

# Reflective Decoding Network for Image Captioning

Lei Ke<sup>1</sup>, Wenjie Pei<sup>2</sup>, Ruiyu Li<sup>2</sup>, Xiaoyong Shen<sup>2</sup>, Yu-Wing Tai<sup>2</sup>

<sup>1</sup>The Hong Kong University of Science and Technology, <sup>2</sup>Tencent

keleiwhu@gmail.com, wenjiecoder@outlook.com, {royryli, dylanshen, yuwingtai}@tencent.com

## Abstract

State-of-the-art image captioning methods mostly focus on improving visual features, less attention has been paid to utilizing the inherent properties of language to boost captioning performance. In this paper, we show that vocabulary coherence between words and syntactic paradigm of sentences are also important to generate high-quality image caption. Following the conventional encoder-decoder framework, we propose the Reflective Decoding Network (RDN) for image captioning, which enhances both the long-sequence dependency and position perception of words in a caption decoder. Our model learns to collaboratively attend on both visual and textual features and meanwhile perceive each word’s relative position in the sentence to maximize the information delivered in the generated caption. We evaluate the effectiveness of our RDN on the COCO image captioning datasets and achieve superior performance over the previous methods. Further experiments reveal that our approach is particularly advantageous for hard cases with complex scenes to describe by captions.

## 1. Introduction

The goal of image captioning is to automatically generate fluent and informative language description of an image for human understanding. As an interdisciplinary task connecting Computer Vision and Nature Language Processing, it explores towards the cutting edge techniques of scene understanding [23] and it is drawing increasing interests in recent years.

To build a top captioning system, there are two crucial requirements. First, the captioning model needs to distill representative and meaningful visual representation from an image. Thanks to the success in image classification [20] and object recognition [12, 33], recent methods [2, 26, 45, 47] have shown significant advancements which mostly benefited from the improved quality of extracted visual features. Second, and the relatively neglected requirement, is to make the generated captions coherent and intelligent. Similar to the human language system, it needs to



**Basis decoder:** A black and white photo of a clock tower in the background.

**Ours:** A view of a **bridge** with a clock tower over a **river**.

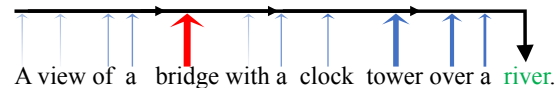


Figure 1. **Top:** Example captions generated by the basis decoder (using traditional LSTM) and our Reflective Decoding Network model. **Bottom:** The reflective attention weight distribution over the past generated hidden states is shown when predicting the word ‘river’. The thicker line indicates a relatively larger weight and the red line means the largest contribution to the prediction.

infer and reason during the generation process based on what has been generated and watched. Typically, this process is achieved by RNN (specifically, LSTM [14]) in storing the sequential information during caption decoding.

The traditional LSTM model, however, tends to focus more on the relatively closer vocabulary while neglecting the farther one. For example, in Figure 1, the word ‘bridge’ has an important hint on predicting the word ‘river’ (which is neglected by the basis decoder), but the two words are separated by 6 words. Current mainstream caption decoder is weak in handling this kind of long-term dependency in sequential sentence, especially when the visual content of an image is complex and hard to describe, which usually leads to a general and less accurate caption description.

In this paper, we propose the Reflective Decoding Network (RDN) for image captioning, which mitigates the drawback of traditional caption decoder by enhancing its long sequential modeling ability. Different from previous methods which boost captioning performance by improving the visual attention mechanism [2, 26, 45], or by improving the encoder to supply more meaningful intermediate representation for the decoder [17, 47, 48, 50], our RDN focuses directly on the target decoding side and jointly apply attention mechanism in both visual and textual domain.

<sup>1</sup>This work was done while Lei Ke was an intern at Tencent.

Besides, we propose to model the positional information of each word within a caption in a supervised way to capture the syntactic structure of natural language. Another advantage in RDN is to visualize how the model inferences and makes word prediction based on the generated words. For instance, our RDN successfully decodes the word ‘river’ in Figure 1 by referring to the previously generated words, especially the most relevant word ‘bridge’.

The main contributions of this paper are four folds:

- We propose the RDN that effectively enhances the long sequential modeling ability of the traditional caption decoder for generating high-quality image captions.
- By considering long-term textual attention, we explicitly explore the coherence between words and visualize the word decision making process in text domain to show how we can interpret the principle and result of the framework from a novel perspective.
- We design a novel positional module to enable our RDN to perceive the relative position of each word in the whole caption and thereby better comprehend the syntactic paradigm of natural language.
- Our RDN achieves state-of-the-art performance on COCO captioning dataset and is particularly superior over existing methods in hard cases with complex scenes to describe by captions.

## 2. Related Work

**Image Captioning.** State-of-the-art captioning methods are mostly driven by advancements in machine translation [7, 36], where the encoder-decoder framework has demonstrated to generate much more novel and coherent sentences compared to the traditional template-based [21, 46] or search-based [9] methods. In [10, 42], the authors introduced a framework which utilizes a pre-trained CNN as an encoder to extract image features, followed by an RNN as a decoder to generate image descriptions. This model was further improved by incorporating high-level semantic attribute information [44, 49] or regularizing the RNN decoder [6]. To distill the salient objects or important regions from an image, different kinds of attention mechanisms were integrated into the captioning framework to exam the relevant image regions when generating sentences [2, 26, 45, 47, 50].

Fusion learning of multiple encoders or decoders forms an essential part of boosting image captioning performance. In [17], the authors utilized multiple CNNs to extract complementary image features, which forms a more informative and integrated representation for decoder. Yao *et al.* [48] proposed GCN-LSTM to build two kinds of graphs to incorporate both semantic and spatial relations into the framework. The outputs from two different separately trained decoders are linearly fused to produce the final prediction.

