

A SEQUENTIAL GUIDING NETWORK WITH ATTENTION FOR IMAGE CAPTIONING

Daouda Sow¹

Zengchang Qin^{1,2}

Mouhamed Niasse³

Tao Wan¹

¹ Intelligent Computing and Machine Learning Lab, School of ASEE, Beihang University, China

² Keep Labs, Keep Inc., Beijing, China ³ School of EEE, North China Electric Power University, China

{davissow, zcquin, taowan}@buaa.edu.cn

ABSTRACT

The recent advances of deep learning in both computer vision (CV) and natural language processing (NLP) provide us a new way of understanding semantics, by which we can deal with more challenging tasks such as automatic description generation from natural images. In this challenge, the encoder-decoder framework has achieved promising performance when a convolutional neural network (CNN) is used as image encoder and a recurrent neural network (RNN) as decoder. In this paper, we introduce a sequential guiding network that guides the decoder during word generation. The new model is an extension of the encoder-decoder framework with attention that has an additional guiding long short-term memory (LSTM) and can be trained in an end-to-end manner by using image/descriptions pairs. We validate our approach by conducting extensive experiments on a benchmark dataset, i.e., MS COCO Captions. The proposed model achieves significant improvement comparing to the other state-of-the-art deep learning models.

1. INTRODUCTION

Automatically describing the content of images is of one of the hardest tasks of scene understanding—a long standing problem of the field of artificial intelligence (AI). This task is challenging as it requires semantic understanding both of image and natural language, building a good learning models to shrink semantic gaps in different modalities. This unsolved problem has attracted a lot of attentions in AI community [1–14]. An image captioning model must detect all objects/scenes present in an image and sometimes even objects/scenes that are not present in the image but related to its description. Moreover, it should be able to capture the relations between the objects/scenes and express them in a well-formed, human-understandable sentence. This challenge has the significance not only in academic research, but also in various applications such as information retrieval and visual question-answering.

Some recent research of image captioning take inspiration from neural machine translation systems (NMT) [15–18] that successfully use sequence-to-sequence learning for translation. NMT models solve the task of translation by a two-fold pipeline. (1) A RNN is used to encode the source sentence in a fixed-length vector and then (2) a decoding RNN is conditioned on that fixed-length vector to generate a sentence in the target language, one word at each time step. This general encoder-decoder framework is well suited to the image description problem as the task is (sort of) equivalent to translating an image into its corresponding description. However, image captioning appears to be much harder than certain machine translation tasks. For instance, when translating from English to French, the source and target languages share similar sentence structures (similar part-of-speech order). This similarity in structure is very useful for

the translating system as the alignment will be much easier. Instead, for image captioning, the structure of the visual data is way different from the structure of the captions describing them. Moreover, the simple CNN+RNN pipeline squash the whole input image into a fixed-length embedding vector. This constitutes a major limitation of the basic encoder-decoder architecture.

To overcome these limitations both for machine translation and image captioning, some new models were proposed by using the *attention* mechanism [3, 16, 18]. For image captioning, attention mechanisms help the model focus on salient regions of the image while generating descriptions by dynamically updating the image representations at each time step. With this the input image is now represented as a sequence of context vector where the length of the sequence depends on the number of words in the sentence to be generated. Promising results has been published since attention was introduced in [16] then later refined in [18]. Another group of models [4–6] tried to overcome the limitation of the basic encoder-decoder framework by still representing the input image as a fixed-length vector but injecting external guiding information at each time step. The external guiding information can be any attribute features connecting the image to its description. For instance, the attribute features could be semantic information obtained from a multimodal space of images and their descriptions learned using Canonical Correlation Analysis [5] or the prediction of frequent word occurrences in captions [4] or even learned by an additional guiding network [6]. The guiding information is however static in all of these models and couldn't be adjusted during the process of generation. In this work, we investigate how can we take advantages of these encoder-decoder models by constructing a joint neural model with attention that has an extra guiding network. Our approach is more closely related to the work of [6] but instead of learning one magic guiding vector, we propose to learn a sequential guiding network that can adapt its guiding vector during words generation. More specifically, the guiding network is a long short-term memory (LSTM) which outputs a guiding vector at each time step based on previous guiding vectors, current attention vector and attribute features. We use the Luong style of attention [18] which is a refined version of attention mechanism and that to the best of our knowledge, there has not been any published work reporting the performance of an image captioning model that is built following *only* the encoder-decoder pipeline with Luong style of attention. Furthermore, to demonstrate the usefulness of the guiding LSTM, we also compare the performance of our model with and without the the guiding LSTM in experiments.

