound of employing attributes in our framework and lead to large performance gain against LSTM-A₃. Such an upper bound enables us to obtain more insights on the factor accounting for the success of the current attribute augmented architecture and also provides guidance to future research in this direction. More specifically, the results, on one hand, indicate the advantage and great potential of leveraging attributes for boosting image captioning, and on the other, suggest that more efforts are further required towards mining and representing attributes more effectively.

Performance on COCO online testing server We also submitted our best run in terms of METEOR, i.e., LSTM-A₃, to online COCO testing server and evaluated the performance on official testing set. Table 2 shows the performance Leaderboard on official testing image set with 5 reference captions (c5) and 40 reference captions (c40). Please note that here we utilize the outputs of 2,048-way *pool5* layer from ResNet-152 as image representations and train the attribute detectors by ResNet-152 in our final submission. Only the latest top-3 performing methods which have been officially published are included in the table. Compared to the top performing methods, our proposed LSTM-A₃ achieves the best performance across all the evaluation metrics on both c5 and c40 testing sets, and ranks the first on the Leaderboard.




Figure 2: Visualization of prediction changes w.r.t. the additional attribute inputs. The attributes are predicted by MIL method in (Fang et al., 2015) and the output sentences are generated by LSTM and LSTM-A₃.







	<p>Ground Truth: a pile of stuffed animals hanging from the ceiling of a store</p> <p>LSTM-A₁: a bunch of stuffed animals in a room</p> <p>LSTM-A₂: a bunch of stuffed animals in a store</p> <p>LSTM-A₃: a bunch of stuffed animals hanging from a ceiling</p> <p>LSTM-A₄: a group of people standing next to each other</p> <p>LSTM-A₅: a bunch of stuffed animals hanging from a ceiling</p>		<p>Ground Truth: a group of children sitting at a table eating pieces of cake</p> <p>LSTM-A₁: a little girl sitting at a table eating food</p> <p>LSTM-A₂: a group of people sitting around a table eating cake</p> <p>LSTM-A₃: a group of children sitting around a table eating cake</p> <p>LSTM-A₄: a little girl sitting at a table with a cake</p> <p>LSTM-A₅: a group of children sitting at a table eating food</p>
	<p>Ground Truth: some brown cabinets a table light and refrigerator</p> <p>LSTM-A₁: a kitchen with a stove and a counter</p> <p>LSTM-A₂: a kitchen with a table and a stove</p> <p>LSTM-A₃: a kitchen with a wooden table and a refrigerator</p> <p>LSTM-A₄: a kitchen with a stove and a stove</p> <p>LSTM-A₅: a kitchen with a stove and a refrigerator</p>		<p>Ground Truth: elephants walk in a line down the middle of the street</p> <p>LSTM-A₁: a group of people riding elephants down a street</p> <p>LSTM-A₂: a group of people walking down a street</p> <p>LSTM-A₃: a group of elephants walking down a street</p> <p>LSTM-A₄: a group of people riding horses down a street</p> <p>LSTM-A₅: a group of elephants walking down a street</p>
	<p>Ground Truth: many people behind a stand selling bananas</p> <p>LSTM-A₁: a man standing in front of a bunch of bananas</p> <p>LSTM-A₂: a group of people standing around a bunch of bananas</p> <p>LSTM-A₃: a group of people standing around a table with bananas</p> <p>LSTM-A₄: a group of people standing around a table</p> <p>LSTM-A₅: a group of people standing around a table with bananas</p>		<p>Ground Truth: a man sitting on a couch holding a will remote control</p> <p>LSTM-A₁: a man sitting on a couch in a living room</p> <p>LSTM-A₂: a man sitting on a couch with a laptop</p> <p>LSTM-A₃: a man sitting on a couch with a remote</p> <p>LSTM-A₄: a woman sitting on a couch with a laptop</p> <p>LSTM-A₅: a man sitting on a couch with a remote</p>

Figure 3: Sentence examples generated by our five LSTM-A architectures and one ground truth sentence.

