

MAT: A Multimodal Attentive Translator for Image Captioning

Chang Liu^{1,4}, Fuchun Sun¹, Changhu Wang², Feng Wang^{3,4}, Alan Yuille⁴

¹Department of Computer Science, Tsinghua University

²Toutiao AI Lab, ³Department of Electronic Engineering, UESTC

⁴Cognitive Science & Computer Science, Johns Hopkins University

¹{cliu13@mails, fcsun@mail}.tsinghua.edu.cn,

²wangchanghu@toutiao.com, ³feng.wff@gmail.com, ⁴alan.yuille@jhu.edu

Abstract

In this work we formulate the problem of image captioning as a multimodal translation task. Analogous to machine translation, we present a sequence-to-sequence recurrent neural networks (RNN) model for image caption generation. Different from most existing work where the whole image is represented by convolutional neural network (CNN) feature, we propose to represent the input image as a sequence of detected objects which feeds as the source sequence of the RNN model. In this way, the sequential representation of an image can be naturally translated to a sequence of words, as the target sequence of the RNN model. To represent the image in a sequential way, we extract the objects features in the image and arrange them in a order using convolutional neural networks. To further leverage the visual information from the encoded objects, a sequential attention layer is introduced to selectively attend to the objects that are related to generate corresponding words in the sentences. Extensive experiments are conducted to validate the proposed approach on popular benchmark dataset, i.e., MS COCO, and the proposed model surpasses the state-of-the-art methods in all metrics following the dataset splits of previous work. The proposed approach is also evaluated by the evaluation server of MS COCO captioning challenge, and achieves very competitive results, e.g., a CIDEr of 1.029 (c5) and 1.064 (c40).

1 Introduction

Image captioning is a challenging problem. Unlike other computer vision tasks such as image classification and object detection, image captioning requires not only understanding the image, but also knowledge of natural language. Early methods for image captioning either explored template-based, e.g., [Mitchell and et al., 2012; Elliott and Keller, 2013] or retrieval-based approaches [Gong *et al.*, 2014; Kuznetsova *et al.*, 2014]. However, the language models were usually heavily hand-designed, and found it hard to generate novel sentences with new compositions.

Inspired by the success of sequence-to-sequence machine translation [Sutskever *et al.*, 2014], based on recurrent neural networks, recent approaches for image captioning brought new insights by using a two-stage ‘encoding’ and ‘decoding’ technique [Donahue *et al.*, 2015; Karpathy and Fei-Fei, 2015; Vinyals *et al.*, 2015b]. The common idea of these approaches is to use the whole CNN feature of the image as the ‘source’ input, to replace the words of the ‘source language’ in the translation task. The caption is then generated by conditioning the output words on the CNN feature of the image. One problem these approaches suffer from is the imbalance of the encoded visual information (representation of the image) and the language part (representation of words), because in an RNN the image in the ‘source language’ only provides the CNN feature of image at one time step only, i.e., the initial step, but the words in the ‘target language’ contribute at multiple time steps. Since the power of RNN lies in its capability to model the contextual information between time steps [Hochreiter and Schmidhuber, 1997], such image representation weakens the RNN’s memory of the visual information as it contains no temporal concept.

To encode more visual information, some approaches, e.g., [Mitchell and et al., 2012; Yang *et al.*, 2011] represent the input image in terms of the objects it contains. One strategy that some early work leverages is to construct inputs for the language model using object-attribute pairs, with objects and attributes represented by descriptive words, e.g., [Mitchell and et al., 2012; Yang *et al.*, 2011], in which, however, the visual features of objects are actually not utilized. Another strategy some recent methods such as [Fang *et al.*, 2015; Wu *et al.*, 2016; You *et al.*, 2016] proposed is to leverage visual descriptors of image parts to enhance the visual information using weakly supervised approach. More recently, some approaches leverage the attention mechanism to encode more visual information [Xu *et al.*, 2015; You *et al.*, 2016], in which the models learn to fix their gaze on salient regions when generating the corresponding words.