2. RELATED WORK

In the past few years, the advances in training deep neural network models both for CV [19] and NLP [15] give new perspectives for au-

tomatic image description generation. The neural encoder-decoder framework of machine translation has recently been used for generating image captions because of the high level similarity of the two fields. Both fields aim to translate a source language to a target one. The model in [8] was the first to follow the encoder-decoder pipeline for image captioning. Authors of [8] use a CNN to compute image features and a LSTM model to encode the corresponding descriptions. The image features are projected into latent states of the LSTM encoder to construct a multimodal distributed representation learned by optimizing a simple pairwise ranking loss. Image descriptions were generated from the multimodal space using a novel multiplicative neural language model named Structure-Content Neural Language Model (SC-NLM). Their approach gives superior results than all previous models but was later outperformed by the model described in [2], which is a simpler encoder-decoder architecture, again directly inspired by Neural Machine Translation (NMT). [2] pass images through a deep CNN, take the activation of the last fully-connected layer as image features and then initialize the hidden states of a RNN cell with the CNN image features. During training at each time step they input the current word, compute a distribution over all the words of the vocabulary based on the hidden states and maximize the likelihood of the true next word using a negative log likelihood loss. The work in [1] employs a more powerful RNN cell, and they incorporated the image features as first input word instead of using it as initial hidden state. Other similar approaches include [5, 10, 11]. [5] were proposed as an extension of the LSTM model by exploring different kind of semantic information that can be used as extra guiding input to the LSTM during decoding steps. [4] followed this direction by injecting more powerful high-level image attributes into the decoder. In their work, they investigate different architecture for injecting word occurrence prediction attributes [9] into the CNN-RNN framework.

Inspired by the use of attention in sequence-to-sequence learning for machine translation [16, 18], visual attention has been proved to be a very effective way of improving image captioning. Some early research follows this direction, e.g., the model proposed in [3] can focus on important parts of images while generating descriptions. The captioning model in [3] is very similar in spirit to that in [7], in which visual representation is constructed for sentence parts while the description is being emitted. In [13], authors proposed a spatial and channel-wise attention mechanism over 3D CNN feature maps. In addition to spatial attention (standard visual attention), their model can also learn to pay attention over CNN channels, which they argued as extracting semantic attributes. In [6], visual attention is combined with semantic information for generating image captions. Their model can learn an additional guiding vector while learning to focus on image regions. However in contrast to our model, their framework only learns a fix guiding vector that couldn't be adapted during words generation.

A pure sequence-to-sequence architecture for image captioning is proposed in [12]. Different from previous approaches, their model represents images as a sequence of detected objects and a '*sequential attention layer*' is introduced to help the model focus on important objects. While resulting in a more complex architecture, their approach claims state-of-the-art results in all metrics. Instead of training via (penalized) maximum likelihood estimation, some recent works use Policy Gradient (PG) methods to directly optimize the non-differentiable testing metrics, claiming boost in term of performance measure. While [20] optimize for the standard CIDEr metric, [21] proposed to optimize for a new testing metric that is a linear combination of CIDEr [22] and SPICE [23] they called SPIDER, which they found better correlated with human judgment. However

in this line of work, it is not clear yet whether the improvement in testing metrics could result in captions with better quality.

