



# Boost image captioning with knowledge reasoning

Feicheng Huang<sup>1</sup> · Zhixin Li<sup>1</sup> · Haiyang Wei<sup>1</sup> · Canlong Zhang<sup>1</sup> · Huifang Ma<sup>2</sup>

Received: 15 April 2020 / Revised: 21 July 2020 / Accepted: 19 September 2020 /  
Published online: 27 October 2020

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2020

## Abstract

Automatically generating a human-like description for a given image is a potential research in artificial intelligence, which has attracted a great of attention recently. Most of the existing attention methods explore the mapping relationships between words in sentence and regions in image, such unpredictable matching manner sometimes causes inharmonious alignments that may reduce the quality of generated captions. In this paper, we make our efforts to reason about more accurate and meaningful captions. We first propose word attention to improve the correctness of visual attention when generating sequential descriptions word-by-word. The special word attention emphasizes on word importance when focusing on different regions of the input image, and makes full use of the internal annotation knowledge to assist the calculation of visual attention. Then, in order to reveal those incomprehensible intentions that cannot be expressed straightforwardly by machines, we introduce a new strategy to inject external knowledge extracted from knowledge graph into the encoder-decoder framework to facilitate meaningful captioning. Finally, we validate our model on two freely available captioning benchmarks: Microsoft COCO dataset and Flickr30k dataset. The results demonstrate that our approach achieves state-of-the-art performance and outperforms many of the existing approaches.

**Keywords** Image captioning · Word attention · Visual attention · Knowledge graph · Reinforcement learning

## 1 Introduction

Image captioning has recently attracted great attention in the field of artificial intelligence, due to the significant progress of machine learning technologies and the release of a number of large-scale datasets (Hossain et al. 2019; Bai and An 2018; Chen et al. 2017c). The gist of the caption task is to generate a meaningful and natural sentence that describes the

---

Editors: Kee-Eung Kim, Vineeth N Balasubramanian.

---

✉ Zhixin Li  
lizx@gxnu.edu.cn

<sup>1</sup> Guangxi Key Lab of Multi-source Information Mining and Security, Guangxi Normal University, Guilin 541004, China


<sup>2</sup> College of Computer Science and Engineering, Northwest Normal University, Lanzhou 730070, China

most salient objects and their interactions for the given image. Solving this problem has great impact on human community, as it can help visual impaired people understand various scenes and be treated as an auxiliary means of early childhood education (Jiang et al. 2018a, b). Despite its widely practical applications, image captioning has long been viewed as a challenging research, mainly because it needs to explore a suitable alignment between two different modalities: image and text.

The popular image captioning approaches adopt the encoder-decoder framework (Vinyals et al. 2015; Jia et al. 2015; Wu and Cohen 2016; Mathews et al. 2016; Ramanishka et al. 2017). In general, a Convolutional Neural Network (CNN) is often used as the encoder to present the image with a fixed-length representation, while a Recurrent Neural Network (RNN) or Long Short-Term Memory (LSTM) Network is employed to decode this representation into a caption. Attention mechanism has indeed demonstrated significant effectiveness on the task of image captioning (Xu et al. 2015; You et al. 2016; Anderson et al. 2018; Lu et al. 2017; Wang et al. (2017)). This mechanism allows models to attend on image regions relevant to each generated word at every time step, rather than only using the whole image to guide the generation of captions. Although promising results have achieved, it's obvious that the current captioning systems are limited in two constraints:

First, since visual attention in captioning models can be viewed as the mappings from image regions to sentence snippets, however, this mapping procedure usually performs compulsively and unpredictably in a “black box”, and thus ignoring the fact that some words are not related to any entity in the image. As the result, it may cause inharmonious alignments between image regions and sentence snippets that will reduce the quality of generated sentences.

Second, most of the captioning models are built on a large number of paired image-caption data, but each image in the training data only contains several ground-truth captions, which will lack of sufficient cues to reveal the incomprehensible intentions that are not explicitly presented in the image. As shown in Fig. 1. Furthermore, in order to extend the ability to describe new entities out of the training data, more knowledge need to be introduced from external data sources.



**Fig. 1** The ground-truth caption simply describes the low-level content of the image, and doesn't explain why this woman is standing there. By incorporating external knowledge, we speculated that she might be waiting for the bus. In the sentence, the words “woman” and “luggage” are more important than others, as they describe the main aspects of the image

In this paper, we mainly focus on alleviating the aforementioned two constraints to generate more accurate and meaningful captions. As we all know, not all words in the caption have equivalent importance in describing the image (Chen et al. 2017b). Seen from Fig. 1. We capture this perception, and devise a word attention to modulate the alignments between words in sentence and regions in image. Specifically, a score is assigned to each input word based on its significance in describing the image, then we compute the word context vector at each time step to make better use of the bottom-up semantic information to boost the process of visual attention. This makes our model suit the human perception that some salient regions in the image are more likely to be described than non-salient ones. At the same time, we also leverage commonsense knowledge extracted from knowledge graph to achieve better generalization. In stead of fusing the input sentence and external knowledge together to train an RNN, we input the knowledge into the word generation stage to augment the probabilities of some potential words that are likely to be applied to describe the given image. This enables our model to generate more meaningful sentences than other existing models. To sum up, our contributions are shown as follows:

- We propose a new text-dependent attention mechanism called word attention to assist the generation of visual attention, thereby making our model generate more accurate captions.
- We introduce a new strategy to incorporate knowledge graph into our encoder-decoder framework to take better use of external knowledge to facilitate novel and meaningful captioning.
- By combining the aforementioned two proposed scenarios, experiments conducted on MSCOCO and Flickr30k benchmarks show that our approach achieves state-of-the-art performance and outperforms many of the existing approaches.

The rest of the paper is organized as: In Sect. 2, we summarize the existing captioning models into several categories, and review the previous works that are related to this paper. Then, we present the implementation of our approach in Sect. 3. The experiment results and analysis are shown in Sect. 4. And in Sect. 5, we make a conclusion and give a brief prospect of future work.

## 2 Related work

The state-of-the-art image captioning solutions are neural network-based sequence learning methods, which use CNNs to encode an image into a fixed-length image representation, and then adopt RNNs to decode the representation into a meaningful sentence, thus the process of caption generation suits an end-to-end style. Since there is a growing body of captioning algorithms have achieved superior performance (Vinyals et al. 2015; Jia et al. 2015; Mathews et al. 2016), these approaches often suffer from the problems of object missing and misprediction, due to their only use image-level representation to initialize the hidden state of RNN or LSTM.

