

# Improving Image Captioning by Concept-Based Sentence Reranking

Xirong Li<sup>1,2</sup> and Qin Jin<sup>1,2(✉)</sup>

<sup>1</sup> Key Lab of DEKE, Renmin University of China, Beijing, China  
qjin@ruc.edu.cn

<sup>2</sup> Multimedia Computing Lab, Renmin University of China, Beijing, China

**Abstract.** This paper describes our winning entry in the ImageCLEF 2015 image sentence generation task. We improve Google’s CNN-LSTM model by introducing *concept-based sentence reranking*, a data-driven approach which exploits the large amounts of concept-level annotations on Flickr. Different from previous usage of concept detection that is tailored to specific image captioning models, the propose approach reranks predicted sentences in terms of their matches with detected concepts, essentially treating the underlying model as a black box. This property makes the approach applicable to a number of existing solutions. We also experiment with fine tuning on the deep language model, which improves the performance further. Scoring METEOR of 0.1875 on the ImageCLEF 2015 test set, our system outperforms the runner-up (METEOR of 0.1687) with a clear margin.

**Keywords:** Image captioning · Sentence reranking · Neural language modeling · ImageCLEF 2015 benchmark evaluation

## 1 Introduction

In this paper we tackle the challenging task of image captioning. Given an unlabeled image, the task is to automatically generate a natural language sentence that describes main entities and events present in the image. See Fig. 1 for some example sentences.

There has been a considerable progress on the topic in the last few years, thanks to powerful image representation derived from deep convolutional neural networks (CNN) [1] and trainable recurrent neural networks (RNN) capable of modeling long term dependency in natural language [2]. Joint models of CNN and RNN have demonstrated quite promising results for image captioning [3–5].

Since the number of model parameters to be optimized is at a million scale, what also matters are the growing amounts of training images associated with manually written descriptions, e.g., Flickr8k [7], Flickr30k [8], MSCOCO [9] with more than 100k images, and Flickr8k-CN [10] as a bilingual extension of Flickr8k. Even though the number of captioned images has increased from a few thousands to over 100k, sentence-level annotations remain in shortage when compared to



- (1) a plane taking off from a runway
- (2) a plane taking off from the runway
- (3) a plane flying in the sky
- (4) an airplane flying in the sky
- (5) a plane flying in the air
- (6) an airplane is flying in the sky
- (7) a plane taking off into the sky

- (3) a plane flying in the sky  
[ flying | plane | sky ]

(a)



- (1) a plant with green leaves
- (2) a cluster of yellow flowers
- (3) a plant is growing in a garden
- (4) a plant with green leaves and green leaves
- (5) a plant with green leaves and purple flowers
- (6) a cluster of yellow flowers in a garden
- (7) a cluster of yellow flowers in a field

- (5) a plant with green leaves and purple flowers  
[ purple | plant | green ]

(b)

**Fig. 1. Examples illustrating concept-based sentence re-ranking for improving image captioning.** Candidate sentences generated by a CNN-LSTM image captioning model are shown in descending order on the left side, while the chosen sentences are shown on the right side. In square brackets are concepts predicted by the neighbor voting algorithm [6]. The candidate sentence best matching the predicted concepts are chosen as the final description. (Color figure online)

concept-level annotations. For instance, one can easily obtain one million learning examples of ‘dog’ from Flickr<sup>1</sup>. How to exploit such large amounts of (noisy) concept-level annotations to improve image captioning is important.

Works on utilizing concept detection for image captioning exist, but are tailored to specific models. For instance, in the work by Fang *et al.* [11], concept

<sup>1</sup> Over 6 million images tagged with ‘dog’ on Flickr, <https://www.flickr.com/search/?tags=dog>, retrieved on April-29-2016.

detection results are used as part of input features for a maximum entropy language model. Gong *et al.* [12] leverage Flickr data to improve image and text embedding within a kernel Canonical Correlation Analysis framework. These works are inapplicable to other models, e.g., the popular CNN+LSTM architecture.

**Contributions of this work.** For improving image captioning by exploiting concept-level annotations, we propose concept-based sentence reranking. The proposed technique treats the underlying sentence generation model as a black box, making it applicable to a number of existing solutions. Moreover, by deriving concept detectors from Flickr images, better image descriptions are predicted with no extra cost of manual annotation. In addition, we show that a fine tuning on a trained deep language model further improves the performance. Putting all this together, our system is the winning entry in the ImageCLEF 2015 image sentence generation task [13, 14].

