

# Convolutional Auto-encoding of Sentence Topics for Image Paragraph Generation\*

Jing Wang<sup>1</sup>, Yingwei Pan<sup>2</sup>, Ting Yao<sup>2</sup>, Jinhui Tang<sup>1†</sup> and Tao Mei<sup>2</sup>

<sup>1</sup>School of Computer Science and Engineering, Nanjing University of Science and Technology, China

<sup>2</sup>JD AI Research, Beijing, China

{jwang, jinhuitang}@njust.edu.cn, {panyw.ustc, tingyao.ustc}@gmail.com, tmei@live.com

## Abstract

Image paragraph generation is the task of producing a coherent story (usually a paragraph) that describes the visual content of an image. The problem nevertheless is not trivial especially when there are multiple descriptive and diverse gists to be considered for paragraph generation, which often happens in real images. A valid question is how to encapsulate such gists/topics that are worthy of mention from an image, and then describe the image from one topic to another but holistically with a coherent structure. In this paper, we present a new design — Convolutional Auto-Encoding (CAE) that purely employs convolutional and deconvolutional auto-encoding framework for topic modeling on the region-level features of an image. Furthermore, we propose an architecture, namely CAE plus Long Short-Term Memory (dubbed as CAE-LSTM), that novelly integrates the learnt topics in support of paragraph generation. Technically, CAE-LSTM capitalizes on a two-level LSTM-based paragraph generation framework with attention mechanism. The paragraph-level LSTM captures the inter-sentence dependency in a paragraph, while sentence-level LSTM is to generate one sentence which is conditioned on each learnt topic. Extensive experiments are conducted on Stanford image paragraph dataset, and superior results are reported when comparing to state-of-the-art approaches. More remarkably, CAE-LSTM increases CIDEr performance from 20.93% to 25.15%.

## 1 Introduction

The recent advances in Convolutional Neural Networks (CNN) have successfully pushed the limits and improved the state-of-the-art technologies of image understanding. As such, it has become achievable to recognize an image with a pre-defined set of categories. In a further step to describe an image in a natural-language utterance, image captioning [Vinyals *et al.*, 2015; Xu *et al.*, 2015; Rennie *et al.*, 2017]

\*This work was performed at JD AI Research.

†Corresponding Author.

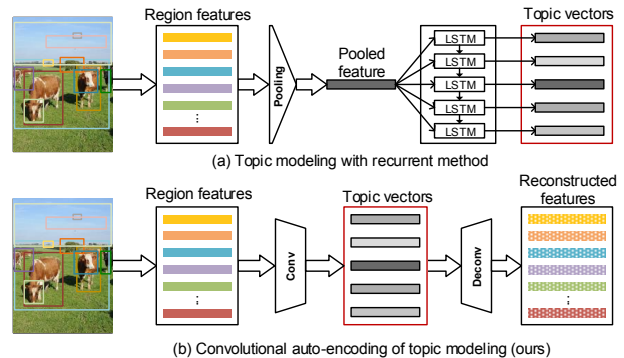


Figure 1: Modeling the topics from an image by employing (a) recurrent method and (b) convolutional auto-encoding (ours).

has expanded the understanding from individual labels to a sentence to reflect the visual content in the image. Nevertheless, considering that a single sentence certainly has an upper-bound of descriptive capacity and thus fails to recapitulate every details in the image, the task of image paragraph generation is recently introduced in [Krause *et al.*, 2017; Liang *et al.*, 2017]. The ultimate goal is to generate a coherent paragraph that describes an image in a finer manner.

The difficulty of image paragraph generation originates from two aspects: 1) how to distill the gists/topics that are worthy of mention from an image? 2) how to describe each topic in one sentence while ensuring the coherence across sentences to form a paragraph? One straightforward way in [Krause *et al.*, 2017] to alleviate the first issue relies on Recurrent Neural Networks (RNN), e.g., Long Short-Term Memory (LSTM) network. The region-level features are encoded into a global vector via mean pooling, which is subsequently input into LSTM to decode the topics, as shown in Figure 1(a). In this case, the inherent structure among all regions is not fully exploited, making it difficult to encapsulate the topics in the image. Furthermore, there is no clear evidence that the outputs of LSTM could characterize the topics well. Instead, we novelly devise Convolutional Auto-Encoding (CAE) structure to model the topics of an image in this paper as illustrated in Figure 1(b). CAE, on one hand, abstracts the topics in the encoder by employing convolutions over the region-level features, and on the other, steers the deconvolutional decoder through reconstruction from topics to features. As

such, the learnt topics are potentially more representative and contain the information needed. To address the second issue, we remould a two-level LSTM-based paragraph generation framework, in which paragraph-level LSTM models the dependency holistically across all the sentences in a paragraph and sentence-level LSTM generates words in sequence conditioning on each learnt topic.