### 2.1 Attention mechanism based models

The problems of object missing and misprediction can be mitigated by introducing attention mechanism into the general captioning models. Motivated by human perception, and

encouraged by recent success in machine translation using attention mechanism, Xu et al. (2015) integrated visual attention into the encoder-decoder framework for image captioning, making their work a new state-of-the-art. Visual attention amounts to learning the latent alignments between words in sentence and regions in image when generating description word-by-word. Following this successful attempt, different attention methods are proposed. You et al. (2016) proposed a semantic attention model which learns to selectively attend to semantic attributes detected from the given image, thereby making better use of the top-down visual information and the bottom-up semantic information. Instead, Li et al. (2017b) combined the global features and local features through a Global-Local attention, where the local features are obtained by Faster R-CNN (Ren et al. 2015). Similarly, Anderson et al. (2018) also implemented their bottom-up attention using Faster R-CNN, and then a top-down attention is constructed to attend to salient regions of the image. Differently from the above methods, Liu et al. (2017a) and Lu et al. (2017) investigated the agreements between image regions and their corresponding words. The former defined a quantitative metric to evaluate the “correctness” of the attention maps generated by the uniform attention model and applied supervision to improve the attention correctness, while the latter proposed an adaptive attention and via a “visual sentinel” to decide when to attend to the visual signals and when to depend on language properties to predict the next word. In particular, Huang et al. (2019) devised an AoA (Attention on Attention) model to filter out the inappropriate attention results, and only used the useful attended information to guide the caption generation process. All of these works have demonstrated the effectiveness of attention mechanism on the image captioning task. In this paper, a new text-dependent word attention is added to the uniform visual attention model. Its calculation only depends on the internal annotation knowledge in the training data, which can provide rich semantic information to guide the generation of visual attention.

## 2.2 Incorporation of external knowledge

While promising advances are presented in exiting captioning methods, they lack the ability to describe novel objects or attributes outside of training corpora, and are unable to express implicit aspects of the image, as the knowledge acquired from ground-truth captions is not sufficient. Such issue can be solved by incorporating knowledge from external resources into the caption generation process. An early study presented in Anne Hendricks et al. (2016) exploited object knowledge from external object recognition datasets or text corpora to facilitate novel object captioning. Recently, Yao et al. (2017a) employed copying mechanism to directly copy novel objects that do not exist in the training corpora but are learnt from object recognition datasets to the output sentence generated by LSTM, thus making their proposed model obtain encouraging performance. Li et al. (2019b) further extended this work by using pointing mechanism to elegantly accommodate the influence between copying mechanism and word generation, in order to generate more accurate and natural sentence. In contrast to using raw external data sources, some studies attempt to incorporate structured knowledge to solve the specific problems. Li et al. (2017a) and Gu et al. (2019) employed knowledge graph for visual question answering and scene graph generation, respectively. These two studies both embedded the knowledge retrieved from external knowledge graph into a common space with other data, making their models flexible to adapt external test instances. Particularly, in Zhou et al. (2019), Zhou et al. proposed to leverage knowledge graph to boost image captioning, which is close to our proposed approach. But unlike (Zhou et al. 2019), which use a knowledge graph to extract indirectly

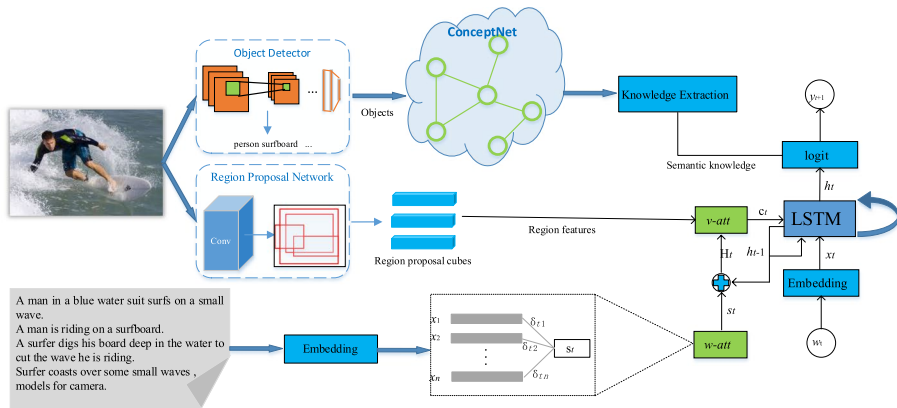
related terms and directly related terms about the entities detected by an object detector to pretrain an RNN, we inject semantically related information of the detected objects into the output stage of the caption generator to augment the probability of some latent meaningful words at each decoding time. This also allows our system to generate more novel and meaningful captions.

## 2.3 Enhanced by reinforcement learning

Moreover, several studies have been proposed to incorporate reinforcement learning technology to address the issues of non-differentiable evaluation metric and exposure bias Ranzato et al. (2015) in image captioning. Ren et al. (2017) treated image captioning as a decision-making task, and proposed a deep reinforcement learning based model to generate a natural description for an image. The model employs a “policy network” and a “value network” to collaboratively determine the next word at each intermediate step. In Rennie et al. (2017), a self-critical sequence training (SCST) approach is proposed to dramatically optimize the training process using the test metrics (especially, the CIDEr metric) at decoding stage. The SCST approach avoids having to estimate the reward signal and consider how to normalize the reward, and it exploits the generated captions in the inference phase as the “baseline” to encourage the generation of more accurate captions. Anderson et al. (2018) and Qin et al. (2019) used the similar manner as Rennie et al. (2017) to improve the performance of captioning models, but different in the way of sequence sampling. It is worth mentioning that the last two works can also be categorized as attention mechanism based method, since the former, as mentioned above, proposed a bottom-up and top-down (Up-Down) attention model, and the latter used this proposed model as the “backbone” to realize its look back and predict forward (LBPF) model. Later on, Yao et al. (2019) proposed to employ instance-level, region-level and image-level features of an image to build a hierarchical structure, thereby giving the caption generator a thorough image understanding. Their proposed architecture is pluggable to many advanced reinforcement learning models, and encouraging performance is achieved. In this work, we also follow these work, and use SCST approach to optimize our model.

## 3 Method

Like most of the captioning methods, we attempt to seed for a suitable and human-like description for a given image. The overview of the proposed image captioning architecture is illustrated in Fig. 2. Compared to previous models, our model consists of two novel ideas. On the one hand, a special word attention is designed to handle the inharmonious matching problem between regions in image and words in caption. On the other hand, we take commonsense knowledge extracted from external knowledge graph into account to facilitate the generation of novel and meaningful captions. Actually, our whole model seems to make full use of the internal annotation knowledge and external knowledge to guide the caption generation. But note that, the two proposed scenarios use knowledge in different ways. In the following, we introduce the implementation of our whole model including the extraction of image feature, the implementation of word attention, the incorporation of knowledge graph and the usage of reinforcement learning.