3. PROPOSED MODEL

Technically, the ultimate goal of the neural CNN+RNN architecture for image captioning is to build an end-to-end model trainable by standard backpropagation that can generate a description S^i given a certain image X^i . Inspired by NMT, such a model can translate an image into a describing sentence. A CNN is first used to obtain image features and a RNN decoder is conditioned on those CNN image features to generate the describing sentence. Given a training dataset consisting of (S^i, X^i) pairs, these models aim to directly maximize the log-probability of generating S^i when X^i is the input. Thus, the optimization problem can be formulated by:

$$\theta^* = \arg \max_{\theta} \sum_i \log p(S^i | X^i; \theta) \quad (1)$$

where θ represents the set of parameters to be learned, X^i is a single image and $S^i = [w_1^i, w_2^i, \dots, w_{N^i}^i]$ is the corresponding caption which is a sequence of N^i words. Because each caption represents a sequence of N^i words, the log probability is calculated using Bayes chain rule:

$$\log p(S^i | X^i; \theta) = \log p(w_1^i | X^i; \theta) + \sum_{t=2}^{N^i} \log p(w_t^i | X^i, w_{1:t-1}^i; \theta) \quad (2)$$

where $p(w_1^i | X^i; \theta)$ is the likelihood of generating the first word w_1^i given only the image X^i and $p(w_t^i | X^i, w_{1:t-1}^i; \theta)$ represents the probability of emitting word w_t^i at time step t conditioned on the image X^i and the words generated so far $w_{1:t-1}^i$. In our work, we model the distribution $p(w_t^i | X^i, w_{1:t-1}^i; \theta)$ with a LSTM cell wrapped with Luong-style attention mechanism [18]. A RNN cell with Luong's attention computes its hidden state \mathbf{h}^t at each time step based on the current input \mathbf{x}^t , the previous hidden state \mathbf{h}^{t-1} and the previous attention vector $\tilde{\mathbf{h}}^{t-1}$. The current hidden state \mathbf{h}^t is combined with the image-side context vector \mathbf{c}^t to form the final output of the cell which is the current attention vector $\tilde{\mathbf{h}}^t$. Finally, the distribution $p(w_t^i | X^i, w_{1:t-1}^i; \theta)$ is computed by applying a Softmax layer on top of the current attention vector $\tilde{\mathbf{h}}^t$:

$$\begin{aligned} p(w_t^i | X^i, w_{1:t-1}^i; \theta) &= \text{Softmax}(\mathbf{W}_s \tilde{\mathbf{h}}^t) \\ \tilde{\mathbf{h}}^t &= \tanh(\mathbf{W}_c [\mathbf{c}^t; \mathbf{h}^t]) \\ \mathbf{h}^t &= \mathcal{R}(\mathbf{x}^t) \end{aligned} \quad (3)$$

where $\mathbf{W}_s, \mathbf{W}_c$ are projection matrices, \mathbf{c}^t is the image-side context vector detailed in section 3.1 and \mathcal{R} is a recursive function whose details are given in Section 3.2.

3.1. Image Features: Convolutional Neural Network

Convolutional neural networks (CNNs) are powerful models in Computer Vision with state-of-the-art performance in image classification and object detection. Let $\mathcal{A} = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_K\}$ denotes the set of annotation vectors (attention states) extracted from the last convolutional layer of the encoding CNN. Here K is the number of neurons in one activation map of the last convolutional layer and each \mathbf{a}_k is a D -dimensional feature vector where D is the number of activation maps of the last convolutional layer. With those attention

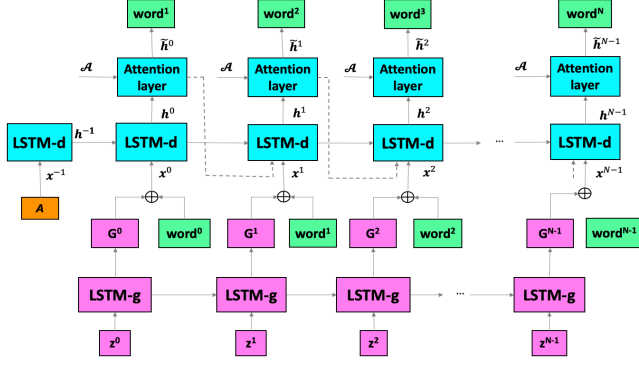


Fig. 1. Unrolled version of our model (better viewed in color).