In this work, we follow the basic strategy of CNN and RNN captioning model, but expand the image representation from one single CNN feature at a single time step in RNN, to a sequence of objects at multiple time steps. In this way, analogous to machine translation, image captioning is formulated as translating a visual language (a sequence of objects) to a natural language (a sequence of words). This idea has several advantages. Firstly, we can leverage object detection

techniques to encode more visual information in the visual language; secondly, the sequential representation of the image accords more with the temporal concept of RNN models, and makes the two sides of translation more balanced. On the basis of the multi-modal sequence translation model, we also introduce a ‘sequential attention’ mechanism, where our model can learn to distribute its attention to different objects in the image. It should be noted that the attention mechanism in our model differs from previous ones because the attention in our model works on the basis of sequential representation of the image, by taking all encoding objects information into account when generating words; while the attention in previous work focuses on the feature map of the whole image representation, and sequential representation of the image is not explored.

To summarize, the main contributions of this work are:

- We present a novel multimodal translation model for image captioning, which translates the ‘visual language’ of a sequence of objects to a natural language with a sequence of words. This sequential formulation ensures the balance between the encoding information and the decoding information of the RNN.

- We introduce a ‘sequential attention’ layer, which learns to consider all the encoder information at each time steps during decoding with different weights.

- We quantitatively validate our model on benchmark dataset, and surpass the state-of-the-art methods in all metrics. For example, on MS COCO Captioning Challenge evaluation server, our model achieves a CIDEr of 1.029 (c5) and 1.064(c40), while the second best achieves 0.965 (c5) and 0.969 (c40). In addition, we also evaluate our method using the more recent metric SPICE [Anderson *et al.*, 2016], which accords more with human judgments, and our model achieves 18.9 (c5).

2 Related Work

CNN+RNN based captioning. A typical captioning way is to combine CNN and RNN, where CNN is used to extract the feature of the whole image, and RNN to construct the language model. For example, Vinyals *et al.* [Vinyals *et al.*, 2015b] proposed an end-to-end model composed of a CNN and an RNN. The model is trained to maximize the likelihood of the target sentence given the CNN feature of the training image at the initial time step. Mao *et al.* [Mao *et al.*, 2015] presented an m-RNN model, where the CNN feature of the image is fed into the multimodal layer after the recurrent layer rather than the initial time step. Similar work that utilizes CNN and RNN to generate descriptions includes [Chen and Lawrence Zitnick, 2015; Karpathy and Fei-Fei, 2015]. However, most of above methods represent the image in a static form, such as a 4096-d CNN feature vector. Although this feature can well represent an image, it is insufficient for the sequential RNN model. That is because such a feature only provides the encoding phase of the RNN model with a single time-step data, leaving the rest of the model to the decoding phase where words in the caption are used.

Sequence based captioning. Several algorithms extend the image representation from a single time-step data to multiple time-step data in image captioning [Donahue *et al.*, 2015]

and video description [Donahue *et al.*, 2015; Venugopalan *et al.*, 2015]. In image captioning, [Donahue *et al.*, 2015] proposed to feed the image feature to the RNN at each time step; and in video description, Venugopalan *et al.* first computed the CNN feature for each video frame, and then fed the mean pooling of these features to the RNN at every time step in the encoding phase. Note that in these algorithms, the inputs to RNN during encoding are the same at each time step; they are the duplicate features of the image [Donahue *et al.*, 2015]. Feeding the same information at each time step cannot make the networks to obtain more contextual information during encoding, but only strengthens the same concept over and over. On the other hand, instead of extending the representation with the same inputs, [Donahue *et al.*, 2015] and [Venugopalan *et al.*, 2015] proposed to extract the CNN feature for each video frame, and feed them to the RNN model one frame per time step, to avoid from using repeated feature inputs. Our work is different from above two approaches, for we focus on image captioning with object-level representation, rather than a video description with frame-level representation.

Attention based captioning. Visual attention has been proved as an effective way for the task of image captioning [Xu *et al.*, 2015; Yang *et al.*, 2016; You *et al.*, 2016; Liu *et al.*, 2016]. These attention based captioning models are capable of learning where to attend in the image when generating the target words. They either learn the distribution of spatial attention from the last convolutional layer of the convolutional neural network [Xu *et al.*, 2015; Yang *et al.*, 2016], or learn the distribution of semantic attention from the visual attributes that are learnt from weakly annotated images from social media [You *et al.*, 2016]. While these methods show the effectiveness of the attention mechanism, the contextual information in the encoding sequence is not explored. Our attention layer differs as it takes a sequential form, where each hidden state of the encoding stage contributes to generating the decoding words.