## 2 Related Work

Depending on whether candidate sentences are given in advance, we see two lines of research, namely retrieval approaches and generative approaches.

The retrieval approaches sort a predefined set of candidate sentences in terms of their relevance with respect to a given image, and then selects the top ranked sentence to annotate the image [7, 12, 15]. To compute cross-media relevance between images and sentences, representing them in a common space is a prerequisite. In [7], this space is derived by Kernel Canonical Correlation Analysis, while normalized linear Canonical Correlation Analysis is used in [12] for scaling to larger training sets. In the DeVISE model [15], the common space is formed by a pretrained *Word2Vec* [16], where the embedding vector of a sentence is obtained by averaging over the vectors of its words. Compared to bag-of-words used in [7, 12], the use of *Word2Vec* enables DeVISE to handle a much larger vocabulary. In [17], the common space is implemented as a visual CNN feature space in order to better capture visual and semantic similarity. One strength of the retrieval approaches is that a predicted sentence is syntactically and grammatically correct. However, they lack the ability to generate novel sentences.

The generative approaches typically consist of two main components, i.e., image encoding and sentence decoding [3–5]. For image encoding, a pretrained deep convolutional neural network (CNN) is employed to project a specific image into a visual feature space. For cross-media comparison, images and words are embedded into a latent common space before they are fed to a recurrent neural network (RNN), which eventually outputs a sequence of words as the caption. Variants of RNN have been investigated, including multi-modal RNN [4], bidirectional RNN [5], and Long Short-Term Memory (LSTM) [3]. Recent benchmark evaluations, e.g., the MSCOCO Captioning Challenge [9] and the ImageCLEF 2015 image sentence generation task [14], have demonstrated outstanding performance of LSTM based solutions [3, 13].

### 3 Our Approach

We aim to improve the performance of an existing image captioning model by re-ranking its generated sentences in terms of concept detection results. More concretely, given an image, the model first generates  $k$  best sentences. Each sentence, indicated by  $hypoSent$ , is associated with a confidence score, denoted as  $sentenceScore(hypoSent)$ . Meanwhile, suppose we have access to a concept detection system that predicts  $m$  concepts deemed to be relevant with respect to the given image. If a detected concept appears in a  $hypoSent$ , we call it a matched concept. Our assumption is that a sentence covering more matched concepts is more likely to be a good description of the image. Hence, the final caption shall be the sentence that maximizes the prediction confidence score and the concept matches. The new sentence confidence score is therefore computed as a linear combination of these two factors, i.e.,

$$newScore(hypoSent) = \theta \cdot concScore(hypoSent) + (1 - \theta) \cdot sentScore(hypoSent), \quad (1)$$

where  $concScore(hypoSent)$  is the averaged confidence score of all the matched concepts in the hypothesis sentence, and  $\theta \in [0, 1]$  is a trade-off parameter. Figure 1 show cases some examples of how this simple strategy helps produce better image descriptions.

The proposed approach treats the image captioning component as a black box, meaning any captioning model can be applied in theory. In this work we adopt Google’s CNN-LSTM framework [3] for its state-of-the-art performance. The framework is described in Sect. 3.1. We then discuss in Sect. 3.2 how to adapt the model to new target data, followed by our choices of concept detection in Sect. 3.3.

#### 3.1 The CNN-LSTM Model for Sentence Generation

To generate the hypothesis sentences for a specific image, we employ the CNN-LSTM model proposed by Vinyals *et al.* [3]. At the heart of the model is an LSTM based recurrent neural network, which computes the posterior probability of a (novel) sentence conditioned on the image.

The network is trained by maximum likelihood estimation using many pairs of image and sentence. As an image and a sentence are of different modalities and thus not directly comparable, an image projection matrix and a word projection matrix are deployed to embed the image and the sentence from their initial representations into a common space before feeding them into the network. We use a CNN feature, i.e., the last fully connected layer of a pretrained 16-layer VGGNet [18], as the initial representation of an image. Each word in a sentence is represented by a one-hot vector.

In the sentence generation stage, a hypothesis sentence of the input image is generated as a sequence of words in a greedy manner. In particular, the CNN vector of the input image, after projection, is fed into the LSTM network to initialize its memory units. The posterior probability over all words is re-estimated

based on the memory status and the words already been chosen in previous iterations. Consequently, the word with the highest probability is selected, while its embedding vector will be fed into the LSTM network in the next iteration. The generation process stops once a predefined STOP token is chosen. Similar to [3], we use beam search to obtain a list of hypothesis sentences.