Similar to [2, 28, 48], our RDN also utilizes the attention mechanism and follows the encoder-decoder framework. However, we explicitly study the coherence between words, which remedies the drawback of current captioning framework in modeling long-term dependency in decoder.

**Language Attention in joint vision and language tasks.** Learning language attention has attracted increasing attention in other joint vision and language problems, such as VQA and grounding referential expressions. In [27], the authors proposed a model to jointly reason both visual and language attentions for visual question answering. Yu *et al.* [51] attentively parsed the expressions into three phrase embeddings to address the task of referring expression comprehension. Different from them, image captioning task is a sequential language generation process. The target description of an image is unknown during inference stage. So, our RDN explores the language attention based on the generated words in previous states. With more time steps, the attended language content will increase dynamically, which enables the word predicted later to capture more useful information for reference.

**Language Attention in NLP tasks.** Our RDN shares some ideas of the self-attention mechanism in machine translation models [30, 39, 40], abstract summarization model [32] and dialogue system [29]. A typical self-attention model such as Transformer [40] aims to learn a latent representation for each position of a sequence by referring to the whole context. In contrast, the Reflective Attention Module (RAM) of our RDN is designed as an attachable module which is seamlessly integrated into the recurrent decoding framework. Thanks to our special two-layer recurrent structure, our RAM collaborates smoothly with the visual attention component of our RDN by sharing the same query value to optimize the captioning process jointly, which is beneficial to ensure our generated captions match with the visual content of an image. To our knowledge, this paper is the first work in jointly exploring both visual and language attention in image captioning.

## 3. Reflective Decoding Network

The overall architecture of our framework is shown in Figure 2. Given an input image, our model first uses Faster R-CNN [33] as Encoder to obtain the visual features of objects in the image. The visual features are then fed to the our Reflective Decoding Network (RDN) to generate caption. Our RDN contains three components: (1) Attention-based Recurrent Module, which attends to the visual features from Encoder; (2) Reflective Attention Module, which employs textual attention to model the compatibility between current and past decoding hidden states, thus it is able to capture more historical and comprehensive information for word decision; (3) Reflective Position Module, which introduces relative position information for each word in the generated caption and helps the model to perceive the syntactic struc-

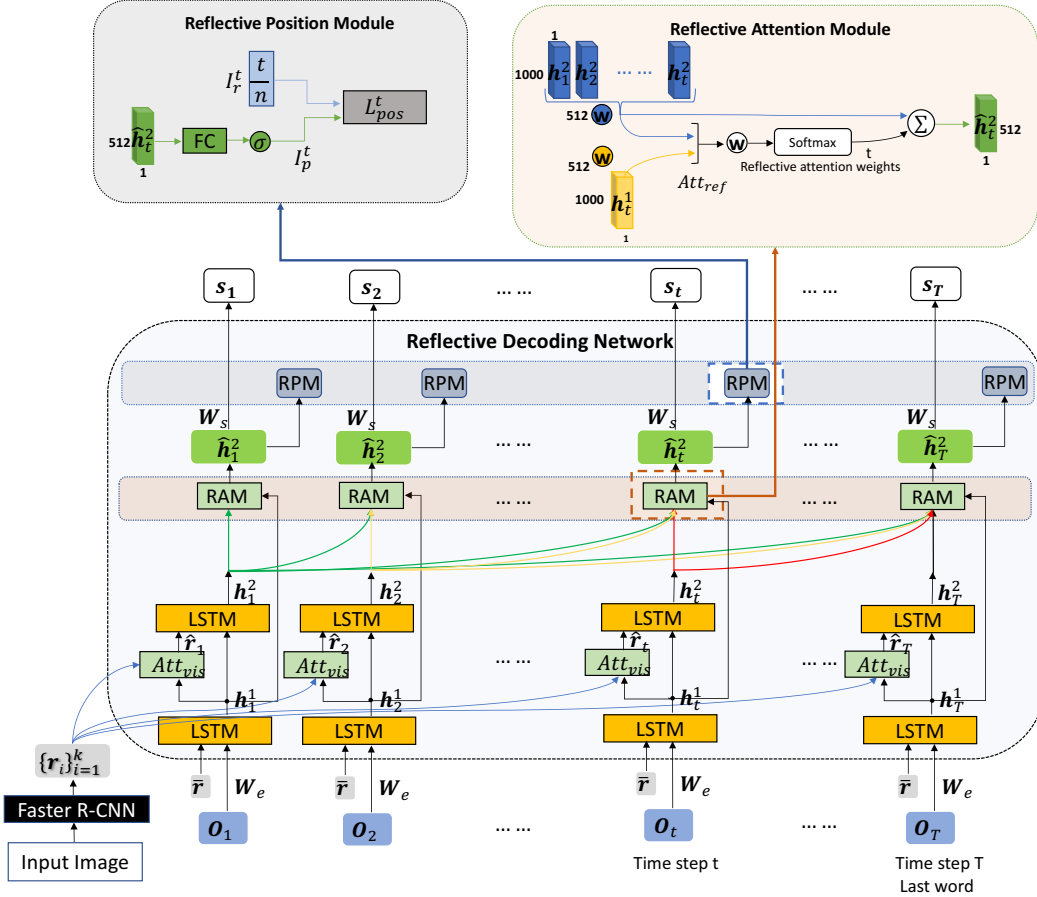


Figure 2. Overview of our framework.  $Att_{ref}$  in RAM (Reflective Attention Module) is the attention layer used to selectively attend to the generated decoding hidden states,  $Att_{vis}$  layer in Attention-based Recurrent Module determines the attention distribution over the detected image regions.  $I_p^t$  and  $I_r^t$  in RPM (Reflective Position Module) are respectively the  $t$ -th predicted and actual relative position in sentence.

ture of sentences. RDN is able to tackle the long-term dependency difficulty in caption decoding.