**Fig. 2** The overview of our captioning framework with word attention and knowledge graph. Specifically, we use objects detected by an object detector to retrieve semantic knowledge from the knowledge graph (here we use ConceptNet) to guide the generation of captions. Meanwhile, to extract more useful information from the given image, a region proposal network is incorporated to generate the region features. Our proposed word attention serves as another pipeline to inject compact textual information to assist the calculation of visual attention. Here, *w-att* and *v-att* represent word attention and visual attention, respectively

### 3.1 Image feature extraction and word embedding

To more effectively use the information in the image to guide the caption generation, we use region proposal features of the image to train our model. Specifically, following the method proposed by Ren et al. (2015), we apply region proposal network to generate a lot of rectangular region proposals. Afterwards, each proposal is feed to ROI pooling layer and 3 full-connected layers to obtain a vector representation  $v_i$  of each image region. Compared to the approach of Xu et al. (2015), which used the  $z = x \times y$  locations of the activation grid to form the image representation, region proposal features provide more useful information of the image content. Please note that, given the image feature  $V = \{v_1, v_2, \dots, v_L\}$ ,  $v_i \in \mathbb{R}^D$ , we use the mean pooling vector  $\bar{v}$  as the global image information, and the  $\bar{v}$  will be feed to initial the LSTM decoder to give an overall understanding of the image.

A common way of word embedding is one-hot encoding. This encoding style sets one element of a vector to 1 and others to 0 to represent a specific word in the dictionary. If there are too many words in the vocabulary, the one-hot vector will become sparse and the problem of dimension explosion will also occur. Besides, the one-hot encoding does not consider the order of words, which is unfavorable to the calculation of our word attention. Thus, in this work, we use a pre-trained word2vec model for word embedding. The word2vec is essentially a neural network, it uses the raw one-hot vectors as inputs to obtain the final embedding vectors. Therefore, such embedding manner fully considers the context between words and can provide more abundant semantic information.

### 3.2 Implementation of word attention

In this subsection, we introduce the implementation of the proposed word attention, and investigate in what manner the word attention contributes to the improvement of visual attention.

Actually, the motivation of word attention comes from the perception that some words are more related to the content of a given image than others. What we need to do is to strengthen this connection, so that these words can play a better guiding role in the training process. Consequently, the model can learn a more suitable mapping pattern between captions and images, which in turn improves the quality of generated captions.

Suppose an image  $I$  to be described by a sentence  $S = \{w_1, w_2, \dots, w_N\}$ , where  $N$  represents the length of the caption. The sequence learning methods usually use RNN or LSTM to generate each word at each time stage, in which LSTM has shown great performance. Following this trend, we add a word attention into the caption generator to form our captioning model. At the training phase, the word attention mainly depends on ground-truth captions, and the operations of the word attention are as follows:

$$\delta_{ii} = f_w(w_i) \quad (1)$$

$$\beta_{ii} = \frac{\exp(\delta_{ii})}{\sum_{k=1}^N \exp(\delta_{ik})} \quad (2)$$

$$s_t = \sum_{i=1}^N \beta_{ii} x_i \quad (3)$$

where  $f_w$  is a function that calculates the weight value allocated to  $w_i$ ,  $x_i$  is the embedding vector of  $w_i$ , and  $s_t$  represents the word context vector at time  $t$ . Note that the  $\delta_{ik}$  stays the same during the generation of each word until the last time step. Here, inspired by the previous works (Kim et al. 2018; Park et al. 2017), we use TF-IDF method as the function  $f_w$ , as this method can measure the importance degree of each word in a sentence or document. The word context vector  $s_t$  is then fused with the previous hidden state  $h_{t-1}$  of the LSTM decoder to combine more compact semantic information to guide the visual attention, calculated as follows:

$$H_t = s_t \odot h_{t-1} \quad (4)$$

$$e_{ii} = W_e^T \tanh(W_v v_i + W_h H_t) \quad (5)$$

$$\alpha_{ii} = \frac{\exp(e_{ii})}{\sum_{k=1}^L \exp(e_{ik})} \quad (6)$$

$$c_t = \sum_{i=1}^L \alpha_{ii} v_i \quad (7)$$

where  $\odot$  is the element-wise multiplication;  $W_e$ ,  $W_v$  and  $W_h$  are learned parameters;  $\tanh$  is the hyperbolic tangent function; the visual context vector  $c_t$  is the weighted sum of all the image region features. Combined with word attention, the decoder LSTM updates for time step  $t$  are:

$$i_t = \sigma(W_i x_t + U_i c_t + Z_i h_{t-1} + b_i) \quad (8)$$

$$f_t = \sigma(W_f x_t + U_f c_t + Z_f h_{t-1} + b_f) \quad (9)$$

$$o_t = \sigma(W_o x_t + U_o c_t + Z_o h_{t-1} + b_o) \quad (10)$$

$$g_t = \sigma(W_g x_t + U_g c_t + Z_g h_{t-1} + b_g) \quad (11)$$

$$m_t = f_t \odot m_{t-1} + i_t \odot g_t \quad (12)$$

$$h_t = o_t \odot \tanh(m_t) \quad (13)$$

$$p_{t+1} = w_{t+1}^T M_g h_t \quad (14)$$

here  $i_t$ ,  $f_t$ ,  $o_t$ ,  $g_t$ ,  $m_t$  and  $h_t$  are input gate, forget gate, output gate, cell gate, cell memory and hidden state of the LSTM, respectively;  $\sigma(\cdot)$  represents the sigmoid function;  $W_*$ ,  $U_*$ ,  $Z_*$ ,  $b_*$  are weight matrices and biases to be learned; Eq. (14) adopts the generation mechanism to predict the next word, where  $M_g$  is the transformation matrix.

Finally, our proposed word attention can be insert into the encoder-decoder framework in a trainable manner, thus making better use of the caption information annotated by humans. As a result, the quality of generated captions will be improved.

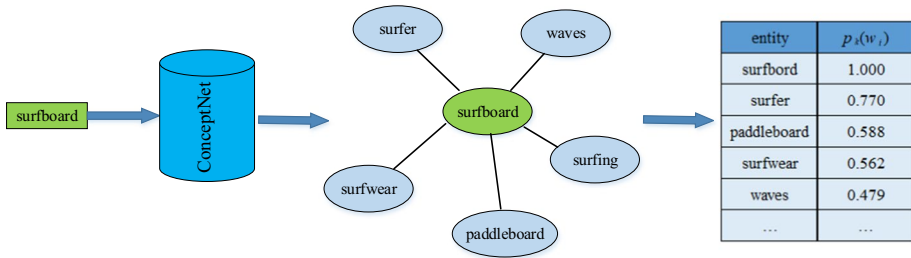
### 3.3 Incorporation of knowledge graph

On the task of image captioning, knowledge is of significantly important, as it provides a lot of cues for generating captions. The ground-truth annotations corresponding to each image in paired image-caption datasets are the knowledge provided by human beings for caption generation, which can be called internal knowledge. However, in many existing datasets, it is impossible to include all the knowledge required for captioning task, thereby limiting research progress. Therefore, we acquire knowledge from external resources to assist the caption generation, so as to improve the generalization performance of the captioning model. In recent years, many knowledge graph have appeared in the field of artificial intelligence. In this paper, we use ConceptNet (Speer et al. 2017), which is an open multilingual knowledge graph containing common sense knowledge closely related to human daily life, to help computers understand human intentions.