### 3.2 Model Adaptation by Fine Tuning

We use stochastic gradient descent to train the LSTM based sentence generation model. Such an incremental learning strategy makes it relatively ease to adjust an existing model with respect to novel data. In particular, we leverage a transfer learning strategy similar to the ones used in the context of object recognition for fine tuning a CNN model [19]. That is, re-train the LSTM model using a novel training dataset but with a relatively lower learning rate.

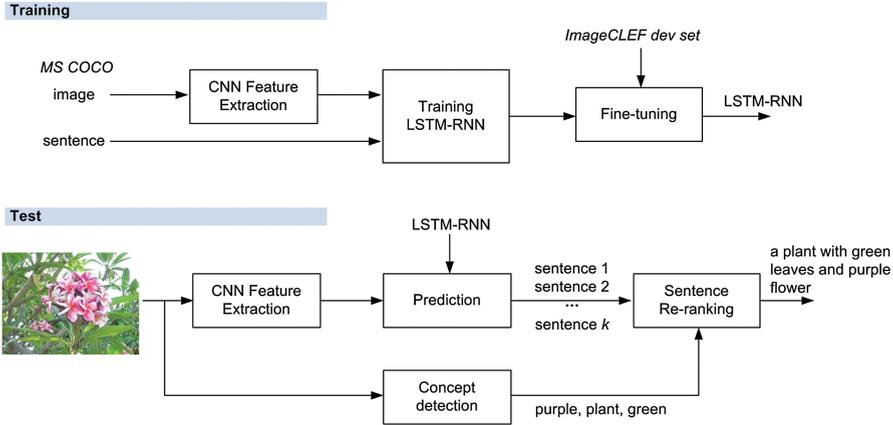
### 3.3 Concept Detection

Notice that the output of the CNN model corresponds to the 1,000 labels defined in the Large Scale Visual Recognition Challenge [20]. We do not directly use this as concept detection results, because the labels are often over specific to be overlapped with words used in image captions. Instead, we consider two methods for concept detection, i.e., neighbor voting [6] and hierarchical semantic embedding [21], both of which are capable of learning from large amounts of weakly labeled Flickr images, and thus predict words used in daily life.

**Neighbor voting** (NeiVote) [6]. Given an image, this method first retrieves a set of neighbor images visually close to the given image, and then count the occurrence of a specific concept in textual annotations associated with the neighbors. The concepts are sorted in descending order by their occurrence and the top  $m$  ranked ones are preserved. Despite its simplicity, a recent comparison [22] shows that the method remains competitive when compared to more complicated alternatives. The visual distance between images are computed in terms of the Euclidean distance between their CNN features.

**Hierarchical Semantic Embedding** (HierSE) [21]. Developed in the context of zero-shot image tagging, HierSE casts concept detection into cross-media relevance computation. In particular, HierSE embeds both concept and image into a *Word2Vec* space. The embedding vector of the concept is obtained by convex combination of the embedding vectors of the concept and its ancestors tracing back to the root in WordNet. The embedding vector of an image is obtained by convex combination of the embedding vectors of the top ten labels predicted by the CNN model. Consequently, the image-concept relevance score is computed as the cosine similarity between the corresponding embedding vectors. The top  $m$  concepts with the largest scores are preserved.

A conceptual diagram of the proposed approach is given in Fig. 2.



**Fig. 2. An illustration of the proposed solution.** We improve the popular CNN+LSTM architecture through (1) model adaption by fine tuning and (2) sentence reranking using auto-detected visual concepts. (Color figure online)

## 4 Evaluation

To verify the effectiveness of our approach, we participated in the ImageCLEF 2015 image sentence generation task [14]. In this task, participants were asked to generate sentence-level textual descriptions for 500k images crawled from the web, from which the task organizers selected a subset of 3,070 images for performance assessment. Notice that we have no access to ground-truth annotations, as the test set is used for blind testing by the organizers only.

In ImageCLEF 2015, a dev set of 2k images with manually written captions are provided for system development. The number of captions per image ranges from 5 to 51 per image, with a mean of 9.5 and a median of 8 descriptions. In order to tune the weighting parameter  $\theta$  in Eq. 1, we randomly split the dev set into three disjoint subsets, i.e., 1,600 images for training, 200 images for validation and the remaining 200 images for internal test.