states, our model computes a context vector c^t which is a weighted sum of attention states and can be seen as a dynamic representation of the image at every time step t :

$$\begin{aligned} s_k^t &= \Phi(a_k, h^t) \\ \alpha_k^t &= \frac{\exp(s_k^t)}{\sum_{j=1}^K \exp(s_j^t)} \\ c^t &= \sum_{k=1}^K \alpha_k^t \times a_k \end{aligned} \quad (4)$$

where $c^t \in \mathbb{R}^D$ and $\alpha_1^t, \alpha_2^t, \dots, \alpha_K^t$ are alignment weights and the function Φ is known as alignment function. In our model, we use the general form described in [18]:

$$\Phi(a_k, h^t) = h^{t\top} W_a a_k \quad (5)$$

where $W_a \in \mathbb{R}^{H \times D}$ is a transformation matrix.

3.2. Sentence Generator: LSTM + Luong's Attention

The form of the recursive function \mathcal{R} (Equation 3) is a critical design choice for generating sequences. In this paper, we use a LSTM cell wrapped with the attention mechanism described in [18] to form \mathcal{R} . LSTM [24] is a powerful form of recurrent neural network that is widely used now because of its ability to deal with issues like vanishing and exploding gradients. The final form of \mathcal{R} is described by the following equations:

$$\begin{aligned} i^t &= \sigma(W_{xh}^i x^t + W_{hh}^i h^{t-1} + W_{hh}^i \tilde{h}^{t-1} + b^i) \\ f^t &= \sigma(W_{xh}^f x^t + W_{hh}^f h^{t-1} + W_{hh}^f \tilde{h}^{t-1} + b^f) \\ o^t &= \sigma(W_{xh}^o x^t + W_{hh}^o h^{t-1} + W_{hh}^o \tilde{h}^{t-1} + b^o) \\ g^t &= \sigma(W_{xh}^g x^t + W_{hh}^g h^{t-1} + W_{hh}^g \tilde{h}^{t-1} + b^g) \\ m^t &= f^t \odot m^{t-1} + i^t \odot g^t \\ h^t &= o^t \odot \tanh(m^t) \end{aligned} \quad (6)$$

here x^t is input signal at time step t , m^t and h^t are respectively memory cell and hidden state of the LSTM cell and \tilde{h}^t represents attention vector. The variables i^t, f^t, o^t, g^t stand respectively for input gate, forget gate, output gate and candidate memory cell. The various W_{xy} and b^z are respectively parameter matrices and bias terms to be optimized. The non-linearity σ is element-wise sigmoid activation and \odot is the element-wise dot product.

3.3. Sequential Guiding Network

While the decoding function can access image features at each time step in the encoder-decoder framework with attention, injecting additional guiding vector to the decoder input signal can lead to higher performance. In our work, we extend the CNN+RNN architecture with attention by inserting an extra guiding network. Different from previous approaches that learn a static guiding vector, we explore the use of a sequential guiding network that can adapt its guiding vector at every time step. We model the sequential guiding network with a LSTM cell and name it LSTM-g. In this way, the guiding vector G^t which is the hidden state of LSTM-g, can be adjusted based on previous guiding vectors and current guiding input signal z^t . We then use the guiding vector G^t to construct the input signal x^t for the decoding cell \mathcal{R} . The guiding input signal z^t at every time step is formed by concatenating the previous attention vector and attribute features. Figure 1 shows an unrolled version of our framework. We referred the LSTM in the decoding cell as LSTM-d.