### 3.1. Object-Level Encoder

The encoder in encoder-decoder framework aims to extract meaningful semantic representation from an input image. We leverage object detection module (Faster R-CNN [33]) with pretrained ResNet-101 [13] to produce the region-level representation. The set of extracted regional visual representation  $R_I$  of an image  $I$  are denoted as  $R_I = \{\mathbf{r}_i\}_{i=1}^k$ ,  $\mathbf{r}_i \in \mathbb{R}^D$ , where  $k$  denotes the number of extracted regions,  $D$  denotes the feature dimension of each region, and  $\mathbf{r}_i$  is the mean pooled convolutional feature within the extracted region. Compared to the conventional uniform meshing method on CNN features, the object-level encoder focuses more on salient objects/regions in an image that is closely related to the perception mechanism in human visual system [3].

### 3.2. Reflective Decoder

Given a set of regional image features  $R_I$  produced by encoder, the goal for the decoder is to generate the caption

$S$ , where  $S = \{s_1, s_2, \dots, s_n\}$  consisting of  $n$  words. The generated caption should not only capture the content information from the image but also be meaningful and coherent. Specifically, in Figure 2, the Attention-based Recurrent Module is employed to selectively attend to the detected regional features and serves the basic function of a captioning decoder while Reflective Attention Module and Reflective Position Module are designed above it as assistants to further enhance captioning quality. Thus, the complete Reflective Decoder is able to take both historical coherence between words and syntactic structure information into consideration while generating image captions.

Attention-based Recurrent Module includes the first LSTM layer and visual attention layer  $Att_{vis}$ , which is designed mainly for top-down visual attention calculation. Its input  $\mathbf{x}_t^1$  at time step  $t$  contains three concatenated parts, the mean-pooled image feature  $\bar{\mathbf{r}} = \frac{1}{k} \sum_{i=1}^k \mathbf{r}_i$ , the embedding vector  $\mathbf{W}_e \mathbf{O}_t$  for current input word  $\mathbf{O}_t$  and the previous output  $\mathbf{h}_{t-1}^2$  from the second LSTM layer, where  $\bar{\mathbf{r}}$  represents the contextual information of the given image,  $\mathbf{W}_e \in \mathbb{R}^{E \times D_o}$  is the embedding matrix for the one-hot vec-

tor  $\mathbf{O}_t$ ,  $D_o$  is the size of the captioning vocabulary and  $E$  is the embedding size. The formula for updating the LSTM units in the first layer is defined as :

$$\mathbf{h}_t^1 = \text{LSTM}(\mathbf{x}_t^1, \mathbf{h}_{t-1}^1), \quad \mathbf{x}_t^1 = [\bar{\mathbf{r}}, \mathbf{W}_e \mathbf{O}_t, \mathbf{h}_{t-1}^2]. \quad (1)$$

For the visual attention layer  $Att_{vis}$ , given the generated  $\mathbf{h}_t^1$  and the set of  $k$  image features  $R_I = \{\mathbf{r}_i\}_{i=1}^k$ , we calculate the normalized attention weight  $\alpha_t^{vis}$  distribution over all the proposed object-level region denotes as :

$$\alpha_{i,t}^{vis} = \mathbf{W}_v^1 \tanh(\mathbf{W}_{rv}^1 \mathbf{r}_i + \mathbf{W}_{hv}^1 \mathbf{h}_t^1), \quad (2)$$

$$\alpha_t^{vis} = \text{softmax}(a_t^{vis}), \quad a_t^{vis} = \{\alpha_{i,t}^{vis}\}_{i=1}^k, \quad (3)$$

where  $\mathbf{W}_v^1 \in \mathbb{R}^{1 \times D_v}$ ,  $\mathbf{W}_{rv}^1 \in \mathbb{R}^{D_v \times D_R}$ ,  $\mathbf{W}_{hv}^1 \in \mathbb{R}^{D_v \times D_h}$  are learned embedding matrices,  $\alpha_t^{vis}$  denotes the calculated attention probability for each regional feature  $\mathbf{r}_i$  at time step  $t$ . So the attended feature is the weighted combination of each subregion,  $\hat{\mathbf{r}}_t = \sum_{i=1}^k \alpha_{i,t}^{vis} \mathbf{r}_i$  based on the weight distribution parameter  $a_t^{vis}$ .

### 3.2.1 Reflective Attention Module.

The Reflective Attention Module contains reflective attention layer  $Att_{ref}$ , combined with the second layer of LSTM, which is designed to output language description. Its input vector is concatenated by the attended feature result  $\hat{\mathbf{r}}_t$  and the hidden state  $\mathbf{h}_t^1$ . Thus the formula for updating the LSTM units in the second layer of LSTM is denoted as :

$$\mathbf{h}_t^2 = \text{LSTM}(\mathbf{x}_t^2, \mathbf{h}_{t-1}^2), \quad \mathbf{x}_t^2 = [\hat{\mathbf{r}}_t, \mathbf{h}_t^1]. \quad (4)$$