**Table 1.** Performance of our solution under varied settings on the ImageCLEF 2015 image sentence generation task.

Training data	Data for fine tuning	Concept detection	METEOR score
ImageCLEF dev	–	–	0.1659
MSCOCO	ImageCLEF dev	–	0.1759
ImageCLEF dev	–	HierSE	0.1781
ImageCLEF dev	–	NeiVote	0.1806
MSCOCO	ImageCLEF dev	HierSE	0.1684
MSCOCO	ImageCLEF dev	NeiVote	<b>0.1875</b>

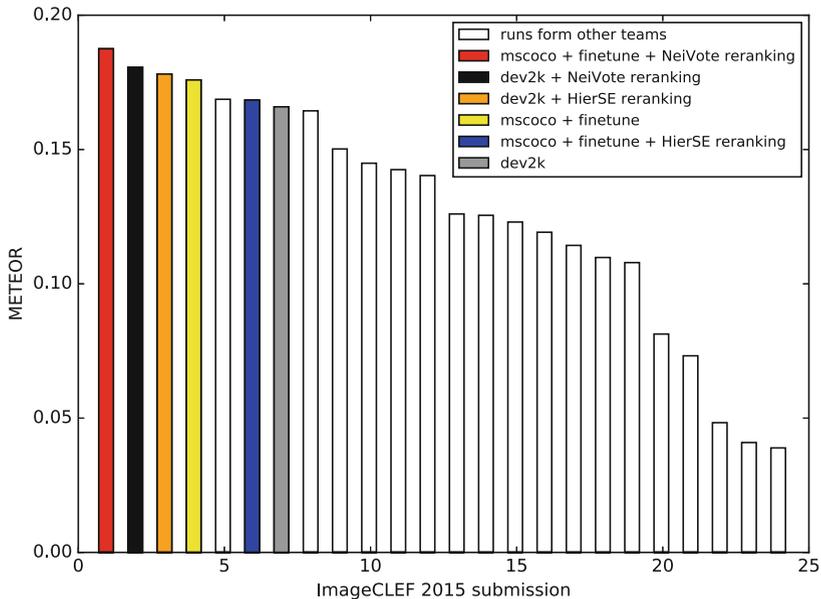
#### 4.1 Experiment 1. The Influence of Fine Tuning

We compare two runs. One run is to train the sentence generation model on the ImageCLEF dev set. The other run is to first train the model on the MSCOCO dataset [9], and then fine tune it on the ImageCLEF dev set using a lower learning rate. Following the protocol [14], we use the METEOR score to assess the performance of image captioning. As shown in the first two rows of Table 1, the fine-tuned model, with METEOR of 0.1759, performs better.

#### 4.2 Experiment 2. The Effect of Concept-Based Sentence Reranking

Given detection results either from NeiVote or from HierSE, we apply concept-based sentence reranking on sentences generated by the CNN-LSTM models with and without fine tuning, respectively. For NeiVote, we retrieve neighbor images from a collection of 2 million Flickr images. For HierSE, we adopt an existing *Word2Vec* model pre-trained on Flickr tags [21].

As we see from Table 1, the performance of the model (without fine tuning) consistently improves after sentence reranking, with the METEOR score increases from 0.1659 to 0.1781 (HierSE) and 0.1806 (NeiVote), respectively. For the fine-tuned model, using HierSE for concept detection causes some performance drop, while NeiVote remains effective, reaching the best METEOR score



**Fig. 3. Comparing with submissions from other teams in the ImageCLEF 2015 image sentence generation task.** Our results given varied settings are highlighted in color bars. The submissions have been sorted in descending order according to their METEOR scores. (Color figure online)



a group of people standing in front of a bridge



a view of a street with a sign in the background



the inside of a car



a deer grazing in a field



a train crossing the road



a room with many chairs and a table



a picture of a poster



a plant with yellow flowers



a view of a beach with a blue sky

**Fig. 4. Some test images with sentences generated by the proposed system.** The images are hand picked that showing our system performs well. (Color figure online)

of 0.1875. The results justify the effectiveness of concept-based sentence reranking, and NeiVote appears to be a better choice for concept detection in this context.

Figure 3 plots the performance of our system in the context of all submissions in ImageCLEf 2015, showing our leading position in the evaluation. Image captioning results of some selected examples are given in Fig. 4. The quantitative and qualitative results demonstrate the potential of the proposed concept-based sentence reranking.

Note that due to the unavailability of the ground truth of the test dataset as aforementioned, we have compared mainly with submissions from other ImageCLEF participants. In our future work we will compare other state-of-the-art methods on other benchmarks, e.g., MSCOCO.