By consolidating the idea of modeling sentence topics for image paragraph generation, we present a new Convolutional Auto-Encoding plus Long Short-Term Memory (CAE-LSTM) architecture, as shown in Figure 2. Specifically, Faster R-CNN is firstly implemented to detect a set of salient image regions. We purely perform convolutions over all the region-level features in the encoder to distill the knowledge and extract the topics in the image. The learnt topics are ensured to capture holistic and representative information through achieving high reconstruction quality by the deconvolutional decoder. After that, we fuse all the region-level features via mean pooling as image representation, which are fed into a two-level LSTM networks for paragraph generation. The paragraph-level LSTM typically explores the dependency recursively throughout the procedure of paragraph generation and outputs a new initial state at a time for sentence-level LSTM with the paragraph history. As such, the generation of one sentence is affected by semantic context from the previous sentences when producing the paragraph. The sentence-level LSTM generates the sentence conditioning on each learnt topic, one word at each time step. Please also note that we uniquely design the metric of Coverage which is to encourage global coverage of objects in the sentence and consider such metric as a reward in the self-critical training strategy. The whole CAE-LSTM framework is jointly optimized.

The main contribution of this work is the proposal of modeling gists/topics in an image to boost image paragraph generation. This also leads to the elegant views of how to distill the knowledge and abstract the topics in an image, and how to nicely integrate the learnt topics into sentence generation and ensure inter-sentence dependency in a paragraph, which are problems not yet fully understood in the literature.

## 2 Related Work

**Image Captioning.** The dominant paradigm in modern image captioning is sequence learning methods [Donahue *et al.*, 2015; Li *et al.*, 2019; Vinyals *et al.*, 2015; Xu *et al.*, 2015; You *et al.*, 2016; Liu *et al.*, 2017; Wang *et al.*, 2018; Yao *et al.*, 2017b; Yao *et al.*, 2018; Yao *et al.*, 2017a; Rennie *et al.*, 2017] which utilize CNN plus RNN model to generate free-form and content-relevant captions. [Vinyals *et al.*, 2015] proposes an end-to-end neural network architecture by utilizing LSTM to generate sentence for an image. Later, by amending the discrepancy between training and inference distributions for sequence modeling, [Rennie *et al.*, 2017] presents a self-critical sequence training strategy to further enhance image captioning. Most recently, a combined bottom-up and top-down attention mechanism [Anderson *et al.*, 2018] is proposed for image captioning, which enables attention at the level of objects and salient image regions.

**Image Paragraph Generation.** The task of image paragraph generation has received increasing attention recently, which describes an image with a long, descriptive, and coherent paragraph. [Krause *et al.*, 2017] is one of the early works that leverages a hierarchical RNN for generating paragraph with the input of region features. In particular, a sentence RNN recursively generates sentence topic vectors conditioned on the global vector (i.e., the mean pooling of region-level features) and a word RNN is subsequently adopted to decode each topic into output sentence. [Liang *et al.*, 2017] extends the hierarchical RNN by involving multi-level adversarial discriminators for paragraph generation. The paragraph generator is thus enforced to produce realistic paragraphs with smooth logical transition between sentence topics. Furthermore, [Chatterjee and Schwing, 2018] augments the hierarchical RNN with coherence vectors, global topic vectors, and a formulation of Variational Auto-Encoders [Kingma and Welling, 2013] to further model the inherent ambiguity of associating paragraphs with images.

**Summary.** In short, our approach focuses on the latter scenario, which produces a coherent paragraph to depict an image. Unlike most of the aforementioned methods that recursively decode gists/topics for paragraph generation via LSTM conditioned on the global vector, our work contributes by exploiting a convolutional and deconvolutional auto-encoding module for gists/topics modeling over region-level features. Such design not only abstracts the topics in the convolutional encoder by exploiting inherent structure among all regions, but also steers the deconvolutional decoder through reconstruction to enforce the distilled topics more representative and informative. Moreover, a reward of Coverage is adopted in the self-critical training strategy to encourage global coverage of objects in the paragraph.

## 3 Approach

We devise our Convolutional Auto-Encoding plus Long Short-Term Memory (CAE-LSTM) architecture to facilitate image paragraph generation by modeling gists/topics in an image. CAE-LSTM firstly utilizes a convolutional encoder to encapsulate region-level features into the topics, which are endowed with the holistic and representative information through achieving high reconstruction quality by deconvolutional decoder. The distilled topics are further integrated into a two-level LSTM-based paragraph generator, enabling the inter-sentence dependency modeling in a paragraph via a paragraph-level LSTM and topic-oriented sentence generation through sentence-level LSTM. The overall training of CAE-LSTM is performed by exploring the reconstruction loss to pursue high-quality reconstruction from topics to region-level features, and the sequence-level reward (e.g., CIDEr) plus coverage reward in self-critical training to encourage the maximum coverage of objects in the paragraph. An overview of our model is depicted in Figure 2.