3 Proposed Method

In the following sections, we first formulate the problem as a multimodal translation problem 3.1. We then present model in detail in Section 3.2. After which we elaborate the attention layer in Section 3.3. And Section 3.4 states the training and inferences details of the model.

3.1 Formulation

In previous work with CNN+RNN solutions, the core idea is usually to maximize the probability of the description given the input image:

$$\log p(S|I) = \sum_{t=1}^N \log p(S_t|I, S_{1:t-1}) \quad (1)$$

where I represents the image, S_i is the i th word in sentence S , and $p(S_t|I, S_{1:t-1})$ is the probability of generating word S_t given the image and previous words $S_{1:t-1}$. A common representation of the image is a CNN feature vector, and the recursive language part is usually modeled with recurrent neural networks, where an RNN unit considers the following two

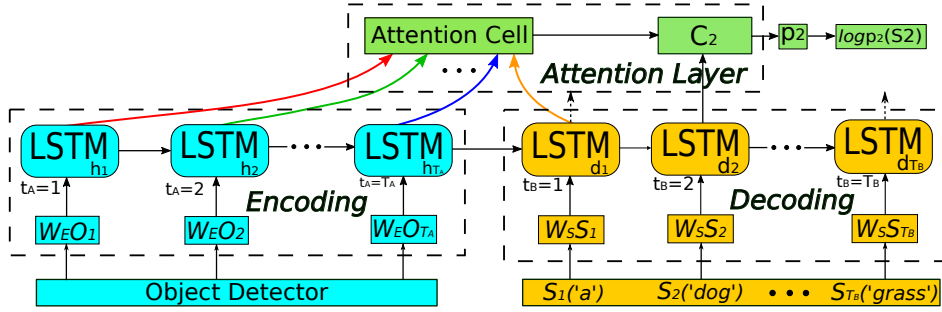


Figure 1: Overview of the Model. The source sequence is represented with the embedding of object in a hidden space; the target sequence is represented with the embedding of the words in the same space. The mapping from the source to the target is modeled with two LSTM units in an unrolled version. An attention layer computes the context vector C_t from all the encoding hidden states and decoding hidden state from last timestep. The target word is generated by a softmax over the vocabulary given the attention context vector. Better view in color.

data as inputs: (1) input at current time step t , and (2) output from the previous time step $t - 1$.

In this work, we formulate image captioning as a multimodal translation task, i.e., translating a visual language (a sequence of objects) to a natural language (a sequence of words). The core idea is to use RNN to model the translation process, by feeding one object at a time to an RNN unit during encoding, and one word at a time during decoding. Besides this multimodal translation scheme, we also introduce an attention context vector C_t to generate the word during decoding, which comes from the proposed sequential attention layer. The attention layer shares similar spirit from the literature of machine translation [Vinyals *et al.*, 2015a; Luong *et al.*, ; Bahdanau *et al.*, 2015], where the core motivation is to consider all the hidden states of the encoding when computing the attention vector. Thus, we call our model Multimodal Attentive Translator, MAT for short.

To be specific, given an image I , we use $seq(I)$ to denote its sequential representation, which contains a sequence of representations $seq(I) = \{O_1, O_2, \dots, O_{T_A}\}$, where O_1 to O_{T_A-1} are the object representations and the last item O_{T_A} encodes the global environmental information by feeding CNN feature of the whole image. Then the RNN takes in $seq(I)$ by encoding each object into a fixed length vector, and is recursively activated at each time step. We denote the attention vector as C_{t_B} at decoding time step t_B , where t_B ranges from 1 to T_B , i.e., the decoding length. The sentence is generated by conditioning the outputs given attention context vector C_{t_B} (Eqn. 2), where C_{t_B} is computed from the attention layer (Eqn. 3).