In general, each piece of knowledge in the knowledge graph can be view as a tripe (*subject*, *rel*, *object*), where *subject* and *object* represent two entities or concepts in the real world, and *rel* is the relationship between them. To obtain informative knowledge that are relevant to the given image, we first use Faster R-CNN (Ren et al. 2015) to detect a series of objects or visual concepts, and then use these objects or concepts to retrieve semantically similar knowledge from the knowledge graph. Figure 3 gives an illustration of using a detected word “*surfboard*” to retrieve semantic knowledge from the ConceptNet. As we have seen, each knowledge entity corresponds to a probability  $p_k$  that represents the degree of correlation. For each detected object or concept, we select the relevant knowledge entities for captioning task. Such we get a small semantic knowledge corpus  $W_k$  containing the most relevant knowledge.

And the problem is that how to apply the important semantic knowledge extracted from the knowledge graph to guide the process of caption generation, which needs to be





**Fig. 3** The illustration of knowledge extraction using an object “surfboard”. For convenience, we simply give a part of the results. Note that each relevant semantic entity corresponds to a probability  $p_k$  that represents the degree of correlation

carefully considered. Since unnecessary inputs may cause noise in the training phase, thus reducing the performance of the model. Thereby, we do not directly input the semantic knowledge into the input layer of LSTM for training, instead we focus on changing the *logit* layer of the caption generation network, and we augment the probability of some potential words that are appear in the constructed semantic knowledge corpus  $W_k$  when we predict the next word. After using Back Propagation to train the whole model, a more robust system is obtained. For this purpose, we changed the Eq. (14) as:

$$p_{t+1} = \begin{cases} w_{t+1}^T M_g h_t + \lambda p_k(w_{t+1}), & w_{t+1} \in W_K \\ w_{t+1}^T M_g h_t, & \text{otherwise} \end{cases} \quad (15)$$

where  $\lambda$  is a hyper parameter that measures the degree of introducing external semantic knowledge. If the word  $w_{t+1}$  exists in the constructed semantic knowledge corpus  $W_k$ , the prediction probability of the word will be determined by the prediction probability of the generation mechanism and the corresponding retrieval probability  $p_k(w_{t+1})$ . In general, a softmax function is used after Eq. (15) to obtain a normalized word probability distribution. By adding an additional probability to each possible word, the model will be able to discover some implicit cues, thus leading to generate more novel and meaningful captions.

### 3.4 Sequence generation based on reinforcement learning

Here we discuss the caption generation process of our model. It is sure that in the training and inference phase, our proposed two scenarios can not only be appropriately combined together to guide the caption generation process, but each can perform solely to address the different deficiencies existing in the previous models. Our approach suits the popular sequence learning based methods. In other word, the sequence is generated word-by-word. The state-of-the-art captioning models are typically trained with cross entropy loss:

$$L(\theta) = - \sum_{t=1}^T \log(p_\theta(w_t^* | w_1^*, \dots, w_{t-1}^*)) \quad (16)$$

where  $p_\theta$  is a captioning model with all the parameters  $\theta$ , and the sequence  $(w_1^*, \dots, w_T^*)$  is the ground-truth caption. However, the cross entropy objective function dose not perform well with the problem of “exposure bias” (Ranzato et al. 2015). The issue can be mitigated by introducing “scheduled sampling” (Bengio et al. 2015) at the decoding stage. However, the “scheduled sampling” strategy seems to be statistically inconsistent. Another

alternative solution to “exposure bias” is reinforcement learning (Liu et al. 2016a, b). In this paper, we adopt SCST method to optimize our model. Note that, the caption generator (LSTM) can be viewed as an “agent”, and the caption words and image features serve as “environment”. In addition,  $p_\theta$  defines a “policy”, that will result in generating the next best word. In this case, we minimize the negative expected reward to train the model:

$$L_r(\theta) = -\mathbb{E}_{w_{1:T}p_\theta} [r(w_{1:T})] \quad (17)$$

where  $r(w_{1:T})$  is the score function, here we use CIDEr as Rennie et al. (2017) and Anderson et al. (2018). The gradient of this loss then will be approximated as:

$$\nabla_\theta L_r(\theta) \approx -(r(w_{1:T}^s) - r(w_{1:T}^m)) \nabla_\theta \log p_\theta(w_{1:T}^s) \quad (18)$$

where  $w_{1:T}^s$  and  $w_{1:T}^m$  denote the sampled sequence and the result of greedy decoding, respectively. In particular, we set the baseline as  $r(w_{1:T}^m)$ , this has made our model achieve significant gains in the performance.

The core idea of this reinforcement learning based training approach is to take the reward obtained by the inference algorithm used by the current model during testing as the baseline of the reinforce algorithm. This approach keeps the model consistent during training and inference, thereby significantly improving the quality of the generated captions. In the later, we demonstrate the effectiveness of our model combining with reinforcement learning based training manner.

## 4 Experiments

In this section, extensive experiments are conducted to evaluate the effectiveness of our proposed model. We first introduce the datasets and evaluation metrics. Then, the implementation details of our experiments are followed. Finally, we make a comparison with other state-of-the-art models and give brief analysis of the results.

### 4.1 Datasets and evaluation metrics

We mainly use the popular MSCOCO 2014 dataset to validate the performance of our proposed models. This large dataset contains 123,287 images, with at least 5 ground-truth sentences are annotated per image for image captioning. For fair comparison with other methods, we adopt the ‘Karpathy’ splits (Johnson et al. 2016) that have been widely used in previous work. Thus, we get 113,287 images for training, and 5,000 images for validation and testing respectively. Compared to MSCOCO 2014 dataset, flickr30k dataset is smaller, and it contains 31,000 images. Since it does not provide official split, we also follow the split presented in work (Johnson et al. 2016), i.e., 29,000 images for training, 1,000 images for validation, and 1,000 images for testing. In the training phase, we select the most 8000 common words in the COCO captions to build our vocabulary, and each word in the vocabulary is presented as a 512-dimensional vector.