4.5 HUMAN EVALUATION

To better understand how satisfactory are the sentences generated from different methods, we also conducted a human study to compare our LSTM-A₃ against three approaches, i.e., CaptionBot, L-RCN and LSTM. A total number of 12 evaluators (6 females and 6 males) from different education backgrounds, including computer science (4), business (2), linguistics (2) and engineering (4), are invited and a subset of 1,000 images is randomly selected from testing set for the subjective evaluation. The evaluation process is as follows. All the evaluators are organized into two groups. We show the first group all the four sentences generated by each approach plus the five human-annotated sentences and ask them the question: Do the systems produce captions resembling human-generated sentences? In contrast, we show the second group once only one sentence generated by different approach or human annotation and they are asked: Can you determine whether the given sentence has been generated by a system or by a human being? From evaluators' responses, we calculate two metrics: 1) M1: percentage of captions that are evaluated as better or equal to human caption; 2) M2: percentage of captions that pass the Turing Test. Table 3 lists the result of the user study. Overall, our LSTM-A₃ is clearly the winner for all two criteria. In particular, the percentage achieves 62.8% and 72.2% in terms of M1 and M2, respectively, making the absolute improvement over the best competitor CaptionBot by 4.6% and 5.9%.

4.6 QUALITATIVE ANALYSIS

Visualization of prediction changes w.r.t. the additional attribute inputs. Figure 2 shows two image examples to illustrate the word prediction changes with respect to the additional attribute








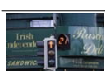
	Attributes: boat: 1 water: 0.92 river: 0.645 small: 0.606 boats: 0.562 dog: 0.555 body: 0.527 floating: 0.484	Generated Sentences: LSTM: a group of people on a boat in the water CaptionBot: I think it's a man with a small boat in a body of water. LSTM-A₃: a man and a dog on a boat in the water	Ground Truth: ① an image of a man in a boat with a dog ② a person on a rowboat with a dalmatian dog on the boat ③ old woman rowing a boat with a dog
	Attributes: bananas: 1 market: 0.995 outdoor: 0.617 bunch: 0.553 table: 0.51 flowers: 0.454 people: 0.431 yellow: 0.377	Generated Sentences: LSTM: a group of people standing around a market CaptionBot: I think it's a bunch of yellow flowers. LSTM-A₃: a group of people standing around a bunch of bananas	Ground Truth: ① bunches of bananas for sale at an outdoor market ② a person at a table filled with bananas ③ there are many bananas layer across this table at a farmers market
	Attributes: man: 0.669 herd: 0.583 standing: 0.496 animals: 0.493 walking: 0.471 cows: 0.427 street: 0.427 road: 0.414	Generated Sentences: LSTM: a man riding a skateboard down a street CaptionBot: I think it's a group of people walking down the road. LSTM-A₃: a man walking down a road with a sheep	Ground Truth: ① a man walks while a large number of sheep follow ② a man leading a herd of sheep down the sheep ③ the man is walking a herd of sheep on the road through a town
	Attributes: computer: 0.764 keyboard: 0.748 screen: 0.627 holding: 0.558 person: 0.515 phone: 0.434 hand: 0.404 remote: 0.395	Generated Sentences: LSTM: a cell phone sitting on top of a table CaptionBot: I think it's a laptop that is on the phone. LSTM-A₃: a person holding a cell phone in their hand	Ground Truth: ① a smart phone being held up in front of a lap top ② the person is holding his cell phone while on his laptop ③ someone holding a cell phone in front of a laptop
	Attributes: flying: 0.877 plane: 0.598 airplane: 0.528 lake: 0.495 water: 0.462 sky: 0.443 red: 0.426 small: 0.365	Generated Sentences: LSTM: a group of people flying kites in the sky CaptionBot: I think it's a plane is flying over the water. LSTM-A₃: a red and white plane flying over a body of water	Ground Truth: ① a plane with water skies for landing gear coming in for a landing at a lake ② a plane flying through a sky above a lake ③ a red and white plane is flying over some water
	Attributes: snow: 0.996 standing: 0.841 zebra: 0.828 enclosure: 0.629 zoo: 0.492 ground: 0.490 area: 0.417 walking: 0.322	Generated Sentences: LSTM: a zebra is standing in a dirt area CaptionBot: I think it's a zebra is standing in the snow. LSTM-A₃: a zebra standing in the snow near a fence	Ground Truth: ① one zebra standing in snow near a stone wall ② a zebra is standing in a snowy field ③ a zebra stands in snow in front of a wall
	Attributes: holding: 0.958 walking: 0.725 people: 0.493 umbrella: 0.548 woman: 0.421 person: 0.361 standing: 0.289 playing: 0.271	Generated Sentences: LSTM: a man walking down a street holding a surfboard CaptionBot: I am not really confident, but I think it's a man walking down a sidewalk holding an umbrella. LSTM-A₃: a group of people walking down a street holding umbrellas	Ground Truth: ① a couple of kids walking with umbrellas in their hands ② two women walking side by side holding umbrellas ③ an image of two girls walking with umbrellas
	Attributes: traffic: 0.746 sign: 0.690 street: 0.555 light: 0.446 signs: 0.374 building: 0.344 stop: 0.313 pole: 0.296	Generated Sentences: LSTM: a street sign that is on a pole CaptionBot: I think it's a sign hanging off the side of a building. LSTM-A₃: a street sign with a traffic light on it	Ground Truth: ① a traffic light in front of some business signs ② a traffic light near a store front ③ a traffic light atop a post in a business district