3.3.1. High-Level Image Attributes

In addition to the CNN features, our model also integrates other high-level attributes of the input image in the decoding phase. The probability distribution over most frequent words in captions has been shown to be powerful and very informative for image description [4, 9]. We explore the use of this kind of image attributes in our model and denote by $A \in \mathbb{R}^{D_a}$ the detected attribute representations. However, the high-level image attributes could be any additional attribute features connecting the image to its describing sentence. The attributes vector is used to construct initial state for the decoding LSTM and as additional guiding information for the guiding LSTM.

3.3.2. Complete Updating Procedure

Given the input image represented by the set of annotation vectors A and attribute features A and the describing sentence represented by the sequence $[w_0^i, w_1^i, \dots, w_{N^i}^i]$, the decoding cell \mathcal{R} update its hidden state h^t at each time step following the procedure:

$$\begin{aligned} x^{-1} &= W_{ax} A \\ x^t &= W_{wx} w_t^i + W_{gx} G^t, \quad t \in \{0, \dots, N^i - 1\} \\ h^t &= \mathcal{R}(x^t), \quad t \in \{-1, 0, \dots, N^i - 1\} \end{aligned} \quad (7)$$

where $W_{ax} \in \mathbb{R}^{D_x \times D_a}$, $W_{wx} \in \mathbb{R}^{D_x \times D_w}$, $W_{gx} \in \mathbb{R}^{D_x \times D_g}$ are projection matrices of the attribute space, word embedding space and guiding vector space to the LSTM-d input space. D_x, D_w, D_a, D_g denote the dimension of LSTM-d input space, word embedding space, attributes vector and guiding vector, respectively. The vector w_t^i is the distributed representation of the t -th word in caption S^i . We padded each caption with w_0^i to the left and $w_{N^i}^i$ to the right, where w_0^i and $w_{N^i}^i$ represent respectively the distributed representations of start-of-sentence token < sos > and end-of-sentence token < eos >. The sequential guiding network (LSTM-g) update its guiding vector G^t with the following the procedure:

$$\begin{aligned} z^0 &= W_z [h^{-1}; A] \\ z^t &= W_z [\tilde{h}^{t-1}; A], \quad t \in \{1, \dots, N^i - 1\} \\ g^t &= f^{lstm}(z^t), \quad t \in \{0, \dots, N^i - 1\} \end{aligned} \quad (8)$$

where f^{lstm} is the recursive function within a LSTM cell and W_z is a projection matrix.




Fig. 2. Luong’s attention applied on image regions.

4. EXPERIMENTAL STUDIES

4.1. Dataset and Preprocessing

MS COCO Captions [25, 26] is a large scale benchmark dataset for image captioning. We work only with the MS COCO *c5* dataset which annotates each image with 5 human-produced captions. It contains 82,783 images for training and 40,504 images for validation and 40,775 images for testing. We use the splits publicly available¹ in previous works [2, 5]. These splits contain 5000 validation images and 5000 testing images taken from the original validation set. The attribute detectors are trained with the same training set and the 1000 most common words in training captions are used to form the attributes. We follow the same data preprocessing in [2]. That is all captions are transformed to lowercase, non-alphanumeric characters are discarded and all words occurring less than 5 times in the training captions are filtered and replaced by special token `<unk>`. The preprocessing results in a vocabulary of 8791 words.

4.2. Configurations and Implementation

To obtain image annotations \mathcal{A} , we use Inception-V3 [27] vision model. We take the $8 \times 8 \times 2048$ activations map of the last convolutional layer (Mixed_7c in TensorFlow) as annotations. That is \mathcal{A} has dimensions of 64×2048 . To avoid overfitting, we did not train the Inception-V3 from scratch but a model pre-trained on ImageNet is used. We did not fine-tune the weights of the vision model though it could give a small performance boost. The dimension of all layers in the decoding LSTM is set to 1024. For the guiding LSTM, the dimension is set to 512. Both LSTM cells are wrapped with dropout to avoid overfitting. We use word vectors with dimension of 512 and initialize all model parameters (except the CNN parameters) randomly with uniform distribution in $[-0.1, 0.1]$. We built our model based on TensorFlow [28] and used the publicly available code of Google’s *NIC* model as base code.