$$\log p(S|seq(I)) = \sum_{t_B=1}^{T_B} \log p(S_{t_B}|C_{t_B}, S_1, \dots, S_{t_B-1}) \quad (2)$$

$$C_{t_B} = ATT(H, d_{t_B-1}), \quad (3)$$

where H denotes all the encoding hidden states h_1, h_2, \dots, h_{T_A} , and d_{t_B-1} denotes the decoding hidden state of last time step. And we have encoding states and decoding states computed from:

$$h_{t_A} = RNN_{en}(seq(I)_{t_A}, h_{t_A-1}), t_A = 1, 2, \dots, T_A \quad (4)$$

$$d_{t_B} = RNN_{de}(S_{t_B}, d_{t_B-1}), t_B = 1, 2, \dots, T_B \quad (5)$$

3.2 Seq2Seq Multimodal Translator

Source sequence representation. The source sequence is represented by the objects from the image. We first use object detectors to locate the objects and extract D_o -dimension CNN features, denoted as $CNN(O_{t_A})$ for the t_A -th object O_{t_A} . Then object features are mapped into an H -dimension hidden space with embedding matrix W_E , with $\mathbb{R}^{H \times D_o}$ dimension. The source sequence of the RNN is represented as:

$$x_{t_A} = W_E CNN(O_{t_A}), \quad t_A \in \{1, 2, \dots, T_A\}, \quad (6)$$

where t_A is the encoding timestep of the network, and T_A is the total length of encoding state.

In general, any object detection method can be leveraged to generate the previous mentioned object features. The feature extractor is shown in Figure 2, and it can be inserted to any RPN based detection network. In practical, we adopt the R-FCN [Jifeng *et al.*, 2016] detection network, where we pool the object features from the last convolutional layer using *roi pooling*. Details refer to 3.4.

Target sequence representation. The target sequence is represented with a set of words S_1, \dots, S_N in the sentence S . Each word S_{t_B} is represented as a ‘one-hot’ vector, with a dimension \mathbb{R}^{D_s} equal to the vocabulary size. We add special START word and END word to denote the start and the end of the sentence. The start word and end word are also included into the vocabulary of the model. Then the words are mapped to the same hidden space with word embedding matrix W_S , with a dimension of $\mathbb{R}^{H \times D_s}$. Thus the target sequence is represented by:

$$x_{t_B} = W_S S_{t_B}, \quad t_B \in \{1, 2, \dots, T_B\}, \quad (7)$$

where t_B is the decoding timestep of the network, and T_B is the total length of decoding stage.

RNN translation from source to target. To model the translation from the source sequence to the target sequence, we leverage the long-short term memory (LSTM), a specific unit of RNN, to avoid the gradient exploding and vanishing problem of the network [Hochreiter and Schmidhuber, 1997; Hochreiter *et al.*, 2001].




Figure 2: Object feature extractor of the model. We extract object features from the last convolutional layer of detection network using roi pooling. The pooled features are selected by the detection score (i.e., cls) from highest to lowest.

LSTM takes in the output of the previous time step, as well as the input at the current time step, as the inputs of the current unit. To better illustrate the idea of this recursive process, we unroll the encoding $LSTM_1$ and decoding $LSTM_2$ along the time dimension, by copying the LSTM unit at each time step, as shown in Fig. 1. The core of an LSTM unit is a memory cell c , which is controlled by several gates. The activation of each gate determines whether the corresponding input is accepted or rejected. Thus the mapping from the source (Eqn. 6) to the target (Eqn. 7) is formulated by the following equations:

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i), \quad (8)$$

$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f), \quad (9)$$

$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o), \quad (10)$$

$$g_t = \sigma(W_{xg}x_t + W_{hg}h_{t-1} + b_g), \quad (11)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot g_t, \quad (12)$$

$$h_t = o_t \odot \phi(c_t), \quad (13)$$

where t ranges from the start of the source sequence to the end of the target sequence in all equations; i_t , f_t , o_t , and g_t represent the input gate, forget gate, output gate, and input modulation gate at time step t , respectively; c_t and h_t are the memory cell and the hidden state; W_{ij} represents the connection matrix and b_j is the bias; σ is the sigmoid non-linearity operator, and ϕ is the hyperbolic tangent non-linearity; \odot is the element-wise multiplication operator.