### 3.1 Notation

Suppose we have an image  $I$  to be described by a paragraph  $\mathcal{P}$ , where  $\mathcal{P} = \{s_k\}_{k=1}^K$  consisting of  $K$  sentences. Each sentence  $s_k = \{w_{(k,t)}\}_{t=1}^{T_k}$  is composed of  $T_k$  words and each

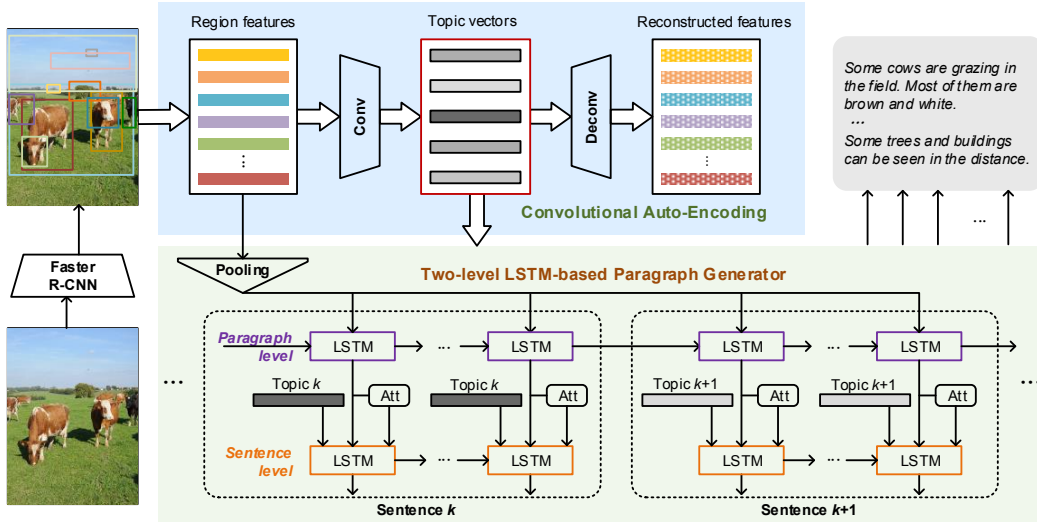


Figure 2: An overview of our Convolutional Auto-Encoding plus Long Short-Term Memory (CAE-LSTM) architecture for image paragraph generation. All region-level features of image regions detected by Faster R-CNN are firstly injected into Convolutional Auto-Encoding (CAE) module for triggering the modeling of topics in the image. In particular, a convolutional encoder is leveraged to distill the knowledge from region-level features and extract the topics. Next, a deconvolutional decoder is steered to reproduce the region-level features from the distilled topics. As such, the learnt topics are ensured to capture the holistic and representative information that are worthy of mention from all regions. After that, we fuse all the region-level features via mean pooling as the image feature, which is fed into a two-level LSTM-based paragraph generator with attention (Att) for paragraph generation. Here the paragraph-level LSTM typically explores the inter-sentence dependency in a paragraph with the recursive input of mean-pooled image feature and the sentence-level LSTM generates the sentence conditioning on each learnt topic. Moreover, we include a reward of Coverage to encourage global coverage of objects in the paragraph for self-critical training.

word  $w_{(k,t)}$  is represented as  $D_s$ -dimensional textual feature  $\mathbf{w}_{(k,t)}$ . Faster R-CNN [Ren *et al.*, 2015] is firstly leveraged to produce a set of detected objects  $\mathcal{V} = \{r_m\}_{m=1}^M$  with  $M$  image regions in  $I$ .  $\mathbf{v}_m^0 \in \mathbb{R}^{D_0}$  denotes the  $D_0$ -dimensional feature of each image region  $r_m$ .

### 3.2 CAE for Topic Modeling

The most typical way to distill the topics from an image is to encode region-level features into a global vector via mean pooling and a LSTM-based decoder is subsequently utilized to recursively output topics conditioning on the global vector, while the inherent structure among all regions is not fully exploited. Furthermore, there is no clear evidence that the outputs of LSTM could characterize the topics well. Here we devise a Convolutional Auto-Encoding (CAE) module for topic modeling. The spirit behind follows the philosophy that the holistical topics abstraction over all region-level features via convolutional encoder and the reconstruction from topics to features via deconvolutional decoder can enforce the learnt topics to be more representative and informative.

**Convolutional Encoder.** Given the set of detected regions  $\mathcal{V} = \{r_m\}_{m=1}^M$ , a convolutional encoder is leveraged to encapsulate all the image regions into  $K$  topic vectors, which is purely applied with convolutions. Specifically, each region  $\mathbf{v}_m^0$  is firstly embedded into a  $D_1$ -dimensional region-level feature  $\mathbf{v}_m$  through a linear layer. A region feature map  $\mathbf{V} \in \mathbb{R}^{M \times D_1 \times 1}$  is thus constructed by concatenating all region-level features (i.e.,  $\mathbf{v}_m$  is the  $m$ -th column of  $\mathbf{V}$ ), which will be set as the input of convolutional encoder. Here  $M$ ,  $D_1$ , and 1 denote the width, height, and the number of

channels for the input feature map  $\mathbf{V}$ , respectively. Next, we utilize a convolutional layer (*conv*) to encode the feature map  $\mathbf{V}$  into the topic feature map  $\mathbf{V}^s$  consisting of topic vectors:

$$\mathbf{V}^s = \text{conv}(\mathbf{V}) \in \mathbb{R}^{1 \times D_2 \times K}. \quad (1)$$

For the convolutional layer *conv*, the size of the convolutional filter is set as  $M \times C_1$ , the stride size is  $C_2$ , and the filter number is  $K$ . Here each vector along channel dimension within the topic feature map  $\mathbf{V}^s$  can be regarded as the  $k$ -th distilled topic vector  $\mathbf{v}_k^s \in \mathbb{R}^{D_2}$ .