4.3. Model Comparisons

We validate the effectiveness of our proposed framework by comparing it to several state-of-the-art captioning models based on CNN+RNN architecture such as NIC [1], Soft-Attention [3], LSTM-A5 [4], LTG-Soft-Attention [6] and LTG-Review-Net [6]. We measure the performance of our proposed model with four popular evaluation metrics: BLEU [29], METEOR [30], ROUGE [31] and CIDEr [22]. To compute these metrics, we use the official MS

Image Captioning Models	B@4	MTR	RGE	CDr
Google NIC	20.3	-	-	-
Soft-Attention	24.3	23.9	-	-
LSTM-A5	32.5	25.1	53.8	98.6
LTG-Soft-Attention	32.3	25.9	53.7	102.3
LTG-Review-Net	33.6	26.1	54.8	103.9
SGN+Luong-Attention (ours)	34.0	26.3	55.2	103.6
Luong-Attention (ours)	27.1	24.5	51.7	86.8
ATT+Luong-Attention (ours)	32.8	25.7	54.1	101.9

Table 1. Single model evaluation results on the 5000 testing images. B@4, MTR, RGE, CDr are respectively short for BLEU@4, METEOR, ROUGE-L and CIDEr.

COCO evaluation toolkit² that is made publicly available. Evaluation results across these metrics are shown in Table 1. Our framework is referenced as SGN+Luong-Attention. Note that for all models we report only single model performance knowing that ensemble of multiple models can always give few extra performance points. As shown in the table, our proposed model outperforms all other CNN+RNN architectures. It is worth noting that our model can be seen as an extension of all previous methods and is more complex than most of them. However, the use of a sequential guiding network can be given credit by the fact that our proposed method outperforms the models LTG-Soft-Attention and LTG-Review-Net which are as complex as ours. In contrast to our model, in LTG-Soft-Attention and LTG-Review-Net, the learned guiding vector was fixed and couldn’t be adjusted during words generation.

To further investigate the usefulness of a sequential guiding network, we also compare our SGN+Luong-Attention captioning model with two other baseline that we built: Luong-Attention is a CNN+RNN captioning model based solely on Luong attention mechanism and does not have any extra guiding information at all. ATT+Luong-Attention is a version of SGN+Luong-Attention where we discard the guiding LSTM and replace the guiding vector at each time step by the high-level image attributes. It can be seen from the table that our SGN+Luong-Attention captioning model significantly outperforms ATT+Luong-Attention, again giving credit to the sequential guiding network. It is also worth noting that the Luong-Attention captioning model achieves better performance than the Soft-Attention model in [3], which proves the advantage of using the Luong style of attention mechanism.

5. CONCLUSION

In this paper, we have extended the encoder-decoder framework for image captioning by inserting a guiding network. By modeling the guiding network with a Long Short-Term Memory, the guiding vector can be adjusted at each time step based on the current context and high-level image attributes. We have also explored a natural way of applying Luong’s attention over image regions and demonstrated its effectiveness. We then combined these two strategies in a single joint model. Experiments were conducted on the MS COCO Captions dataset and showed that the proposed model achieves superior performance over existing models based on the CNN+RNN architecture trained using maximum likelihood estimation. In future work, we will consider to apply this model to cross-modal information retrieval.