Figure 4: Attributes and sentences generation results on COCO. The attributes are predicted by MIL method in (Fang et al., 2015) and the output sentences are generated by 1) LSTM, 2) CaptionBot², 3) our LSTM-A₃, and 4) Ground Truth: randomly selected three ground truth sentences.


inputs. Take the first image as an example, the predicted subject is “a cake” in LSTM model. By additionally incorporating the detected attributes, e.g., “candles” and “birthday,” the output subject in the sentence by our LSTM-A₃ changes into “a birthday cake with candles,” demonstrating the advantage of the auxiliary attribute inputs.

Sentence generation comparison between five LSTM-A architectures. The examples of sentence generated by our five LSTM-A architectures are further illustrated in Figure 3. In general, the sentences generated by LSTM-A₃ and LSTM-A₅ are very comparable and more accurate than those by LSTM-A₁, LSTM-A₂ and LSTM-A₄. For instance, LSTM-A₃ and LSTM-A₅ produce the sentence of “a bunch of stuffed animals hanging from a ceiling,” which describes the first image very precisely and finely.

Sentence generation comparison across different approaches. Figure 4 showcases a few sentence examples generated by different methods, the detected high-level attributes, and human-annotated ground truth sentences. From these exemplar results, it is easy to see that all of these automatic methods can generate somewhat relevant sentences, while our proposed LSTM-A₃ can predict more relevant keywords by jointly exploiting high-level attributes and image representations for image captioning. For example, compared to subject term “a group of people” and “a man” in the sentence generated by LSTM and CaptionBot respectively, “a man and a dog” in our LSTM-A₃ is more precise to describe the image content in the first image, since the keyword “dog” is one of the detected attributes and directly injected into LSTM to guide the sentence generation. Similarly, verb term “holding” which is also detected as one high-level attribute presents the fourth image more exactly. Moreover, our LSTM-A₃ can generate more descriptive sentence by enriching the semantics with high-level attributes. For instance, with the detected adjective “red,” the generated sentence “a red and white plane flying over a body of water” of the fifth image depicts the image content more comprehensive.