**Deconvolutional Decoder.** For deconvolutional decoder, we perform deconvolution (i.e., the conjugate operation of convolution) to decode the distilled topic feature map,  $\mathbf{V}^s$ , back to the region feature map. In particular, given the distilled topic feature map  $\mathbf{V}^s$ , a deconvolutional layer (*deconv*) is adopted to reproduce the region feature map  $\tilde{\mathbf{V}}$ :

$$\tilde{\mathbf{V}} = \text{deconv}(\mathbf{V}^s) \in \mathbb{R}^{M \times D_1 \times 1}. \quad (2)$$

The filter size and the stride size in *deconv* are set the same as the convolutional layer *conv*. Here the  $m$ -th column of the reconstructed region feature map  $\tilde{\mathbf{V}}$  corresponds to the reconstructed region-level feature for  $m$ -th image region.

**Reconstruction Loss.** To measure the quality of reconstruction for our CAE module, a global reconstruction loss is defined as the  $L_1$  distance between the input region feature map  $\mathbf{V}$  and reconstructed region feature map  $\tilde{\mathbf{V}}$ :

$$L_{rec}(\tilde{\mathbf{V}}, \mathbf{V}) = \|\tilde{\mathbf{V}} - \mathbf{V}\|_1. \quad (3)$$

We also experimented with  $L_2$  distance as the reconstruction loss but that did not make a major difference. By minimizing

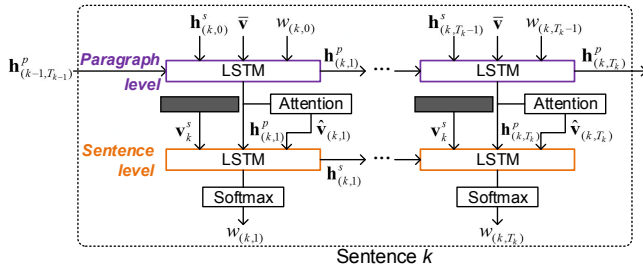


Figure 3: Illustration of two-level LSTM-based paragraph generator.

the reconstruction loss, the distilled topic vectors are enforced to capture holistic and representative information from all regions through achieving higher reconstruction quality.

Note that besides acting as the gists to guide the sentence generation in the followed two-level LSTM networks, the distilled topic vectors can also be utilized to determine the number of generated sentences. Specifically, each topic vector is additionally injected into a linear layer to obtain a distribution over two states  $\{\text{CONTINUE}=0, \text{STOP}=1\}$ , which determines whether the sentence is the last one in the paragraph.

### 3.3 Two-level LSTM-based Paragraph Generator

Inspired by the two-layer LSTM structure with top-down attention mechanism in [Anderson *et al.*, 2018], we remould a two-level LSTM-based paragraph generator consisting of a paragraph-level LSTM for inter-sentence dependency modeling and a sentence-level LSTM for sentence generation conditioning on each distilled topic, as illustrated in Figure 3.

To be specific, at each time step  $t$  for generating the  $k$ -th sentence in a paragraph, the input vector  $\mathbf{x}^p_{(k,t)}$  of paragraph-level LSTM is defined as the concatenation of the previous output of sentence LSTM  $\mathbf{h}^s_{(k,t-1)} \in \mathbb{R}^H$ , the mean-pooled image feature  $\bar{\mathbf{v}} = \frac{1}{M} \sum_{m=1}^M \mathbf{v}_m$ , and the embedding of previously generated word  $\mathbf{w}^{(k,t-1)}$ :

$$\mathbf{x}^p_{(k,t)} = [\mathbf{h}^s_{(k,t-1)}, \bar{\mathbf{v}}, \mathbf{W}_s \mathbf{w}^{(k,t-1)}], \quad (4)$$

where  $\mathbf{W}_s$  is transformation matrix for input word. Such input collects the maximum contextual information for paragraph-level LSTM to model inter-sentence dependency.