Based on the current hidden state  $\mathbf{h}_t^2$  at the time step  $t$  and the past hidden states set  $\{\mathbf{h}_1^2, \mathbf{h}_2^2, \dots, \mathbf{h}_{t-1}^2\}$ , the reflective attention layer  $Att_{ref}$  calculates the normalized weight distribution  $\alpha_t^{ref}$  above all the generated  $t$  hidden states as shown in the top right of Figure 2. The formula is defined as :

$$\alpha_{i,t}^{ref} = \mathbf{W}_h^2 \tanh(\mathbf{W}_{h_2h}^2 \mathbf{h}_i^2 + \mathbf{W}_{h_1h}^2 \mathbf{h}_t^1), \quad (5)$$

$$\alpha_t^{ref} = \text{softmax}(a_t^{ref}), \quad a_t^{ref} = \{\alpha_{i,t}^{ref}\}_{i=1}^t, \quad (6)$$

where  $\mathbf{W}_h^2 \in \mathbb{R}^{1 \times D_f}$ ,  $\mathbf{W}_{h_2h}^2 \in \mathbb{R}^{D_f \times D_h}$ ,  $\mathbf{W}_{h_1h}^2 \in \mathbb{R}^{D_f \times D_h}$  are three trainable matrices parameters,  $\alpha_t^{ref}$  denotes the generated attention probability set for each hidden state  $\mathbf{h}_i$  in the set  $\{\mathbf{h}_i^2\}_{i=1}^t$  at time step  $t$  and  $\alpha_{i,t}^{ref}$  reflects the relevance between the past predicted word at  $i$ -th step and current prediction (at  $t$ -th step) by measuring the compatibility between their corresponding hidden states. So we can calculate the attended hidden state result  $\hat{\mathbf{h}}_t^2 = \sum_{i=1}^t \alpha_{i,t}^{ref} \mathbf{h}_i^2$ .

The reflective decoding output  $\hat{\mathbf{h}}_t^2$  of the top attention layer  $Att_{ref}$  is utilized to predict the word  $s_t$  under the conditional probability distribution :

$$p(s_t | s_{1:t-1}) = \text{softmax}(\mathbf{W}_s \hat{\mathbf{h}}_t^2 + \mathbf{b}_s), \quad (7)$$

where  $\mathbf{W}_s \in \mathbb{R}^{D_o \times D_h}$  are the trainable weights and  $\mathbf{b}_s \in \mathbb{R}^{D_o}$  are the biases. By calculating  $s_t$  in this way, all the generated hidden states  $\{\mathbf{h}_i\}_{i=1}^t$  play a role in word precision and their extent of contributions can be clearly visualized, as will be demonstrated in section 4.3.2.

It should be noted that our proposed Reflective Attention Module models the dependencies between pairs of words at different time steps explicitly, taking into account the corresponding hidden states. In contrast, LSTM memorizes the historical sequence information by balancing the overall relevance of all time steps instead of modeling the dependency for each pair of words specifically.

### 3.2.2 Reflective Position Module.

It is often the case that many of the words have relatively fixed positions in a sentence due to the syntactic structure in natural language. For example, the numeral and subject words, *i.e.* ‘a man’ or ‘a woman’, mostly appear at the beginning of the sentence while the predicates tend to occupy the middle position. So we propose the Reflective Position Module by injecting the word position information during training as a guidance for the sequence decoding model to perceive its relative position or progress in the whole sentence. When decoding the  $t$ -th word, its actual relative position  $I_r^t$  and the predicted relative position  $I_p^t$  are calculated as :

$$I_r^t = \frac{t}{n}, \quad I_p^t = \sigma(\mathbf{W}_l \hat{\mathbf{h}}_t^2), \quad (8)$$

where  $n$  is the length of the sentence,  $\sigma$  is the sigmoid function and  $\mathbf{W}_l \in \mathbb{R}^{1 \times D_h}$  is the trainable relative position embedding matrix, respectively. The reflective position module shown in top left of Figure 2 aims to minimize the difference between  $I_r^t$  and  $I_p^t$  by designing a loss function, which refines the attended hidden state  $\hat{\mathbf{h}}_t^2$  mentioned in 3.2.1 by enabling it to perceive more sequential information of its relative position.

It is different from the popular position embedding methods [11, 40], which add the absolute position embedding to the corresponding input features in each dimension. Our Reflective Position Module models the relative position information individually in a supervised way. A key benefit of this design is that it can avoid the potential inter-pollution between the regular input feature and the position embedding, and equip our model with a strong perception of relative position for each word in the caption. Thus, the syntactic structure in natural language can be well preserved.

## 3.3. Training and Inference

**Training.** Two kinds of losses are utilized for optimizing our RDN model. The first is the cross entropy loss in traditional captioning training, which is to minimize the negative log likelihood:

$$L_{xe} = -\log p(S^* | I) = -\sum_{t=2}^n \log p(s_t^* | s_{1:t-1}^*), \quad (9)$$

where  $I$  is the given image,  $S^*$  is the ground truth caption, formula for calculating  $p(s_t^* | s_{1:t-1}^*)$  is defined in equation 7 and  $s_0^*$  is the start of the sentence.

The second loss is defined as the Position-Perceptive Loss  $L_{pos}$ :

$$L_{pos} = \sum_{t=1}^n \|I_r^t - I_p^t\|^2, \quad (10)$$

where  $I_r^t$  and  $I_p^t$  are the actual relative position and predicted relative position defined in Equation 8 and  $L_{pos}$  is designed to minimize the gap between them.

The objective function for optimizing our RDN is defined as :

$$L = L_{xe} + \lambda L_{pos}. \quad (11)$$

The trade-off parameter  $\lambda$  balances the contribution between the traditional caption loss in encoder-decoder framework and the Position-Perceptive Loss.

**Inference.** During the inference stage, since the length of the whole predicted sentence is unknown, the relative position information is removed from the input. As the discrepancy problem [22] between training and inference, which means the previous ground truth captioning token is not available for inference, we use the previously predicted word as input instead of ground truth word as in [5, 45]. This method is called teacher forcing algorithm [43]. Also, we adopt the popular beam search strategy which iteratively selects the top-k best sentences at time step  $t$  as candidates to generate the new top-k sentence at time  $t + 1$  in our experiment instead of greedy search.