To automatically evaluate the quality of machine-generated captions remains a great challenge, mainly because of the fact that machines lack of the ability to make a suitable judgment independently as humans. Most of the existing evaluation metrics for image captioning task attempt to calculate a quantitative value according to the consistency of ground-truth annotations and generated captions. Here, we briefly introduce the evaluation

metrics used in the experiment, including BLEU (Papineni et al. 2002), METEOR (Banerjee and Lavie 2005), CIDEr-D (Vedantam et al. 2015) and ROUGE-L (Lin and Hovy 2003). BLEU is the most common metric for the evaluation of machine generated sentence. Since BLEU is based on n-gram precision, we choose BLEU-1, BLEU-2, BLEU-3 and BLEU-4 to evaluate the performance. METEOR is designed to make up for the deficiency of BLEU metric, and it takes sentence stems and synonyms into account to evaluate the generated sentence. CIDEr-D metric aims to measure the consensus between human annotations and generated captions, it is mainly designed for automatically image captioning. In addition, ROUGE-L metric pays more attention on recall rate, and is also employed to our experiments.

## 4.2 Implementation details

In this work, we use Faster R-CNN as the object detector to detect a number of objects. The Faster R-CNN is pretrained on the Visual Genome dataset, and then fine-tuned on the MSCOCO dataset. The RPN network, which is part of the Faster R-CNN, is also used to generate the region features of the given image. As a result, for each  $256 \times 256$  size image, We get 36 2048-dimensional image region vectors. In the decoding stage, we use the LSTM network as the decoder, and the input and hidden layers are both set to 512. The dimension of word embedding is also set to 512. Since complex network structure usually leads to overfitting, we adopt dropout method to randomly inactivate some neurons, here the dropout rate is set to 0.5. The hyper parameter  $\lambda$  is empirically set to 0.2, and we will discuss the selection of  $\lambda$  in the later.

The training process is divided into the following two stages: In the first stage, the model is trained under cross-entropy, and the mini-batch size is set to 64. In particular, we use the Adam optimization function to optimize our network, with the initial learning rate of  $5 \times 10^{-4}$  and the momentum of 0.9. At every 5 epochs, the learning rate is annealed by 0.7. In order to obtain a more generalized model, we select the BLEU-4 metric to monitor the training process. Early stopping is applied when the BLEU-4 score continues decrease in 5 consecutive epochs, and the maximum iteration is set to 30 epochs. In the second stage, the reinforcement learning optimization algorithm is run to further optimize the model. At this stage, the training epoch is set to 20, the training batch size is adjusted to 32, the learning rate is fixed at  $1 \times 10^{-4}$ , and other parameters remain unchanged. During the inference phase, we use beam search technology to select the most appropriate caption from candidate captions, and the beam size is set to 3. The maximum length of generated sentence here is set to 16.

## 4.3 Experimental results and analysis

### 4.3.1 Results on MSCOCO and flick30k datasets

As mentioned above, we empirically evaluate the effectiveness of our model on MSCOCO and flickr30k datasets. In the following, we compare our model with other state-of-the-art models, including attention based models, knowledge incorporated models and reinforcement learning Enhanced models.

Table 1 shows the comparison results of our proposed model with other state-of-the-art models on MSCOCO dataset. The best result of each metric in the table is shown in bold, as are the tables below. Except for the UP-Down model, our model outperforms all the

**Table 1** Performance comparison with other state-of-the-art models on the MSCOCO benchmark

Method	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr-D
Google NIC Vinyals et al. (2015)	0.666	0.461	0.329	0.246	–	–	–
Hard-Attention Xu et al. (2015)	0.718	0.504	0.357	0.250	0.230	–	–
ATT You et al. (2016)	0.731	0.565	0.424	0.316	0.250	0.535	0.943
Review Net Wu and Cohen (2016)	0.720	0.550	0.414	0.313	0.256	0.533	0.965
Areas-Att Pedersoli et al. (2017)	–	–	–	0.307	0.245	–	0.938
Adaptive Lu et al. (2017)	0.742	0.580	0.439	0.332	0.266	–	1.085
Saliency-Att Cornia et al. (2018)	0.708	0.536	0.391	0.284	0.248	0.521	0.898
CNet-NIC Zhou et al. (2019)	0.731	0.549	0.405	0.299	0.256	0.539	1.072
SCST:Att2in Rennie et al. (2017)	–	–	–	0.333	0.263	0.553	1.114
Up-Down Anderson et al. (2018)	<b>0.798</b>	–	–	0.363	<b>0.277</b>	0.569	1.204
Ours	0.793	<b>0.638</b>	<b>0.490</b>	<b>0.373</b>	0.273	<b>0.574</b>	<b>1.212</b>

Bold values indicate the best results of each metric

compared models. Even compared with the state-of-the-art Up-Down model, our model obtains superior results on several metrics, especially for BLEU, ROUGE-L and CIDEr-D metrics, we achieve BLEU-2 / BLEU-3 / BLEU-4 / ROUGE-L / CIDEr-D scores of 0.638, 0.490, 0.373, 0.574 and 1.212, respectively. The Up-Down model uses bottom-up image features acquired by Faster R-CNN and CIDEr-D optimization technology to train the network, which is similar to ours. Besides, this model especially uses two LSTMs to generate captions, which can increase the gains to some extent. In contrast, we only employ one LSTM at the decoding stage, this may reduce the performance of the model. However, by incorporating word attention and knowledge graph, our full model can achieve comparable results even better results than the UP-Down model.

To evaluate the effectiveness of our each design, we implement our several models with different components. Let RL denotes our reinforcement learning baseline model, RL+WA only incorporates word attention to boost captioning, RL+KG only introduces the knowledge graph into the baseline architecture, and RL+MA+KG represents the full model that combined with word attention and knowledge graph. We can see from Table 2, the RL+WA model and RL+KG model show better results over several metrics

**Table 2** Quantitative results on the use of different components

RL	WA	KG	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr-D
√			0.773	0.620	0.470	0.364	0.270	0.561	1.132
√	√		0.775	0.618	0.476	0.358	0.269	0.566	1.165
√		√	0.784	0.627	0.482	0.363	0.270	0.563	1.178
√	√	√	<b>0.793</b>	<b>0.638</b>	<b>0.490</b>	<b>0.373</b>	<b>0.273</b>	<b>0.574</b>	<b>1.212</b>

Bold values indicate the best results of each metric

than the baseline model (RL), and our RL+WA+KG model achieves best performance compared to other models. The results demonstrate that not only the performance of the model can be improved by using word attention and knowledge graph alone, but also the combination of these two designs can make the model get better performance. In addition, the use of knowledge graph can bring more benefits than word attention, indicating that through the incorporation of external knowledge, our model is able to discover more important cues to describe the given image.

Table 3 shows the comparison results on Flickr30K dataset. By incorporating reinforcement learning technology to optimize our model, it is not surprising that we outperform all the compared models. The flickr30k dataset contains less training data, thereby leading to a small reduction in the performance. Even that we achieve superior performance on all the standard evaluation metrics. Especially, we show the comparison results on the online MSCOCO evaluation server. As shown in Table 4. Compared to the state-of-the-art models, we also achieve better results on almost all the metrics.