Next, given the output  $\mathbf{h}^p_{(k,t)} \in \mathbb{R}^H$  of paragraph-level LSTM and the corresponding topic vector  $\mathbf{v}_k^s \in \mathbb{R}^{D_2}$ , a normalized attention distribution over all regions is measured as

$$\begin{aligned} a^m_{(k,t)} &= \mathbf{W}_{att} [\tanh(\mathbf{W}_v \mathbf{v}_m + \mathbf{W}_h \mathbf{h}^p_{(k,t)} + \mathbf{W}_t \mathbf{v}_k^s)], \\ \alpha_{(k,t)} &= \text{softmax}(\mathbf{a}_{(k,t)}), \end{aligned} \quad (5)$$

where  $a^m_{(k,t)}$  is the  $m$ -th element of  $\mathbf{a}_{(k,t)}$ ,  $\mathbf{W}_{att} \in \mathbb{R}^{1 \times D_3}$ ,  $\mathbf{W}_v \in \mathbb{R}^{D_3 \times D_1}$ ,  $\mathbf{W}_h \in \mathbb{R}^{D_3 \times H}$ , and  $\mathbf{W}_t \in \mathbb{R}^{D_3 \times D_2}$  are transformation matrices, respectively. Here the  $m$ -th element  $\alpha^m_{(k,t)}$  of the attention distribution  $\alpha_{(k,t)}$  denotes the attention probability of  $\mathbf{v}_m$ . The attended image feature  $\hat{\mathbf{v}}_{(k,t)} = \sum_{m=1}^M \alpha^m_{(k,t)} \mathbf{v}_m$  is thus calculated by aggregating all region-level features weighted by attention.

We further set the concatenation of the attended image feature  $\hat{\mathbf{v}}_{(k,t)}$ , the corresponding topic vector  $\mathbf{v}_k^s$  and the output  $\mathbf{h}^p_{(k,t)}$  of paragraph-level LSTM, as the input  $\mathbf{x}^s_{(k,t)}$  of

sentence-level LSTM for topic-oriented sentence generation:

$$\mathbf{x}^s_{(k,t)} = [\hat{\mathbf{v}}_{(k,t)}, \mathbf{v}_k^s, \mathbf{h}^p_{(k,t)}]. \quad (6)$$

The output  $\mathbf{h}^s_{(k,t)}$  of sentence-level LSTM is fed into a softmax layer to predict the next word  $w_{(k,t)}$ .

Note that the initial state of sentence-level LSTM is set to zero at every beginning of sentence generation. The state of paragraph-level LSTM is initialized with zero only at the beginning of the entire paragraph generation since paragraph-level LSTM should explore inter-sentence dependency recursively throughout whole procedure of paragraph generation.

### 3.4 Self-critical Training with Coverage Reward

To further boost our paragraph generation model by amending the discrepancy between training and inference, we adopt the self-critical training strategy [Rennie *et al.*, 2017] to optimize two-level LSTM-based paragraph generator. Despite having high quantitative scores, qualitative analysis shows that the captioning models solely optimized with sequence-level loss/reward are often limited to describing very generic information of objects, or rely on prior information and correlations from training examples, and resulting frequently in undesired effects such as object hallucination [Lu *et al.*, 2018]. As a result, besides the sequence-level reward (e.g., CIDEr), we additionally include a coverage reward in self-critical training strategy to encourage the global coverage of objects in the paragraph and thus improve the paragraph.

Specifically, we select the top- $1k$  frequent objects from the vocabulary in the training data as the high-frequency objects. The coverage reward is defined as the coverage ratio of high-frequency objects for the generated paragraph relative to ground-truth paragraph:  $R^c = \frac{|Q_g \cap Q_{gt}|}{|Q_{gt}|}$ , where  $|\cdot|$  denotes the number of elements in a set. Here  $Q_g$  and  $Q_{gt}$  represent the set of high-frequency objects mentioned in the generated paragraph and ground-truth paragraph respectively. Accordingly, the final reward in self-critical training strategy is measured as the combination of CIDEr reward ( $R^d$ ) and coverage reward:  $R = \beta R^c + R^d$ , where  $\beta$  is the tradeoff parameter.

## 4 Experiments

### 4.1 Datasets and Experimental Settings

**Dataset.** We conducted the experiments and evaluated our CAE-LSTM on Stanford image paragraph dataset (Stanford) [Krause *et al.*, 2017], a benchmark in the field of image paragraph generation. The dataset contains 19,551 images and there is one human-annotated paragraph per image. On average, each paragraph has 67.5 words and each sentence consists of 11.91 words. In our experiments, we follow the widely used settings in [Krause *et al.*, 2017] and take 14,575 images for training, 2,487 for validation and 2,489 for testing.

**Compared Methods.** We compare the proposed method with the following state-of-the-art methods: (1) **Image-Flat** [Karpathy and Fei-Fei, 2015] is a standard image captioning model which directly decodes an image into a paragraph word by word, via a single LSTM. (2) **Regions-Hierarchical** [Krause *et al.*, 2017] adopts a hierarchical LSTM to generate a paragraph, sentence by sentence. (3) **RTT-GAN** [Liang *et al.*

*et al.*, 2017] integrates sentence attention and word attention into the hierarchical LSTM, coupled with adversarial training strategy. (4) **CapG-RevG** [Chatterjee and Schwing, 2018] leverages coherence vectors/global topic vectors to generate coherent paragraphs and maintains the diversity of the paragraphs by a variational auto-encoder formulation. (5) **CAE-LSTM** is the proposal in this paper. Moreover, one degraded version of CAE-LSTM, i.e., **LSTM-ATT**, is devised to model topics via LSTM (instead of CAE) and adopt the same two-level LSTM architecture with attention for paragraph generation, which is trained without self-critical training strategy.