4.7 ANALYSIS OF THE BEAM SIZE k

In order to analyze the effect of the beam size k in testing stage, we illustrate the performances of our two top performing architectures LSTM-A₃ and LSTM-A₅ with the beam size in the range of $\{1, 2, 3, 4, 5\}$ in Figure 5. To make all performances fall into a comparable scale, all scores are normalized by the highest score of each evaluation metric. As shown in Figure 5, we can see that

Figure 5: The effect of beam size k on (a) LSTM-A₃ and (b) LSTM-A₅.

almost all performances in terms of each evaluation metric are like the “^” shapes when beam size k varies from 1 to 5. Hence, we set the beam size k as 3 in our experiments, which can achieve the best performance with a relatively small beam size.

5 DISCUSSIONS AND CONCLUSIONS

We have presented Long Short-Term Memory with Attributes (LSTM-A) architectures which explores both image representations and high-level attributes for image captioning. Particularly, we study the problem of augmenting high-level attributes from images to complement image representations for enhancing sentence generation. To verify our claim, we have devised variants of architectures by modifying the placement and moment, where and when to feed into the two kinds of representations. Experiments conducted on COCO image captioning dataset validate our proposal and analysis. Performance improvements are clearly observed when comparing to other captioning techniques and more remarkably, the performance of our LSTM-A to date ranks the first on COCO image captioning Leaderboard.

Our future works are as follows. First, more attributes will be learnt from large-scale image benchmarks, e.g., YFCC-100M dataset, and integrated into image captioning. We will further analyze the impact of different number of attributes involved. Second, how to generate free-form and open-vocabulary sentences with the learnt attributes is also expected.

REFERENCES

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *ECCV*, 2016.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- Jacob Devlin, Hao Cheng, Hao Fang, Saurabh Gupta, Li Deng, Xiaodong He, Geoffrey Zweig, and Margaret Mitchell. Language models for image captioning: The quirks and what works. In *ACL*, 2015.
- Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.
- Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, C. Lawrence Zitnick, and Geoffrey Zweig. From captions to visual concepts and back. In *CVPR*, 2015.
- Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. Every picture tells a story: Generating sentences from images. In *ECCV*, 2010.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *MM*, 2014.
- Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.
- Ryan Kiros, Ruslan Salakhutdinov, and Rich Zemel. Multimodal neural language models. In *ICML*, 2014.
- Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. Babytalk: Understanding and generating simple image descriptions. *IEEE Trans. on PAMI*, 2013.
- Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *ACL Workshop*, 2004.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, and Alan L. Yuille. Explain images with multimodal recurrent neural networks. In *NIPS Workshop on Deep Learning*, 2014.
- Margaret Mitchell, Xufeng Han, Jesse Dodge, Alyssa Mensch, Amit Goyal, Alex Berg, Kota Yamaguchi, Tamara Berg, Karl Stratos, and Hal Daumé III. Midge: Generating image descriptions from computer vision detections. In *EACL*, 2012.
- Vicente Ordonez, Girish Kulkarni, and Tamara L Berg. Im2text: Describing images using 1 million captioned photographs. In *NIPS*, 2011.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002.
- Devi Parikh and Kristen Grauman. Relative attributes. In *ICCV*, 2011.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- Kenneth Tran, Xiaodong He, Lei Zhang, Jian Sun, Cornelia Carapcea, Chris Thrasher, Chris Buehler, and Chris Sienkiewicz. Rich image captioning in the wild. *arXiv preprint arXiv:1603.09016*, 2016.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015.
- Qi Wu, Chunhua Shen, Lingqiao Liu, Anthony Dick, and Anton van den Hengel. What value do explicit high level concepts have in vision to language problems? In *CVPR*, 2016.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- Yezhou Yang, Ching Lik Teo, Hal Daumé III, and Yiannis Aloimonos. Corpus-guided sentence generation of natural images. In *EMNLP*, 2011.
- Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *CVPR*, 2016.
- Luowei Zhou, Chenliang Xu, Parker Koch, and Jason J Corso. Image caption generation with text-conditional semantic attention. *arXiv preprint arXiv:1606.04621*, 2016.