## 4. Experiments

### 4.1. Datasets and Experimental Settings

**COCO Dataset.** COCO captions dataset [4] contains 82,783 images for training and 40,504 images for validation. Each image has five corresponding human-annotated captions. Also, we adopt the ‘Karpathy’ splits setting [18], which includes 113,287 training images, 5K validation images and 5K testing images for offline evaluation. For the online server evaluation, the entire images and captions in dataset is used for training. Following the text preprocessing in [2], we convert all the captions to lower case and remove the less frequent words which occur less than 5 times, obtaining a captioning vocabulary of 10,010 words.

**Visual Genome Dataset.** Visual Genome [19] is a large dataset for modeling the interactions and relations between objects within an image. The dataset consists of 108K images with densely annotated objects, attributes and pairwise relations. Compared to [48], we only utilize the annotated object and attribute data from the dataset to pretrain the object-level encoder and discard the pairwise relation data. We follow the same data split setting in [2] to include 98K

images for training, 5K images respectively for validation and testing. After cleaning these annotated object and attribute strings, we obtain a dataset including 400 attributes and 1,600 objects classes to train our Faster R-CNN model.

**Evaluation Metrics.** To objectively evaluate the performance of our captioning model, we use five widely accepted automatic evaluation metrics, including CIDEr [41], SPICE [1], BLEU [31], METEOR [8] and ROUGE-L [24].

**Implementation Details.** We implement our RDN using Caffe [16]. To train the object-level encoder, we use the Faster R-CNN with ResNet-101 pre-trained for image classification on ImageNet [35] and further refine it on the Visual Genome dataset. For each image, we set the IoU thresholds for region proposal suppression and object prediction to 0.7 and 0.3 respectively. For the remaining image subregions, we set a filter threshold 0.2. We rank the left-over boxes by their confidence scores from high to low and choose no more than top 100 as the final feature representations. Each region with dimension number 2,048 is the global average pooling result of the layer Res5c.

We set the word embedding size and the hidden size in each LSTM layer to 1,000. The dimensions for attention layers  $Att_{vis}$  and  $Att_{ref}$  are set to 512 respectively. During training, the initial learning rate is set to 0.01 and the polynomial decay strategy is adopted to decline the effective learning rate to zero by 70k iterations using a batch size 100. We tune the trade-off parameter  $\lambda$  on the ‘Karpathy’ validation split to obtain the best performance and finally set it to 0.02. For data augmentation, we adopt it only during the online test server submission to boost performance by flipping the original image and randomly cropping 90%. During decoding process, the beam search size is set to 5.

### 4.2. Ablation Study on Reflective Modules

To study the effects of Reflective Attention Module and Reflective Position Module in our model, an ablation experiment is designed to compare the performance with following combinations: (1) Baseline: the baseline denotes the RDN without Reflective Attention Module and Reflective Position Module; (2)  $RDN_{pos}$ : the RDN with the Reflective Attention Module removed, with only position module reserved, the number of attention layers in decoder is reduced to one; (3)  $RDN_{ref}$ : the RDN with the Reflective Position Module removed, cutting down the relative position information input; (4) RDN: the complete RDN implementation.

In Table 1, it can be observed that both the Reflective Position Module and Reflective Attention Module are important for our model and RDN improves the caption performance over all the metrics compared to baseline. The fact that  $RDN_{pos}$  outperforms baseline model validates the contribution of Reflective Position Module to enhance the quality of decoding hidden states during caption generation. Also, by injecting the Reflective Attention Module,

Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
Baseline	77.0	61.3	47.2	36.1	26.8	56.1	113.2	20.1
RDN <sub>pos</sub>	77.4	61.6	47.5	36.3	27.0	56.5	114.3	20.4
RDN <sub>ref</sub>	<b>77.6</b>	61.6	47.4	36.3	27.1	56.7	115.0	<b>20.5</b>
RDN	77.5	<b>61.8</b>	<b>47.9</b>	<b>36.8</b>	<b>27.2</b>	<b>56.8</b>	<b>115.3</b>	<b>20.5</b>

Table 1. Ablation study on COCO ‘Karpathy’ test split on single model. Our experiments show the contribution for reflective attention and position module, respectively. Results are obtained with beam size 5 without CIDEr optimization. All value reported in percentage (%).

Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
Review Net [47]	-	-	-	29.0	23.7	-	88.6	-
LSTM-A3 [49]	73.5	56.6	42.9	32.4	25.5	53.9	99.8	18.5
Att2in [34]	-	-	-	31.3	26.0	54.3	101.3	-
Adaptive [26]	74.2	58.0	43.9	33.2	26.6	-	108.5	-
Up-Down [2]	77.2	-	-	36.2	27.0	56.4	113.5	20.3
RFNet [17]	76.4	60.4	46.6	35.8	<b>27.4</b>	56.5	112.5	<b>20.5</b>
RDN	<b>77.5</b>	<b>61.8</b>	<b>47.9</b>	<b>36.8</b>	27.2	<b>56.8</b>	<b>115.3</b>	<b>20.5</b>

Table 2. Performance comparison on MSCOCO ‘Karpathy’ test split on single model. All image captioning models trained without optimizing CIDEr metric. (-) indicates the metric is not provided.

Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr	SPICE
NIC [42]	-	-	-	32.1	25.7	-	99.8	-
Att2in [34]	-	-	-	32.8	26.7	55.1	106.5	-
Review Net [47]	76.7	60.9	47.3	36.6	27.4	56.8	113.4	20.3
RFNet [17]	77.4	61.6	47.9	37.0	<b>27.9</b>	57.3	116.3	<b>20.8</b>
RDN	<b>77.6</b>	<b>62.2</b>	<b>48.6</b>	<b>37.8</b>	27.5	<b>57.4</b>	<b>117.3</b>	20.6

Table 3. Performance comparison on MSCOCO ‘Karpathy’ test split on ensemble models trained with cross entropy loss. Our model is the ensembling result of 6 single models initialized with different random seeds.