**Settings.** For each image, we apply Faster R-CNN to detect objects within this image and select top  $M = 50$  regions with highest detection confidences to represent the image. Each region is represented as the 4,096-dimensional output of fc7 layer after RoI pooling from the conv5\_3 feature map of Faster R-CNN (backbone: VGG16 [Simonyan and Zisserman, 2014]). The Faster R-CNN is pre-trained over Visual Genome [Krishna *et al.*, 2017], similar to [Anderson *et al.*, 2018]. To build the vocabulary, all the words from the training set are converted to lowercase and those which appear less than 4 times are omitted. Each word is represented as “one-hot” vector (binary index vector in a vocabulary). The maximum sentence number  $K$  is 6 and the maximum word number in a sentence is 20 (padded where necessary). For our CAE, the convolutional filter size in the convolutional layer is set as  $C_1 = 26$  with stride  $C_2 = 2$ . The dimensions of the embedded region-level feature and distilled topic vector are set as  $D_1 = 1,024$  and  $D_2 = 500$ . For the two-level LSTM networks, the dimension of hidden state in each LSTM is  $H = 1,000$ . The dimension of the hidden layer for measuring attention distribution is  $D_3 = 512$ . Three popular metrics are adopted: METEOR [Banerjee and Lavie, 2005], CIDEr [Vedantam *et al.*, 2015], and BLEU-4 [Papineni *et al.*, 2002]. We compute all metrics with the released code from Microsoft COCO Evaluation Server [Chen *et al.*, 2015].

**Implementation Details.** We follow the two-phrase training recipe in [Rennie *et al.*, 2017] to train our CAE-LSTM. For the first phrase, we set the learning rate as  $1 \times 10^{-4}$  and the training of CAE-LSTM is thus performed by jointly leveraging the reconstruction loss in CAE and the cross entropy loss in two-level LSTM paragraph generator. Here we evaluate the model at each epoch on the validation set and select the model with the best CIDEr score as an initialization for the next training phrase. For the second phrase of self-critical training, the learning rate is set as  $5 \times 10^{-6}$  and CAE-LSTM is further optimized with the combination of CIDEr reward and coverage reward. During inference, we adopt inference constraint in [Melas-Kyriazi *et al.*, 2018] to penalize trigram repetition. The tradeoff parameter  $\beta$  is set as 8 according to the validation performance. Note that Batch normalization [Ioffe and Szegedy, 2015] and dropout [Srivastava *et al.*, 2014] (dropout rate: 0.5) are applied in our experiments.

One issue of training our CAE-LSTM is to determine the order of the regions in the concatenation of region-level features as the input of CAE module. In the experiments, the region orders are determined according to the objectiveness score of each image region. We additionally studied the re-

Method	C	M	B-4
Image-Flat [Karpathy and Fei-Fei, 2015]	11.06	12.82	7.71
Regions-Hierarchical [Krause <i>et al.</i> , 2017]	13.52	15.95	8.69
RTT-GAN [Liang <i>et al.</i> , 2017]	20.36	18.39	9.21
CapG-RevG [Chatterjee and Schwing, 2018]	20.93	18.62	9.43
LSTM-ATT	20.17	17.72	8.72
CAE-LSTM	<b>25.15</b>	<b>18.82</b>	<b>9.67</b>

Table 1: Performance of our CAE-LSTM and other state-of-the-art methods on Stanford, where C, M, and B-4 are short for CIDEr, METEOR, and BLEU-4. All values are reported as percentage (%).

gion orders by sorting all the regions with regard to the confidence score of bounding box, or simply using fixed random order. The CIDEr score constantly fluctuates within 0.008 when using different orders of regions for concatenation in our CAE-LSTM. The results indicate that the performance is not sensitive to the selection of region order.

## 4.2 Performance Comparison and Analysis

**Quantitative Analysis.** The performances of different models on the Stanford dataset are shown in Table 1. Overall, the results across three evaluation metrics consistently indicate that our proposed CAE-LSTM achieves better performances against other state-of-the-art techniques including non-attention models (Image-Flat, Regions-Hierarchical, and CapG-RevG) and attention-based approaches (LSTM-ATT and RTT-GAN). Specifically, the CIDEr and METEOR scores of our CAE-LSTM can achieve 25.15% and 18.82%, which makes the relative improvement over the best competitor CapG-RevG by 20.2% and 1.1%, respectively. As expected, by additionally modeling topics/gists in an image via LSTM, Regions-Hierarchical exhibits better performance than Image-Flat which leaves inter-sentence dependency unexploited. Moreover, LSTM-ATT leads to a performance boost over Regions-Hierarchical, which directly encodes an image as a global representation by performing mean pooling over all region-level features. The results basically indicate the advantage of region-level attention mechanism in the two-level LSTM networks by learning to focus on the image regions that are most indicative to infer the next word. Most specifically, RTT-GAN and CapG-RevG by modeling reality and diversity of paragraphs with Generative Adversarial Networks and Variational Auto-Encoders, improves LSTM-ATT. However, the performances of RTT-GAN and CapG-RevG are still lower than our CAE-LSTM which exploits the inherent structure among all image regions for topic modeling in a convolutional and deconvolutional auto-encoding framework.