### 4.3.2 The analysis of introducing external knowledge

Acquiring the knowledge from the previous work Li et al. (2019a), we find that each image in the captioning datasets (e.g. MSCOCO, Flickr30k) usually contains a small number of objects. Therefore, in the experiment, we select the top-3 objects according to the detected score, and inject them into the ConceptNet to retrieve semantic knowledge to boost the caption generation. Hence, given an image we get a series of relevant knowledge, with each piece of knowledge contains two parts: semantic entity and the relevant probability. Then, we consider to use  $\lambda$  to control the degree of introducing external knowledge. We intuitively choose the value of  $\lambda$  from 0 to 0.9. Figure 4 shows the change of BLEU-1, BLEU-2 and ROUGE-L scores conditioned on the selection of parameter  $\lambda$ . We can see that when  $\lambda = 0.2$  the BLEU-1 score and ROUGE-L score reach the peak simultaneously, while the BLEU-2 score reaches its maximum at  $\lambda = 0.3$ . With the  $\lambda$  value continues to increase, the scores of these three metrics gradually decrease. We speculate that the large  $\lambda$  may reduce the stability of the training model, while a smaller  $\lambda$  provides little benefits to the caption generation process. To bring a bigger benefit, we set the  $\lambda$  value to 0.2 in other experiments.

**Table 3** performance comparison with other state-of-the-art models on the flickr30k benchmark


Method	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr-D
Google NIC Vinyals et al. (2015)	0.663	0.423	0.277	0.183	–	–	–
Hard-Attention Xu et al. (2015)	0.669	0.439	0.296	0.199	0.185	–	–
ATT You et al. (2016)	0.647	0.460	0.324	0.230	0.189	–	–
SCA-CNN Chen et al. (2017a)	0.662	0.468	0.325	0.223	0.195	–	–
Adaptive Lu et al. (2017)	0.677	0.494	0.354	0.251	0.204	–	0.531
Att-Region Wu et al. (2017)	0.730	0.550	0.400	0.280	–	–	–
ours	<b>0.745</b>	<b>0.566</b>	<b>0.417</b>	<b>0.313</b>	<b>0.221</b>	<b>0.498</b>	<b>0.716</b>

Bold values indicate the best results of each metric

**Table 4** Performance comparison with other state-of-the-art models on the online MSCOCO server

Method	BLEU-1		BLEU-4		METEOR		ROUGE-L		CIDEr-D	
	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40
Google NIC Vinyals et al. (2015)	0.713	0.895	0.309	0.587	0.254	0.346	0.530	0.682	0.943	0.946
Hard-Attention Xu et al. (2015)	0.705	0.881	0.277	0.537	0.241	0.322	0.516	0.654	0.865	0.893
ATT You et al. (2016)	0.731	0.900	0.316	0.599	0.250	0.335	0.535	0.682	0.943	0.958
Review Net Wu and Cohen (2016)	0.720	0.900	0.313	0.597	0.256	0.347	0.533	0.686	0.965	0.969
MSM Yao et al. (2017b)	0.739	0.919	0.330	0.632	0.256	0.350	0.542	0.700	0.984	1.003
Adaptive Lu et al. (2017)	0.748	0.920	0.336	0.637	0.264	0.359	0.550	0.705	1.042	1.059
SCST:Att2in Rennie et al. (2017)	0.781	0.937	0.352	0.645	0.270	0.355	0.563	0.707	1.147	1.167
Up-Down Anderson et al. (2018)	<b>0.802</b>	<b>0.952</b>	0.369	0.685	0.276	0.367	0.571	0.724	1.179	1.205
ours	0.798	0.950	<b>0.376</b>	<b>0.688</b>	<b>0.284</b>	<b>0.376</b>	<b>0.580</b>	<b>0.730</b>	<b>1.218</b>	<b>1.229</b>

Bold values indicate the best results of each metric



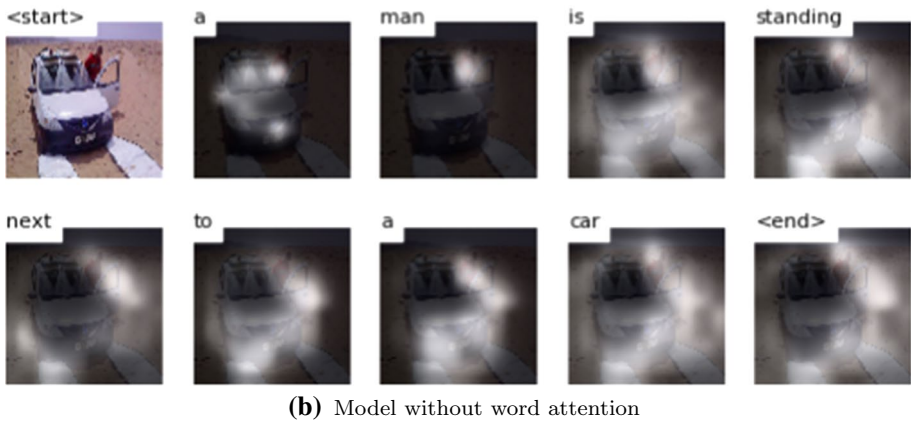
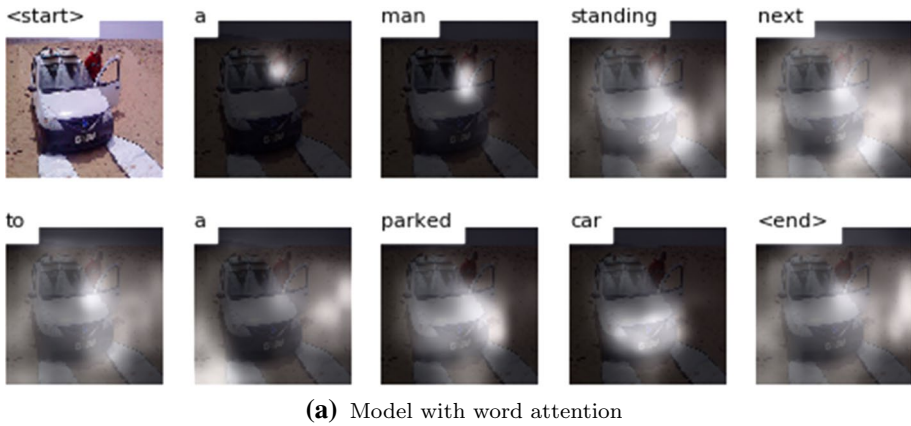
**Fig. 4** The change of BLEU-1, BLEU-2 and ROUGE-L scores after selecting different  $\lambda$  values. This experiment is conducted on MSCOCO benchmark

### 4.3.3 Attention analysis

As the previous work (Xu et al. 2015), we visualize the attention on individual pixels, thereby better revealing the correctness of the visual attention. Figure 5 shows an example of word generation when attending to different regions of the image. (a) and (b) are models with and without word attention respectively. As indicated by this example, when generating the descriptive words, especially the more important words, like “man” and “car”, the model with word attention can more accurately focus on the appropriate positions of the image. Compared to the model without word attention, the model with word attention can predict the word “parked”, indicating that more fine-grained captions can be generated by our proposed model. The result shows that our proposed word attention combining with the standard visual attention can perform well with the word generation process, and facilitate to generate more accurate and fine-grained captions.