RDN<sub>ref</sub> performs obviously better than baseline, which shows the importance of model’s ability to capture long-term dependency between words. In particular, with a suitable combination of the two modules, RDN achieves the best result with CIDEr score 115.7, BLEU-4 score 37.0 and BLEU-3 score 47.9, yielding the improvement over our baseline model by 2.0%, 2.2% and 1.5% respectively, which is a considerable advancement over the benchmark. Compared the baseline model with total 1.15B parameters, RDN has only 0.84% more in model size, which is neglectable.

### 4.3. Performance Comparison and Analysis

We compare our proposed RDN with other state-of-the-art image captioning methods considering different aspects both in offline and online situation. Latest and representative works include: (1) **Adaptive** [26] which proposes the adaptive attention through designing a visual sentinel gate for captioning model to decide whether to attend to the image feature or just rely on the sequential language model, (2) **LSTM-A3** [49] which incorporates the high level semantic attribute information to the encoder-decoder model, (3) **Up-Down** [2] which introduces the bottom-up and top-down attention mechanism to enable attention calculated at the level of objects or salient subregions and (4) **RFNet** [17] which uses multiple kinds of CNNs to extract complementary image feature and generate a more informative repre-

sentation for the decoder.

For fair comparison, our model and the baseline use standard ResNet-101 as basic architecture for encoder and all the reported results on test portions of MSCOCO ‘Karpathy’ splits are trained without additional CIDEr optimization [34]. GCN-LSTM [48] is not included because it uses the additional densely annotated pair-wise relation data between objects to pretrain semantic relation detector and build convolutional graphs. We only adopt CIDEr optimization strategy for the online server submission, since directly optimizing the CIDEr metric has little effect on perceived caption quality during human evaluation [25] and small difference in its optimization implementation would influence the caption performance a lot.

#### 4.3.1 Quantitative Analysis

**Offline Evaluation.** For offline evaluation, we compare performance of different models on ‘Karparthy’ split dataset both in single and ensemble model situations. In Table 2, it can be observed that our single RDN achieves the best results among all existing captioning methods across the six evaluation metrics including all the BLEU entries, ROUGE-L and CIDEr, performs on par with RFNet in SPICE and is slightly inferior to it in METEOR. Different from the previous captioning models (Up-Down, RFNet, Review Net, Adaptive, etc.) that boost performance through

Model	BLEU-1		BLEU-4		METEOR		ROUGE-L		CIDEr		SPICE	
	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40
NIC [42]	71.3	89.5	30.9	58.7	25.4	34.6	53.0	68.2	94.3	94.6	18.2	63.6
Review Net [47]	72.0	90.0	31.3	59.7	25.6	34.7	53.3	68.6	96.5	96.9	18.5	64.9
LSTM-A3 [49]	78.7	93.7	35.6	65.2	27.0	35.4	56.4	70.5	116.0	118.0	-	-
Adaptive [26]	74.8	92.0	33.6	63.7	26.4	35.9	55.0	70.5	104.2	105.9	19.7	67.3
Att2all [34]	78.1	93.7	35.2	64.5	27.0	35.5	56.3	70.7	114.7	116.7	-	-
Up-Down [2]	80.2 <sub>2</sub>	95.2 <sub>2</sub>	36.9 <sub>3</sub>	68.5 <sub>3</sub>	27.6 <sub>3</sub>	36.7 <sub>3</sub>	57.1 <sub>3</sub>	72.4 <sub>3</sub>	117.9 <sub>3</sub>	120.5 <sub>3</sub>	-	-
RFNet [17]	80.4 <sub>1</sub>	95.0 <sub>3</sub>	38.0 <sub>1</sub>	69.2 <sub>2</sub>	28.2 <sub>1</sub>	37.2 <sub>2</sub>	58.2 <sub>1</sub>	73.1 <sub>2</sub>	122.9 <sub>1</sub>	125.1 <sub>2</sub>	-	-
RDN	80.2 <sub>2</sub>	95.3 <sub>1</sub>	37.3 <sub>2</sub>	69.5 <sub>1</sub>	28.1 <sub>2</sub>	37.8 <sub>1</sub>	57.4 <sub>2</sub>	73.3 <sub>1</sub>	121.2 <sub>2</sub>	125.2 <sub>1</sub>	-	-

Table 4. Performance comparison of published image captioning models on COCO Leaderboard. RDN achieves superior performance when comparing to other state-of-the-art methods. Top-3 rankings are indicated by red footnote for each metric.

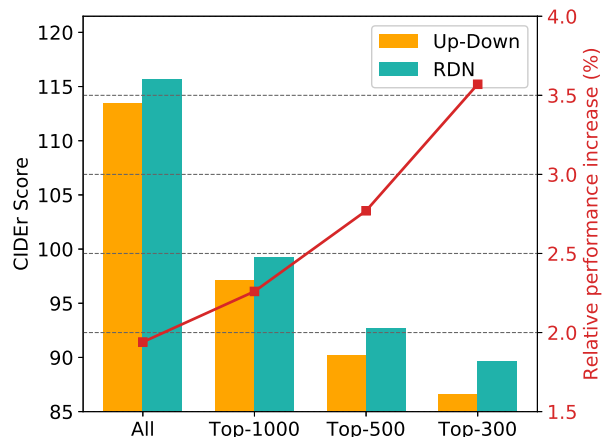


Figure 3. Performance comparison between our RDN model and Up-Down [2] on hard Image Captioning as a function of average length of annotations (ground truth captions). We rank the ‘Karpathy’ test set according to their average length of annotations in descending order and extract four different size of subsets. Smaller subset corresponds to averagely longer annotations and harder captioning. It reveals that our model exhibits more superiority over Up-Down in harder cases.

extracting more indicative and compact visual representation, the enhancement of our captioning model only attributes to a better reasoning and inference ability of the decoder directly on the target side. Moreover, from the Table 3, we can see that our ensemble RDN outperforms other ensemble models in most of evaluation metrics, with the highest CIDEr score 117.3, and performs only inferiorly to RFNet in METEOR and SPICE entry. RDN is the ensemble result of 6 single models with different random seed initialization while RFNet is composed of 4 RFNets with a total of 20 groups of different image representations.