**Qualitative Analysis.** Figure 4 shows several paragraph examples generated by LSTM-ATT, CAE-LSTM and one human-annotated Ground Truth (GT) paragraph. From these exemplar results, it is easy to see that all of these paragraph generation models can produce somewhat relevant paragraphs, while our proposed CAE-LSTM can generate coherent and accurate paragraphs by learning to distill the gists/topics from an image via a convolutional auto-encoding module to guide paragraph generation. For example, compared to LSTM-ATT which repeatedly generates the phrase “holding a tennis racket” in the 2-th and 3-th sentences for the first image, the pro-


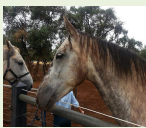
Images	LSTM-ATT	CAE-LSTM	GT
	<ul style="list-style-type: none"> <li>. A man is <u>standing on a tennis court</u>.</li> <li>. The man is <u>holding a tennis racket</u>.</li> <li>. The man is <u>holding a tennis racket in his hand</u>.</li> <li>. The court is green and green with white lines.</li> <li>. The court is surrounded by a tall green tree with green leaves.</li> <li>. There are many trees in the background and a <u>fence</u> in the background.</li> </ul>	<ul style="list-style-type: none"> <li>. A man is <u>standing on a tennis court</u>.</li> <li>. He is wearing a white baseball cap <u>a white shirt and blue shorts</u>.</li> <li>. The man is <u>holding a tennis racket in his hand</u>.</li> <li>. There are trees behind the man.</li> <li>. Part of a tall tree can be seen <u>behind the fence</u>.</li> <li>. A tall palm tree is standing behind the court.</li> </ul>	<ul style="list-style-type: none"> <li>. There is a person <u>standing on top of a tennis court</u>.</li> <li>. The person is wearing <u>a white shirt and blue shorts</u>.</li> <li>. The person is wearing a white hat and dark colored shoes.</li> <li>. They are <u>holding a blue tennis racket in their hand</u>.</li> <li>. There is a green fence behind them.</li> <li>. There is a bush <u>behind the fence</u>.</li> <li>. There are cars parked in a lot behind the bush.</li> </ul>
	<ul style="list-style-type: none"> <li>. A large brown <u>horse</u> is standing next to a metal fence.</li> <li>. The horse is brown and has brown spots on its mane.</li> <li>. <u>There are trees in the background</u> and the horse is standing on a metal fence.</li> <li>. <u>There are many trees and bushes in the background</u>.</li> <li>. <u>The sky above is blue</u> with white clouds.</li> <li>. <u>The sky is blue</u> in color.</li> </ul>	<ul style="list-style-type: none"> <li>. Two <u>horses</u> are standing next to a fence.</li> <li>. The horse on the right is looking to the left.</li> <li>. There is a metal fence behind the horse.</li> <li>. The horse is standing <u>near a metal gate</u>.</li> <li>. <u>There are trees in the background</u>.</li> <li>. The sky is blue and the sky is clear.</li> </ul>	<ul style="list-style-type: none"> <li>. It is the middle of the day during the summer.</li> <li>. <u>A horse is standing behind a silver gate</u> at the african like setting.</li> <li>. There is dirt and twigs on the ground.</li> <li>. <u>In the background there is trees</u>.</li> <li>. There is a person with a blue shirt behind the horse and a second horse on the opposite side of the gate.</li> <li>. There are other objects in the area such as poles and wires.</li> </ul>

Figure 4: Examples for paragraphs generated by LSTM-ATT, CAE-LSTM and human-annotated Ground Truth paragraphs (GT) on Stanford dataset (better viewed in color). The words coded with same colors indicate the accurate semantic matches between the generated paragraphs and ground truth paragraphs. We underline the semantically similar phrases/sentences in a paragraph.




Figure 5: The performance improvement by considering one more design (i.e., Convolutional Encoding (CE), Reconstruction Loss (RL), and Self-critical training with Coverage (SC)) in CAE-LSTM, over the performance of LSTM-ATT across different metrics.

duced paragraph for our CAE-LSTM is more coherent to describe the diverse and accurate topics in the image. The results again verify that the convolutional auto-encoding module in our CAE-LSTM could characterize the topics better than LSTM. Moreover, by additionally measuring the coverage of objects as a reward in the self-critical training strategy, our CAE-LSTM is encouraged to produce paragraphs which cover more objects in images, leading to more descriptive paragraph with objects, e.g., “shirt” and “shorts.”