### 4.3.4 Qualitative result analysis

Furthermore, to evaluate the quality of captions generated by the entire model and further verify the effectiveness of our proposed two scenarios, we provide a qualitative analysis of the results. The visualization results are shown in Fig. 6. It can be seen that by combining word attention and knowledge graph, our full model can generate more fine-grained captions, as well as reveal more implicit aspects of the image, which are not easily to be discovered by machines, but it seems not difficult for humans. For example, look at the first picture, our model can predict the object “court” even if it does not straightforwardly appear in the ground-truth captions. In addition, the model prefers to use detected objects to describe the image, of course these objects also appear in our constructed semantic knowledge corpus. However, like most existing models, limited by the caption length, the proposed model does not perform well for complex images with multiple objects. See the last picture, the model can’t predict the object “table”. It is notable that the multi-object captioning involves another more challenging study of artificial intelligence, i.e., dense



**Fig. 5** An example illustrating the effectiveness of word attention. < start > and < end > are tokens representing the beginning and end of sentence respectively

captioning, here we only consider generating a caption of a simple image. In general, our model is suitable for the description of most scenes and achieves performance compared to the state-of-the-art models.

## 5 Conclusion

In this paper, we explore to incorporate more useful semantic knowledge, including the internal annotation knowledge and the external knowledge extracted from knowledge graph, to reason about more accurate and meaningful captions. We first propose a new text-dependent attention mechanism, which we call word attention, to improve the correctness of basic visual attention when generating sequential descriptions word-by-word. The special word attention provides important semantic information to the calculation of visual attention. We demonstrate that our proposed model incorporating word attention as well as visual attention can significantly improve the agreement between





**Detected objects:**  
person; tennis racket; sport ball

**Knowledge entities:**  
man; woman; tennis; pickleball;  
badminton;; racquet; racket; court;  
hardcourt


**Generated:**  
A tennis player is playing tennis in the court.  
**Annotations:**  
A man in blue tennis shorts and a white polo shirt braces to hit the tennis ball.  
A tennis player locked on to the tennis ball with his racket and eyes .  
A middle-aged man is about to hit a tennis ball with his racket .  
A tennis player attempts a backhanded shot .  
A man in a white shirt is playing tennis .



**Detected objects:**  
person; surfboard; beach

**Knowledge entities:**  
man; woman; surfboard; wave;  
surfing; surfwear; surfer;  
paddleboard; water; overwater

**Generated:**  
A man riding a surfboard on top of a wave.  
**Annotations:**  
Surfer performing one of his moves on a nice sunny day with a scene in the background  
A skilled surfer grabbing the tip of his board while the rides a wave .  
One male surfer doing a turn on a surfboard in the ocean .  
A surfer in a black wetsuit catches a small wave .  
Man surfing a wave and pulling off a trick .



**Detected objects:**  
person; cat; aeroplane

**Knowledge entities:**  
man; woman; girl; cat; dog; kitten;  
sharp teeth; aeroplane; airplane; fly  
machine

**Generated:**  
A girl is playing a toy airplane next to a cat.  
**Annotations:**  
A little girl dressed in pink on the floor playing with a toy airplane with a cat looking on.  
A cat is watching the movement of a young girl who is playing with a toy airplane.  
A girl and her cat are curious to see how this airplane is built with legos.  
A small child plays with her airplane as a cat looks on.  
A cat is watching a girl construct a Lego airplane .



**Detected objects:**  
person; chair; dining table

**Knowledge entities:**  
man; woman; table; tabletop;  
worktable; placemat; chair; sofa; bed;  
food;

**Generated:**  
A man in a black shirt is speaking using powerpoint.  
**Annotations:**  
A man in a black polo and khakis is giving a powerpoint presentation.  
Man presenting a powerpoint presentation for several people.  
A man in a black shirt is addressing a group of people.  
A gentlemen is giving a presentation via powerpoint.  
A man is standing before a classroom and speaking.

**Fig. 6** Some examples of captions generated by our full model. The words with green color are the important words to be described

regions in image and words in sentence. Then, in order to facilitate meaningful captioning and overcome the problem of misprediction, we introduce a new strategy to incorporate commonsense knowledge extracted from knowledge graph into the encoder-decoder framework. This has indeed enhanced the generalization of our captioning model. Furthermore, we exploit reinforcement learning to optimize our training process, thereby making a significant improvement of the captioning performance. By combining the above mentioned several strategies, we achieve state-of-the-art performance on several standard evaluation metrics.

In the future work, we expect to independently construct a more compact knowledge graph by using sentence level and image level semantic information of a given instance to boost image captioning.

**Acknowledgements** This work is supported by the National Natural Science Foundation of China (Nos. 61966004, 61663004, 61866004, 61762078), the Guangxi Natural Science Foundation (Nos. 2019GXNS-FDA245018, 2018GXNSFDA281009, 2017GXNSFAA198365), the Guangxi “Bagui Scholar” Teams for Innovation and Research Project, the Guangxi Talent Highland Project of Big Data Intelligence and Application, Guangxi Collaborative Innovation Center of Multi-source Information Integration and Intelligent Processing.