3.3 Sequential Attention Layer

Although the decoding hidden state of LSTM unit comes from previous encoding hidden state, the states that are far from the current decoding state may contribute little to generate the word. In other words, objects that appear at the very beginning of the source sequence may be well related to the word at the decoding step. For example, a dog detected with highest score will appear at the first time step of the encoding procedure, but the word "dog" may be at the end of the decoding time step, if the sentence is "A man is playing with a dog." Driven by the idea, we design an attention context vector, which lets the model generate words considering all hidden encoding states, computed from the proposed attention layer. The attention layer thus takes inputs from hidden states H ,




Figure 3: Attention Layer. At decoding timestep $t = t_B$, the attention layer computes the attention context vector C_{t+1} from all the encoding hidden states h_1, \dots, h_{T_A} , as well as previous decoding state d_t . Better view in color.

as well as previous decoding state d_{t_B} , and outputs the attention context vector, as shown in Figure. 3. Specifically, the distribution p_{t_B+1} is computed from the context vector C_{t+1} , which derives as (for all the equations we omit bias term for simplicity):

$$p_{t_B+1} = \text{Softmax}(W_C(out_{t_B+1})) \quad (14)$$

$$out_{t_B+1} = \text{Concat}(d_{t_B+1}, d_{t_B}') \quad (15)$$

$$d_{t_B}' = \sum_{i=1}^{T_A} a_i^{t_B} h_i \quad (16)$$

$$a_i^{t_B} = \text{Softmax}(u_{t_B})_i \quad (17)$$

$$u_{t_B} = V^T \tanh(W_H h_i + W_D d_{t_B}) \quad (18)$$

where V , W_H , W_D and W_C are the parameters of the layer. $V \in \mathbb{R}^{H \times 1}$ dimension vector; $W_H, W_D \in \mathbb{R}^{H \times H}$; $W_C \in \mathbb{R}^{2H \times H}$. The vector u_{t_B} is of length T_A , which represents how much attention the model pays to each encoding hidden state h_{t_A} when generating the attention context vector. u_{t_B} is then used to compute a new state vector, d_{t_B}' . We then concatenate d_{t_B}' and current decoding hidden state d_{t_B+1} , with which we feed as the hidden state to make prediction over the vocabulary using softmax to generate the word.

3.4 Training and Inference

Training. To obtain the sequential representation of the image, we run object detectors on the image. Specifically, we leverage R-FCN[Jifeng *et al.*, 2016] trained on MS COCO dataset, using the Resnet101 convolutional architecture[He *et al.*, 2016]. We use a threshold of 0.5 to detect the objects in the image, and the objects are ordered according to their detection scores. We use two different LSTMs for encoding and decoding. The hidden state size is set to 512. To cope with variable length of both the source sequence and the target sequence for batch training, we leverage a bucket and padding method, where the sequences are split to different buckets and zero padded to bucket length according to the length of the source sequence as well as the length of the target sequence. Specifically, in training, we use four buckets, i.e., $\{(2, 10), (4, 15), (6, 20), (8, 30)\}$. For example, given a training sample with 5 objects and 10 words, it belongs to the third bucket (6,20) since although $10 \text{ words} < 15 \text{ threshold}$, the object number 5 exceeds threshold 4, thus the second bucket doesn't

Methods	MS COCO, 5000 testing images							
	CIDEr	METEOR	RougeL	BLEU-1	BLEU-2	BLEU-3	BLEU-4	SPICE
Berkeley LRCN [Donahue <i>et al.</i> , 2015]	-	-	-	0.628	0.442	0.304	0.210	-
DeepVS [Karpathy and Fei-Fei, 2015]	0.660	0.195	-	0.625	0.450	0.321	0.230	-
Attention [Xu <i>et al.</i> , 2015]	-	0.230	-	0.718	0.504	0.357	0.250	-
Reviewnet [Yang <i>et al.</i> , 2016]	0.886	0.237	-	-	-	-	0.290	-
Semantic [You <i>et al.</i> , 2016]	-	0.243	-	0.709	0.537	0.402	0.304	-
MAT(ours)	1.058	0.258	0.541	0.731	0.567	0.429	0.323	18.9

Table 1: Comparison results on MSCOCO on 5000 testing images following previous work. The highest score is labeled as bold. Our model is noted as ‘MAT(ours)’ in gray. All metrics are reported using c5 references. Note SPICE metric is recently proposed, and lack refereed methods to compare on.