**Ablation study.** Here we study how each design in our CAE-LSTM influences the overall performance. Convolutional Encoding (CE) replaces the LSTM for topic modeling in LSTM-ATT with the convolutional encoder in our devised CAE module. Reconstruction Loss (RL) enforces the distilled topics to capture holistic and representative information via reconstruction from topics to region-level features. Self-critical training with Coverage (SC) further encourages global coverage of objects with a coverage reward in self-critical training. Figure 5 shows the performance improvement by considering one more design for paragraph generation in our CAE-LSTM, over the performance of LSTM-ATT across different metrics. The results across different metrics consistently indicate that the topic modeling via convolutional encoding in CE leads to a performance boost compared to LSTM-ATT which capitalizes on LSTM to explore dependency among topics. Furthermore, by integrating reconstruction loss into the learning of topics, CE+RL exhibits better performance than CE. The results demonstrate the advantage of reconstruction from topics to region-level features in our CAE. By

additionally exploiting coverage reward in self-critical training, CE+RL+SC further boosts up the performances.

**Human Evaluation.** To better understand how satisfactory are the paragraphs generated from different methods, we conducted Turing Test to evaluate our CAE-LSTM against the baseline LSTM-ATT. In particular, six evaluators are invited and a subset of 1,000 images is randomly selected from testing set for human evaluation. We show all the evaluators once only one paragraph generated by different approach or human annotation and they are asked: can you determine whether the given paragraph has been generated by a system or by a human being? From evaluators’ responses, we calculate the percentage of paragraphs that pass the Turing Test. The results of Turing Test for Human, CAE-LSTM, and LSTM-ATT are 88.5%, 39.8%, and 14.7%, which clearly show that our CAE-LSTM outperforms LSTM-ATT.

## 5 Conclusion

We have present Convolutional Auto-Encoding plus Long Short-Term Memory (CAE-LSTM), which explores the modeling of sentence topics to boost image paragraph generation. Particularly, we study the problem from the viewpoint of topic distillation in an auto-encoding manner. To verify our claim, we have devised a purely convolutional structure, in which an encoder abstracts topics by employing convolutions over region-level image features and a deconvolutional decoder is to validate the topics through high-quality reconstruction from topics to features. A two-level LSTM-based paragraph generator is then remould, on one hand, to model dependency across sentences in a paragraph, and on the other, to generate the sentence conditioning on each learnt topic. Experiments conducted on Stanford image paragraph dataset verify our proposal and analysis. Performance improvements are observed when comparing to state-of-the-art image paragraph generation techniques.

## Acknowledgements

This work was partially supported by the National Natural Science Foundation of China (Grant No. 61732007 and U1611461).

## References

- [Anderson *et al.*, 2018] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and vqa. In *CVPR*, 2018.
- [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.
- [Chatterjee and Schwing, 2018] Moitrey Chatterjee and Alexander G Schwing. Diverse and coherent paragraph generation from images. In *ECCV*, 2018.
- [Chen *et al.*, 2015] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- [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.
- [Ioffe and Szegedy, 2015] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [Karpathy and Fei-Fei, 2015] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015.
- [Kingma and Welling, 2013] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [Krause *et al.*, 2017] Jonathan Krause, Justin Johnson, Ranjay Krishna, and Li Fei-Fei. A hierarchical approach for generating descriptive image paragraphs. In *CVPR*, 2017.
- [Krishna *et al.*, 2017] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *I-JCV*, 2017.
- [Li *et al.*, 2019] Yehao Li, Ting Yao, Yingwei Pan, Hongyang Chao, and Tao Mei. Pointing novel objects in image captioning. In *CVPR*, 2019.
- [Liang *et al.*, 2017] Xiaodan Liang, Zhiting Hu, Hao Zhang, Chuang Gan, and Eric P Xing. Recurrent topic-transition gan for visual paragraph generation. In *ICCV*, 2017.
- [Liu *et al.*, 2017] Siqi Liu, Zhenhai Zhu, Ning Ye, Sergio Guadarrama, and Kevin Murphy. Optimization of image description metrics using policy gradient methods. In *ICCV*, 2017.
- [Lu *et al.*, 2018] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Neural baby talk. In *CVPR*, 2018.
- [Melas-Kyriazi *et al.*, 2018] Luke Melas-Kyriazi, Alexander Rush, and George Han. Training for diversity in image paragraph captioning. In *EMNLP*, 2018.
- [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.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [Rennie *et al.*, 2017] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Srivastava *et al.*, 2014] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *JMLR*, 2014.
- [Vedantam *et al.*, 2015] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015.
- [Vinyals *et al.*, 2015] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015.
- [Wang *et al.*, 2018] Jing Wang, Jianlong Fu, Jinhui Tang, Zechao Li, and Tao Mei. Show, reward and tell: Automatic generation of narrative paragraph from photo stream by adversarial training. In *AAAI*, 2018.
- [Xu *et al.*, 2015] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.
- [Yao *et al.*, 2017a] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Incorporating copying mechanism in image captioning for learning novel objects. In *CVPR*, 2017.
- [Yao *et al.*, 2017b] Ting Yao, Yingwei Pan, Yehao Li, Zhao-fan Qiu, and Tao Mei. Boosting image captioning with attributes. In *ICCV*, 2017.
- [Yao *et al.*, 2018] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *ECCV*, 2018.
- [You *et al.*, 2016] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. Image captioning with semantic attention. In *CVPR*, 2016.