¹<http://cs.stanford.edu/people/karpathy/deepimagesent/>

²<https://github.com/tylin/coco-caption>

6. REFERENCES

- [1] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan, “Show and tell: Lessons learned from the 2015 mscoco image captioning challenge,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, pp. 652–663, 2017.
- [2] Andrej Karpathy and Li Fei-Fei, “Deep visual-semantic alignments for generating image descriptions,” *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3128–3137, 2015.
- [3] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio, “Show, attend and tell: Neural image caption generation with visual attention,” in *ICML*, 2015.
- [4] Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, and Tao Mei, “Boosting image captioning with attributes,” *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 4904–4912, 2017.
- [5] Xu Jia, Efstratios Gavves, Basura Fernando, and Tinne Tuytelaars, “Guiding long-short term memory for image caption generation,” *CoRR*, vol. abs/1509.04942, 2015.
- [6] Wenhao Jiang, Lin Ma, Xinpeng Chen, Hanwang Zhang, and Wei Liu, “Learning to guide decoding for image captioning,” in *AAAI*, 2018.
- [7] Xinlei Chen and C. Lawrence Zitnick, “Mind’s eye: A recurrent visual representation for image caption generation,” *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2422–2431, 2015.
- [8] Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel, “Unifying visual-semantic embeddings with multimodal neural language models,” *CoRR*, vol. abs/1411.2539, 2014.
- [9] Hao Fang, Saurabh Gupta, Forrest N. Iandola, Rupesh Kumar Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C. Platt, C. Lawrence Zitnick, and Geoffrey Zweig, “From captions to visual concepts and back,” in *CVPR*, 2015.
- [10] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, and Alan L. Yuille, “Deep captioning with multimodal recurrent neural networks (m-rnn),” *CoRR*, vol. abs/1412.6632, 2014.
- [11] Jeff 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,” *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2625–2634, 2015.
- [12] Chang Liu, Fuchun Sun, Changhu Wang, Feng Wang, and Alan L. Yuille, “Mat: A multimodal attentive translator for image captioning,” in *IJCAI*, 2017.
- [13] Long Chen, Hanwang Zhang, Jun Xiao, Liqiang Nie, Jian Shao, and Tat-Seng Chua, “Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning,” *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6298–6306, 2017.
- [14] Raffaella Bernardi, Ruken Cakici, Desmond Elliott, Aykut Erdem, Erkut Erdem, Nazli Ikizler-Cinbis, Frank Keller, Adrian Muscat, and Barbara Plank, “Automatic description generation from images: A survey of models, datasets, and evaluation measures,” in *IJCAI*, 2017.
- [15] Kyunghyun Cho, Bart van Merriënboer, aglar Gülehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio, “Learning phrase representations using rnn encoder-decoder for statistical machine translation,” in *EMNLP*, 2014.
- [16] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “Neural machine translation by jointly learning to align and translate,” *CoRR*, vol. abs/1409.0473, 2014.
- [17] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le, “Sequence to sequence learning with neural networks,” in *NIPS*, 2014.
- [18] Thang Luong, Hieu Pham, and Christopher D. Manning, “Effective approaches to attention-based neural machine translation,” in *EMNLP*, 2015.
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems 25*, pp. 1097–1105, 2012.
- [20] Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel, “Self-critical sequence training for image captioning,” *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1179–1195, 2017.
- [21] Siqi Liu, Zhenhai Zhu, Ning Ye, Sergio Guadarrama, and Kevin Murphy, “Improved image captioning via policy gradient optimization of spider,” *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 873–881, 2017.
- [22] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh, “Cider: Consensus-based image description evaluation,” *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4566–4575, 2015.
- [23] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould, “Spice: Semantic propositional image caption evaluation,” in *ECCV*, 2016.
- [24] Hochreiter Sepp and Schmidhuber Jürgen, “Long short-term memory,” *Neural Comput.*, pp. 1735–1780, 1997.
- [25] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick, “Microsoft coco: Common objects in context,” in *ECCV*, 2014.
- [26] 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,” *CoRR*, vol. abs/1504.00325, 2015.
- [27] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna, “Rethinking the inception architecture for computer vision,” *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2818–2826, 2016.
- [28] Martín Abadi, Ashish Agarwal, and et al., “Tensorflow: Large-scale machine learning on heterogeneous distributed systems,” *CoRR*, vol. abs/1603.04467, 2015.
- [29] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *ACL*, 2002.
- [30] Satanjeev Banerjee and Alon Lavie, “Meteor: An automatic metric for mt evaluation with improved correlation with human judgments,” in *ACL*, 2005, pp. 65–72.
- [31] Chin-Yew Lin, “Rouge: a package for automatic evaluation of summaries,” July 2004.