## 5 Conclusions

We present in this paper *concept-based sentence reranking*, a data-driven approach to improve image captioning. As demonstrated by our participation in the ImageCLEF 2015 benchmark evaluation, the proposed approach is found to be effective for improving the popular CNN-LSTM image captioning model. In essence the improvement is gained by exploiting the large amount of noisy concept-level annotations associated with Flickr images. In addition, fine tuning on the deep language model helps its generalization to novel data.

**Acknowledgements.** The authors are grateful to the ImageCLEF coordinators for the benchmark organization efforts [14,23]. This research was supported by the Fundamental Research Funds for the Central Universities and the Research Funds of Renmin University of China (No. 16XNQ013).

## References

1. Krizhevsky, A., Sutskever, I., Hinton, G.: ImageNet classification using deep convolutional neural networks. In: Proceedings of NIPS (2012)
2. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: Proceedings of NIPS (2014)
3. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: a neural image caption generator. In: Proceedings of CVPR (2015)
4. Mao, J., Xu, W., Yang, Y., Wang, J., Huang, Z., Yuille, A.: Deep captioning with multimodal recurrent neural networks (m-RNN). In: Proceedings of ICLR (2015)
5. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: Proceedings of CVPR (2015)
6. Li, X., Snoek, C., Worring, M.: Learning social tag relevance by neighbor voting. IEEE Trans. Multimedia **11**(7), 1310–1322 (2009)
7. Hodosh, M., Young, P., Hockenmaier, J.: Framing image description as a ranking task: data, models and evaluation metrics. J. Artif. Int. Res. **47**(1), 853–899 (2013)
8. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: new similarity metrics for semantic inference over event descriptions. TACL **2**, 67–78 (2014)
9. Lin, T., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. CoRR abs/1405.0312 (2014)
10. Li, X., Lan, W., Dong, J., Liu, H.: Adding Chinese captions to images. In: Proceedings of ICMR (2016)
11. Fang, H., Gupta, S., Iandola, F., Srivastava, R., Deng, L., Dollar, P., Gao, J., He, X., Mitchell, M., Platt, J., Zitnick, L., Zweig, G.: From captions to visual concepts and back. In: Proceedings of CVPR (2015)
12. Gong, Y., Wang, L., Hodosh, M., Hockenmaier, J., Lazebnik, S.: Improving image-sentence embeddings using large weakly annotated photo collections. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8692, pp. 529–545. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10593-2\\_35](https://doi.org/10.1007/978-3-319-10593-2_35)
13. Li, X., Jin, Q., Liao, S., Liang, J., He, X., Huo, Y., Lan, W., Xiao, B., Lu, Y., Xu, J.: RUC-tencent at ImageCLEF 2015: concept detection, localization and sentence generation. In: CLEF Working Notes (2015)

14. Gilbert, A., Piras, L., Wang, J., Yan, F., Dellandrea, E., Gaizauskas, R., Villegas, M., Mikolajczyk, K.: Overview of the ImageCLEF 2015 scalable image annotation. In: CLEF Working Notes, Localization and Sentence Generation Task (2015)
15. Frome, A., Corrado, G., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: DeViSE: a deep visual-semantic embedding model. In: Proceedings of NIPS (2013)
16. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Proceedings of NIPS (2013)
17. Dong, J., Li, X., Snoek, C.G.M.: Word2VisualVec: cross-media retrieval by visual feature prediction. CoRR abs/1604.06838 (2016)
18. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Proceedings of ICLR (2015)
19. Agrawal, P., Girshick, R., Malik, J.: Analyzing the performance of multilayer neural networks for object recognition. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8695, pp. 329–344. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10584-0\\_22](https://doi.org/10.1007/978-3-319-10584-0_22)
20. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. *Int. J. Comput. Vision* **115**(3), 211–252 (2015)
21. Li, X., Liao, S., Lan, W., Du, X., Yang, G.: Zero-shot image tagging by hierarchical semantic embedding. In: Proceedings of SIGIR (2015)
22. Li, X., Uricchio, T., Ballan, L., Bertini, M., Snoek, C.G.M., Del Bimbo, A.: Socializing the semantic gap: a comparative survey on image tag assignment, refinement and retrieval. *ACM Comput. Surv.* **49**(1), 14:1–14:39 (2016)
23. Villegas, M., et al.: General overview of ImageCLEF at the CLEF 2015 Labs. In: Mothe, J., Savoy, J., Kamps, J., Pinel-Sauvagnat, K., Jones, G.J.F., SanJuan, E., Cappellato, L., Ferro, N. (eds.) CLEF 2015. LNCS, vol. 9283, pp. 444–461. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-24027-5\\_45](https://doi.org/10.1007/978-3-319-24027-5_45)