**Online Evaluation on COCO Testing Server.** We also compare our model with the published state-of-the-art captioning systems on COCO Testing Server with 5(c5) and 40(c40) reference sentences as shown in Table 4. Using the ensemble of 9 CIDEr optimized models, our RDN achieves leading performance over all metrics while performing on par with RFNet [17]. Surprisingly, RFNet has a much better performance in online evaluation compared to the of-

line case, in which it performs much poorer than our model (even poorer than Up-Down [2] in some cases) shown in Table 2. Since the code of RFNet is not released, it is hard to investigate the inconsistency. Nevertheless, our RDN achieves the superior performance in all the c40 entries. Compared to c5, c40 has far more reference sentences and existing evaluation experiments show it achieves higher correlation with human judgement [4, 41]. Moreover, our model is more simple and elegant with only one encoder-decoder in single model compared to RFNet, which utilizes multiple encoders (ResNet, DenseNet [15], Inception-V3 [38], Inception-V4, and Inception-ResNet-V2 [37]) to extract 5 groups of features and includes time-consuming feature fusion steps to produce the final thought vectors. Besides, our RDN boosts captioning performance by optimizing the decoding stage while RFNet mainly focuses on improving the encoder. Thus, it is a promising extension to apply the encoding mechanisms of RFNet to our RDN. Compared to Up-Down [2], which uses traditional LSTM and object-level encoder, the CIDEr-c40, CIDEr-c5 and METEOR-c40 are improved by 3.9%, 2.7% and 2.9%.

**Evaluation on hard Image Captioning.** We further investigate the effect of the average length of annotations (ground truth captions) on the captioning performance, since generally the images with averagely longer annotations contain more complex scenes and thus are harder for captioning. Specifically, we rank the whole ‘Karpathy’ testset (5000 images) according to their average length of annotations in descending order and extract four different size of subsets (all set, top-1000, top-500, top-300 respectively). Smaller subset corresponds to averagely longer annotations and implies harder image captioning. Figure 3 shows the comparison between our RDN and Up-Down [2] (main difference of the two models is that Up-Down uses traditional LSTM). It reveals that the performance of both models are decreasing with the increasing average length of annotations, which reflects that the captioning is getting harder. However, our model exhibits more superiority over Up-Down in harder cases, which in turn validates the ability of our RDN to capture the long-term dependencies within captions. We also provide one such comparison between our RDN and Att2in [34] on hard captioning in the supplementary file.


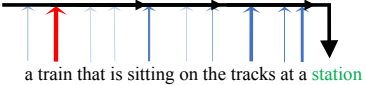

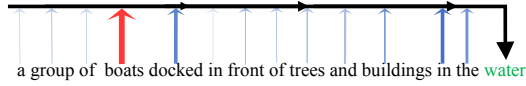

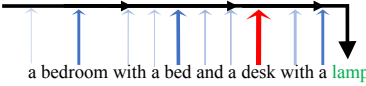

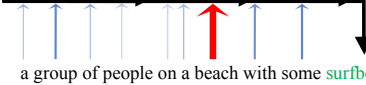
	Generated sentence	Reflective weight visualization
	<b>Basis decoder:</b> a train that is sitting on the tracks <b>Ours:</b> a train that is sitting on the tracks at a <b>station</b>	 a train that is sitting on the tracks at a <b>station</b>
	<b>Basis decoder:</b> a group of boats parked next to each other <b>Ours:</b> a group of boats docked in front of <b>trees and buildings</b> in the <b>water</b>	 a group of boats docked in front of trees and buildings in the <b>water</b>
	<b>Basis decoder:</b> a bedroom with a bed and a desk <b>Ours:</b> a bedroom with a bed and a desk <b>with a lamp</b>	 a bedroom with a bed and a desk with a <b>lamp</b>
	<b>Basis decoder:</b> a group of people are standing in the water <b>Ours:</b> a group of people on a beach <b>with some surfboards</b>	 a group of people on a beach with some <b>surfboards</b>

Figure 4. Examples of captions generated by our RDN compared to the basis decoder (using traditional LSTM) and their reflective attention weight distribution over the past generated hidden states when predicting the key words highlighted in green. The thicker line indicates a relatively larger weight and the red line means the largest contribution. More examples are provided in the supplementary material.

### 4.3.2 Qualitative Analysis

To investigate the physical interpretability of RDN model’s improvement, some qualitative comparisons of captioning results are shown in Figure 4. Compared to basis decoder, our RDN is able to generate more detailed and discriminative descriptions for images. Take the first case in Figure 4 as an example, the basis decoder can provide a general and reasonable caption for the image. However, it cannot recognize the word ‘station’ which actually exists in the image while our model successfully infers it based on the previously generated words, especially the closely related words ‘train’, ‘tracks’ and ‘sitting’. For the reflective weight visualization, the generated words with the largest contribution to the predicting word are usually strong related in vocabulary coherence, such as the correlations between words “boat” and “water”, “beach” and “surfboards”. We show additional results in supplementary material.

In Figure 5, compared to other captioning models, our RDN is able to predict the word and its relative position in the sentence simultaneously during caption generation. The predicted relative position in blue for each word is highly close to its actual relative position value in sentence, which demonstrates a good position-perceptive ability of our model to capture the syntactic structure of a sentence.