Methods	MS COCO Captioning Challenge, 40775 images (c5)						
	CIDEr	METEOR	RougeL	B-1	BLEU-2	BLEU-3	BLEU-4
Google NIC [Vinyals <i>et al.</i> , 2015b]	0.943	0.254	0.530	0.713	0.542	0.407	0.309
MSR Captivator [Devlin <i>et al.</i> , 2015]	0.931	0.248	0.526	0.715	0.543	0.407	0.308
Berkeley LRCN [Donahue <i>et al.</i> , 2015]	0.921	0.247	0.528	0.718	0.548	0.409	0.306
m-RNN [Mao <i>et al.</i> , 2015]	0.917	0.242	0.521	0.716	0.545	0.404	0.299
MSR [Fang <i>et al.</i> , 2015]	0.912	0.247	0.519	0.695	0.526	0.391	0.291
ACVT [Wu <i>et al.</i> , 2016]	0.911	0.246	0.528	0.725	0.556	0.414	0.306
Attention [Xu <i>et al.</i> , 2015]	0.865	0.241	0.516	0.705	0.528	0.383	0.277
Reviewnet [Yang <i>et al.</i> , 2016]	0.965	0.256	0.533	0.720	0.550	0.414	0.313
Semantic [You <i>et al.</i> , 2016]	0.943	0.250	0.530	0.713	0.542	0.407	0.309
MAT(ours)	1.029	0.258	0.540	0.734	0.568	0.427	0.320

Table 2: Comparison results on MS COCO Caption Challenge 2015, using MS COCO evaluation server. All metrics are reported using c5 references.

fit. We then zero padding the words from 10 to 15, and the objects from 5 to 6. Our loss is the sum of the negative log likelihood of the generated words at decoding time steps, i.e.:

$$Loss = - \sum_{t_B=1}^{T_B} \log p_{t_B}(S_{t_B}) \quad (19)$$

We use SGD with batch size of 64 to train the network. The learning rate is set to 0.1, and halved when training loss stops to decrease. To avoid overfitting, we leverage dropout at 0.5 for all layers, and early stops the training on validation split with 5000 images. On a Titan X Maxwell computer, the training process takes about 12 hours.

Inference. We first run object detectors on the testing image, and the feature representations of the objects are fed in to the encoding parts of the model. We use BeamSearch of size 20, which considers iteratively the best b candidates to generate next word. The sentence generation stops when it generates the special END word.

4 Experimental Results

4.1 Datasets and Evaluation Measurements

MSCOCO [Lin *et al.*, 2014] contains 82,783 training, 40,504 validation and 40,775 testing images, which are withheld in MS COCO server. To compare with previous methods, we follow the split from previous work [Karpathy and Fei-Fei, 2015; Xu *et al.*, 2015], i.e., we use 5000 images for validation and 5000 images for testing from the 40504 validation set. Moreover, we also compare with the state-of-the-art methods listed in the leader board on the MSCOCO website,

where 40,755 images are withheld for testing (ground truth sentences unavailable). We follow previous work to generate the vocabulary of the model, that is, we first count all the occurrence number of all the words and filter the words which occurrence number is less than 5. The final vocabulary size is 8791.

Evaluation measurements: We use the public available MS COCO evaluation toolkit¹ to evaluate our model, which computes BLEU [Papineni *et al.*, 2002], CIDEr [Vedantam *et al.*, 2015], METEOR [Banerjee and Lavie, 2005], and RougeL [Lin, 2004]. In addition, we evaluate our model with the recent metric SPICE [Anderson *et al.*, 2016], which accords more with human judgments. For all the metrics, the higher the better.

4.2 Overall Comparison with the State of the Arts

MSCOCO: The results are shown in Table 1 and Table 2, which show the performance of our model compared to state-of-the-art published methods. From Table 1 and Table 2 it can be seen that our model outperforms state-of-the-art in all evaluation measurements, which implies the validity and effectiveness of proposed model.