## References

- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., & Zhang, L. (2018). Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6077–6086).
- Anne H. L., Venugopalan, S., Rohrbach, M., Mooney, R., Saenko, K., & Darrell, T. (2016). Deep compositional captioning: Describing novel object categories without paired training data. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1–10).
- Bai, S., & An, S. (2018). A survey on automatic image caption generation. *Neurocomputing*, 311, 291–304.
- Banerjee, S., & Lavie, A. (2005). Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization* (pp. 65–72).
- Bengio, S., Vinyals, O., Jaitly, N., & Shazeer, N. (2015). Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in neural information processing systems* (pp. 1171–1179).
- Chen, L., Zhang, H., Xiao, J., Nie, L., Shao, J., Liu, W., & Chua, T.S. (2017a). Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5659–5667).
- Chen, M., Ding, G., Zhao, S., Chen, H., Liu, Q., & Han, J. (2017b). Reference based lstm for image captioning. In *Thirty-first AAAI conference on artificial intelligence*.
- Chen, T.H., Liao, Y.H., Chuang, C.Y., Hsu, W.T., Fu, J., & Sun, M. (2017c). Show, adapt and tell: Adversarial training of cross-domain image captioner. In *Proceedings of the IEEE international conference on computer vision* (pp. 521–530).
- Chunseong Park, C., Kim, B., & Kim, G. (2017). Attend to you: personalized image captioning with context sequence memory networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 895–903).
- Cornia, M., Baraldi, L., Serra, G., & Cucchiara, R. (2018). Paying more attention to saliency: Image captioning with saliency and context attention. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 14(2), 1–21.
- Gu, J., Zhao, H., Lin, Z., Li, S., Cai, J., & Ling, M. (2019). Scene graph generation with external knowledge and image reconstruction (pp. 1969–1978).
- Hossain, M. Z., Sohel, F., Shiratuddin, M. F., & Laga, H. (2019). A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys (CSUR)*, 51(6), 1–36.
- Huang, L., Wang, W., Chen, J., & Wei, X. (2019). Attention on attention for image captioning (pp. 4634–4643).
- Jia, X., Gavves, E., Fernando, B., & Tuytelaars, T. (2015). Guiding the long-short term memory model for image caption generation. In *Proceedings of the IEEE international conference on computer vision* (pp. 2407–2415).
- Jiang, W., Ma, L., Chen, X., Zhang, H., & Liu, W. (2018a). Learning to guide decoding for image captioning. In *Thirty-second AAAI conference on artificial intelligence*.
- Jiang, W., Ma, L., Jiang, Y.G., Liu, W., & Zhang, T. (2018b). Recurrent fusion network for image captioning. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 499–515).
- Johnson, J., Karpathy, A., & Fei-Fei, L. (2016). Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4565–4574).
- Kim, B., Han Lee, Y., Jung, H., & Cho, C. (2018). Distinctive-attribute extraction for image captioning. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 0–0).
- Li, G., Su, H., & Zhu, W. (2017a). Incorporating external knowledge to answer open-domain visual questions with dynamic memory networks. [arXiv: Computer Vision and Pattern Recognition](https://arxiv.org/abs/1708.02877).
- Li, L., Tang, S., Deng, L., Zhang, Y., & Tian, Q. (2017b). Image caption with global-local attention. In *Thirty-first AAAI conference on artificial intelligence*.

- Li, X., Jiang, S., & Han, J. (2019a). Learning object context for dense captioning. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 33, pp. 8650–8657).
- Li, Y., Yao, T., Pan, Y., Chao, H., & Mei, T. (2019b). Pointing novel objects in image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 12497–12506).
- Lin, C.Y., & Hovy, E. (2003). Automatic evaluation of summaries using n-gram co-occurrence statistics. In *Proceedings of the 2003 human language technology conference of the North American chapter of the association for computational linguistics* (pp. 150–157).
- Liu, C., Mao, J., Sha, F., & Yuille, A. (2017a). Attention correctness in neural image captioning. In *Thirty-first AAAI conference on artificial intelligence*.
- Liu, S., Zhu, Z., Ye, N., Guadarrama, S., & Murphy, K. (2016). Optimization of image description metrics using policy gradient methods. arXiv preprint [arXiv:1612.003705](https://arxiv.org/abs/1612.003705)
- Liu, S., Zhu, Z., Ye, N., Guadarrama, S., & Murphy, K. (2017b). Improved image captioning via policy gradient optimization of spider. In *Proceedings of the IEEE international conference on computer vision* (pp. 873–881).
- Lu, J., Xiong, C., Parikh, D., & Socher, R. (2017). Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 375–383).
- Mathews, A.P., Xie, L., & He, X. (2016). Senticap: Generating image descriptions with sentiments. In *Thirtieth AAAI conference on artificial intelligence*.
- Papineni, K., Roukos, S., Ward, T., & Zhu, W.J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics* (pp. 311–318). Association for Computational Linguistics.
- Pedersoli, M., Lucas, T., Schmid, C., & Verbeek, J. (2017). Areas of attention for image captioning. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1242–1250.
- Qin, Y., Du, J., Zhang, Y., & Lu, H. (2019). Look back and predict forward in image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 8367–8375).
- Ramanishka, V., Das, A., Zhang, J., & Saenko, K. (2017). Top-down visual saliency guided by captions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7206–7215).
- Ranzato, M., Chopra, S., Auli, M., & Zaremba, W. (2015). Sequence level training with recurrent neural networks. arXiv preprint [arXiv:1511.06732](https://arxiv.org/abs/1511.06732)
- Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems* (pp. 91–99).
- Ren, Z., Wang, X., Zhang, N., Lv, X., & Li, L.J. (2017). Deep reinforcement learning-based image captioning with embedding reward. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 290–298).
- Rennie, S.J., Marcheret, E., Mroueh, Y., Ross, J., & Goel, V. (2017). Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7008–7024).
- Speer, R., Chin, J., & Havasi, C. (2017). Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-first AAAI conference on artificial intelligence*.
- Vedantam, R., Lawrence Zitnick, C., & Parikh, D. (2015). Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4566–4575).
- Vinyals, O., Toshev, A., Bengio, S., & Erhan, D. (2015). Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 3156–3164).
- Wang, Y., Lin, Z., Shen, X., Cohen, S., & Cottrell, G.W. (2017). Skeleton key: Image captioning by skeleton-attribute decomposition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7272–7281).
- Wu, Q., Shen, C., Wang, P., Dick, A., & van den Hengel, A. (2017). Image captioning and visual question answering based on attributes and external knowledge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6), 1367–1381.
- Wu, Z., & Cohen, R. (2016). Encode, review, and decode: Reviewer module for caption generation. arXiv preprint [arXiv:1605.07912](https://arxiv.org/abs/1605.07912)
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., & Bengio, Y. (2015). Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning* (pp. 2048–2057).
- Yao, T., Pan, Y., Li, Y., & Mei, T. (2017a). Incorporating copying mechanism in image captioning for learning novel objects. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6580–6588).
- Yao, T., Pan, Y., Li, Y., Mei, T. (2019). Hierarchy parsing for image captioning (pp. 2621–2629).

- Yao, T., Pan, Y., Li, Y., Qiu, Z., & Mei, T. (2017b). Boosting image captioning with attributes. In *Proceedings of the IEEE international conference on computer vision* (pp. 4894–4902).
- You, Q., Jin, H., Wang, Z., Fang, C., & Luo, J. (2016). Image captioning with semantic attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4651–4659).
- Zhou, Y., Sun, Y., & Honavar, V. (2019). Improving image captioning by leveraging knowledge graphs. In *2019 IEEE winter conference on applications of computer vision (WACV)* (pp. 283–293). IEEE.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.