## 5. Conclusion

We have presented a novel architecture, Reflective Decoding Network (RDN), which explicitly explores the co-

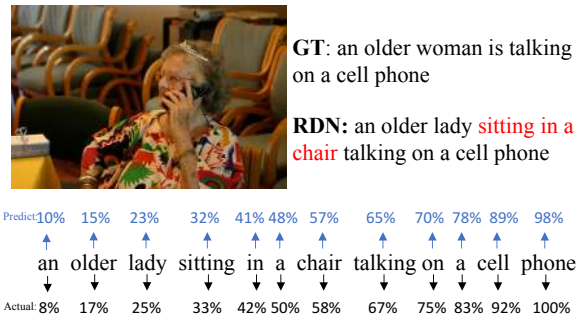


Figure 5. **Top:** Ground-truth caption and the example caption generated by our RDN model. **Bottom:** The predicted relative position value (shown in blue) from the Reflective Position Module and the actual relative position for each word in the sentence. All value reported in integer percentage value (%).

herence between words in the captioning sentence and enhances the long-term sequence inference ability of LSTM. Particularly, the attention mechanism applied in both visual and textual domain and the proposed position-perceptive scheme are to maximize the reference information available for captioning model. We also show how the learned attention in textual domain can provide interpretability during the captioning generation process from a new perspective. Extensive experiments conducted on standard and hard COCO image captioning dataset with superior performance validate the effectiveness of our proposal. For future work, we are interested in extending our model to source code captioning and text summarization.



## References

- [1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *ECCV*, 2016. 5
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and vqa. In *CVPR*, 2018. 1, 2, 5, 6, 7
- [3] Timothy J Buschman and Earl K Miller. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *science*, 2007. 3
- [4] 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. 5, 7
- [5] Xinlei Chen and C Lawrence Zitnick. Mind’s eye: A recurrent visual representation for image caption generation. In *CVPR*, 2015. 5
- [6] Xinpeng Chen, Lin Ma, Wenhao Jiang, Jian Yao, and Wei Liu. Regularizing rnns for caption generation by reconstructing the past with the present. In *CVPR*, 2018. 2
- [7] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 2
- [8] Michael Denkowski and Alon Lavie. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the ninth workshop on statistical machine translation*, 2014. 5
- [9] 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. *arXiv preprint arXiv:1505.01809*, 2015. 2
- [10] 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. 2
- [11] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*, 2017. 4
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *ECCV*, 2014. 1
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997. 1
- [15] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. 7
- [16] 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 *Proceedings of the 22nd ACM international conference on Multimedia*, 2014. 5
- [17] Wenhao Jiang, Lin Ma, Yu-Gang Jiang, Wei Liu, and Tong Zhang. Recurrent fusion network for image captioning. In *ECCV*, 2018. 1, 2, 6, 7
- [18] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. 5
- [19] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 5
- [20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [21] 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. *PAMI*, 2013. 2
- [22] Alex M Lamb, Anirudh Goyal ALIAS PARTH GOYAL, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. Professor forcing: A new algorithm for training recurrent networks. In *NIPS*, 2016. 5
- [23] Li-Jia Li, Richard Socher, and Li Fei-Fei. Towards total scene understanding: Classification, annotation and segmentation in an automatic framework. In *CVPR*, 2009. 1
- [24] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*, 2004. 5
- [25] Siqi Liu, Zhenhai Zhu, Ning Ye, Sergio Guadarrama, and Kevin Murphy. Optimization of image description metrics using policy gradient methods. *CoRR, abs/1612.00370*, 2016. 6
- [26] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *CVPR*, 2017. 1, 2, 6, 7
- [27] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *NIPS*, 2016. 2
- [28] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural baby talk. In *CVPR*, 2018. 2
- [29] Hongyuan Mei, Mohit Bansal, and Matthew R Walter. Coherent dialogue with attention-based language models. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017. 2
- [30] Lesly Miculicich Werlen, Nikolaos Pappas, Dhananjay Ram, and Andrei Popescu-Belis. Self-attentive residual decoder for neural machine translation. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018. 2
- [31] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002. 5
- [32] Romain Paulus, Caiming Xiong, and Richard Socher. A deep reinforced model for abstractive summarization. *arXiv:1705.04304*, 2017. 2
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 1, 2, 3

- [34] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017. 6, 7
- [35] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 5
- [36] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 2014. 2
- [37] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, 2017. 7
- [38] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 7
- [39] Ke Tran, Arianna Bisazza, and Christof Monz. Recurrent memory networks for language modeling. *arXiv preprint arXiv:1601.01272*, 2016. 2
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 2, 4
- [41] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015. 5, 7
- [42] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and Tell: A Neural Image Caption Generator. In *CVPR*, 2015. 2, 6, 7
- [43] Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1989. 5
- [44] 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. 2
- [45] Kelvin Xu, Jimmy Lei Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. *ICML*, 2015. 1, 2, 5
- [46] Yezhou Yang, Ching Lik Teo, Hal Daumé III, and Yiannis Aloimonos. Corpus-guided sentence generation of natural images. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2011. 2
- [47] Zhilin Yang, Ye Yuan, Yuexin Wu, William W Cohen, and Ruslan R Salakhutdinov. Review networks for caption generation. In *NIPS*, 2016. 1, 2, 6, 7
- [48] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *ECCV*, 2018. 1, 2, 5, 6
- [49] Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, and Tao Mei. Boosting image captioning with attributes. In *ICCV*, 2017. 2, 6, 7
- [50] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *CVPR*, 2016. 1, 2
- [51] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. Mattnet: Modular attention network for referring expression comprehension. In *CVPR*, 2018. 2