In consideration of the latent possibility that the metrics used by MS COCO evaluation server may drift away from human judgments, although hiring human to evaluate and compare the generated sentences is a more reliable and solid way, it’s very money and time consuming and thus is not always practical. Nonetheless, in addition to validate on MS COCO evaluation server, we also evaluate our model using the re-

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

Methods	MS COCO, 5000 testing images							
	CIDEr	METEOR	RougeL	BLEU-1	BLEU-2	BLEU-3	BLEU-4	SPICE
Baseline 1	0.927	0.230	0.516	0.705	0.504	0.357	0.250	17.2
Baseline 2	0.961	0.245	0.519	0.711	0.542	0.403	0.298	17.8
MAT(ours)	1.058	0.258	0.541	0.731	0.567	0.429	0.323	18.9

Table 3: Comparison between MAT and baseline methods. Baseline 1 is CNN+RNN method. Baseline 2 is Seq2Seq method, i.e., using MAT model but without attention layer.

cent proposed metric SPICE [Anderson *et al.*, 2016], which is more consistent with human judgments and gives a better idea of the performance of language generating models. We achieve a SPICE of 18.9(c5), using public available SPICE evaluation tool². Given no contemporary work reports SPICE on MS COCO 5000 testing split, and performances reported in the SPICE paper [Anderson *et al.*, 2016] use C40 testing set obtained from MS COCO organizers, which is inaccessible to us and thus makes our C5 performance incomparable, we decide, nonetheless, to publish our C5 SPICE performance on 5000 testing split, in order to provide a comparable number for future work.

4.3 Comparison with Baseline Methods

We compare our MAT model with two baselines: 1) vanilla CNN+RNN captioning model, where single CNN image feature is used; 2) Seq2Seq model without attention layer, i.e. take out attention layer from MAT model. The experiment settings, i.e., training and inference details for baseline 1 and baseline 2 are consistent with that of MAT model. The results are shown in Table 3. Two conclusions can be drawn from the Table. Firstly, from the comparison between baseline 1 and baseline 2, it shows that the sequential representation of the image indeed improves the performance compared to the single CNN feature of the representation, which accords with idea of encoding more visual information of the source sequence of the RNN. Secondly, when comparing MAT results to both baseline 1 and baseline 2, it can be drawn that the proposed sequential attention layer further improves the performance by a large margin. These two facts indicate the effectiveness of our multimodal translator with sequential attention layer.

4.4 Qualitatively Evaluation

We present the captioning results of the proposed method in Figure 4. The results are randomly selected from the unused MS COCO testing 5000 split. We show four sentences for each image: GT shows the ground truth sentence; B1 shows the baseline 1 method, i.e., CNN+RNN model; B2 shows the baseline 2 method, i.e., Seq2Seq model (i.e. MAT w/o attention); and finally, MAT is our translation model with sequential attention layer. From the figure we can see some interesting results. Take the bottom picture for example. The baseline model 1 mistake the two boats for one boat, which MAT model predicts correctly. Moreover, the MAT model correctly predicts actions, e.g., in the middle picture, MAT predicts the action ‘pitching’, rather than ‘swing’ or ‘stand’.

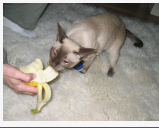




	GT: A cat that is eating some kind of banana. B1: A couple of cats laying on top of a bed. B2: A cat laying next to a pair of scissors. MAT: A close up of a cat eating a banana.
	GT: A large white dog is sitting on a bench beside an elderly man. B1: A dog standing next to a person on a skateboard. B2: A couple of dogs sitting on a bench. MAT: A man sitting on a bench with a dog.
	GT: A baseball player preparing to throw a pitch during a game. B1: A baseball player swinging a bat at a ball. B2: A baseball player standing on top of a field. MAT: A baseball player pitching a ball on a field.
	GT: A man on a bicycle riding next to a train. B1: A man standing in front of a train station. B2: A person riding a bike next to a train. MAT: A man riding a bike next to a train.
	GT: Two kayaks one pink the other yellow on bank of water. B1: A yellow surfboard sitting on top of a sandy beach. B2: A couple of boats parked next to each other. MAT: Two boats are docked in a body of water.

Figure 4: Random selected caption generation results on MS COCO testing 5000 split. GT stands for ground truth sentence. B1, i.e., baseline 1, and B2, i.e., baseline 2 shows CNN+RNN and Seq2Seq baseline methods respectively. MAT shows our final generation results.

5 Conclusion

In this work, we focus on improving the performance of image captioning by introducing two novel ideas. The first is to formulate the task as a multimodal sequence-to-sequence translation task, where the source language is represented as a sequence of objects detected from the image. The second is to introduce a sequential attention layer, which take all encoding hidden states in to consideration when generating each word. The proposed model shows superior performance over state of the art methods, and is quantitatively and qualitatively evaluated to show the effectiveness of the method.

Acknowledgments

This paper is jointly supported by National Natural Science Foundation of China under with Grant No.61621136008, 61327809,61210013,91420302 & 91520201. This paper is also supported by NSF STC award CCF-1231216.

²<https://github.com/peteanderson80/coco-caption.git>

References

- [Anderson *et al.*, 2016] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *ECCV*, 2016.
- [Bahdanau *et al.*, 2015] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *ICLR*, 2015.
- [Banerjee and Lavie, 2005] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop*, 2005.
- [Chen and Lawrence Zitnick, 2015] Xinlei Chen and C. Lawrence Zitnick. Mind’s eye: A recurrent visual representation for image caption generation. In *CVPR*, 2015.
- [Devlin *et al.*, 2015] Jacob Devlin, Hao Cheng, Hao Fang, and et al. Language models for image captioning: The quirks and what works. In *ACL*, 2015.
- [Donahue *et al.*, 2015] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.
- [Elliott and Keller, 2013] Desmond Elliott and Frank Keller. Image description using visual dependency representations. In *EMNLP*, 2013.
- [Fang *et al.*, 2015] Hao Fang, Saurabh Gupta, Forrest Iandola, and et al. From captions to visual concepts and back. In *CVPR*, 2015.
- [Gong *et al.*, 2014] Yunchao Gong, Liwei Wang, Micah Hodosh, Julia Hockenmaier, and Svetlana Lazebnik. Improving image-sentence embeddings using large weakly annotated photo collections. In *ECCV*, 2014.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997.
- [Hochreiter *et al.*, 2001] Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jürgen Schmidhuber. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001.
- [Jifeng *et al.*, 2016] Dai Jifeng, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *NIPS*, 2016.
- [Karpathy and Fei-Fei, 2015] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.
- [Kuznetsova *et al.*, 2014] Polina Kuznetsova, Vicente Ordonez, Tamara L Berg, and Yejin Choi. Treetalk: Composition and compression of trees for image descriptions. *ACL*, 2014.
- [Lin *et al.*, 2014] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*. 2014.
- [Lin, 2004] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL-04 workshop*, volume 8. Barcelona, Spain, 2004.
- [Liu *et al.*, 2016] Chenxi Liu, Junhua Mao, Fei Sha, and Alan Yuille. Attention correctness in neural image captioning. *AAAI*, 2016.
- [Luong *et al.*,] Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- [Mao *et al.*, 2015] Junhua Mao, Xu Wei, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan L. Yuille. Learning like a child: Fast novel visual concept learning from sentence descriptions of images. In *ICCV*, 2015.
- [Mitchell and et al., 2012] Margaret Mitchell and et al. Midge: Generating image descriptions from computer vision detections. In *EACL*, 2012.
- [Papineni *et al.*, 2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002.
- [Sutskever *et al.*, 2014] Ilya Sutskever, Oriol Vinyals, and Quoc VV Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014.
- [Vedantam *et al.*, 2015] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015.
- [Venugopalan *et al.*, 2015] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. Sequence to sequence - video to text. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [Vinyals *et al.*, 2015a] Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. Grammar as a foreign language. In *NIPS*, 2015.
- [Vinyals *et al.*, 2015b] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015.
- [Wu *et al.*, 2016] Qi Wu, Chunhua Shen, Lingqiao Liu, Anthony Dick, and Anton van den Hengel. What value do explicit high level concepts have in vision to language problems? In *CVPR*, 2016.
- [Xu *et al.*, 2015] Kelvin Xu, Jimmy Ba, Ryan Kiros, and et al. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- [Yang *et al.*, 2011] Yezhou Yang, Ching Lik Teo, Hal Daumé III, and Yiannis Aloimonos. Corpus-guided sentence generation of natural images. In *EMNLP*, 2011.
- [Yang *et al.*, 2016] Zhilin Yang, Ye Yuan, Yuexin Wu, William W Cohen, and Ruslan R Salakhutdinov. Review networks for caption generation. In *NIPS*, 2016.
- [You *et al.*, 2016] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *CVPR*, 2016.