

Semantic Compositional Networks for Visual Captioning

Zhe Gan[†], Chuang Gan^{*}, Xiaodong He[‡], Yunchen Pu[†]
 Kenneth Tran[‡], Jianfeng Gao[‡], Lawrence Carin[†], Li Deng[‡]

[†]Duke University, ^{*}Tsinghua University, [‡]Microsoft Research, Redmond, WA 98052, USA

{zhe.gan, yunchen.pu, lcarin}@duke.edu, ganchuang1990@gmail.com

{xiaoh, ktran, jfgao, deng}@microsoft.com

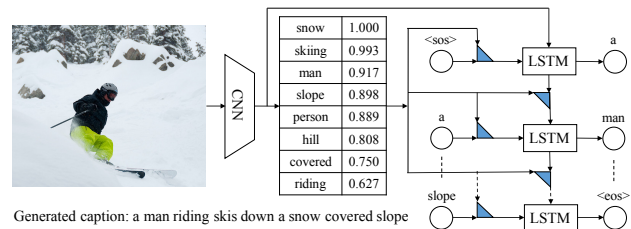
Abstract

A *Semantic Compositional Network (SCN)* is developed for image captioning, in which semantic concepts (*i.e.*, tags) are detected from the image, and the probability of each tag is used to compose the parameters in a long short-term memory (LSTM) network. The SCN extends each weight matrix of the LSTM to an ensemble of tag-dependent weight matrices. The degree to which each member of the ensemble is used to generate an image caption is tied to the image-dependent probability of the corresponding tag. In addition to captioning images, we also extend the SCN to generate captions for video clips. We qualitatively analyze semantic composition in SCNs, and quantitatively evaluate the algorithm on three benchmark datasets: COCO, Flickr30k, and Youtube2Text. Experimental results show that the proposed method significantly outperforms prior state-of-the-art approaches, across multiple evaluation metrics.

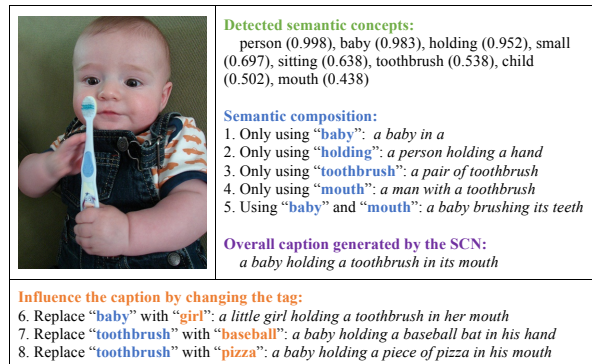
1. Introduction

There has been a recent surge of interest in developing models that can generate captions for images or videos, termed visual captioning. Most of these approaches learn a probabilistic model of the caption, conditioned on an image or a video [29, 47, 13, 20, 48, 52, 10, 46, 32, 56]. Inspired by the successful use of the encoder-decoder framework employed in machine translation [2, 8, 40], most recent work on visual captioning employs a convolutional neural network (CNN) as an encoder, obtaining a fixed-length vector representation of a given image or video. A recurrent neural network (RNN), typically implemented with long short-term memory (LSTM) units [17], is then employed as a decoder to generate a caption.

Recent work shows that adding explicit high-level semantic concepts (*i.e.*, tags) of the input image/video can further improve visual captioning. As shown in [49, 54], detecting explicit semantic concepts encoded in an image, and adding



(a) Overview of the proposed model.



(b) Examples of SCN-based image captioning.

Figure 1: Model architecture and illustration of semantic composition. Each triangle symbol represents an ensemble of tag-dependent weight matrices. The number next to a semantic concept (*i.e.*, a tag) is the probability that the corresponding semantic concept is presented in the input image.

this high-level semantic information into the CNN-LSTM framework, has improved performance significantly. Specifically, [49] feeds the semantic concepts as an *initialization* step into the LSTM decoder. In [54], a model of semantic attention is proposed which selectively attends to semantic concepts through a soft attention mechanism [2]. On the other hand, although significant performance improvements were achieved, integration of semantic concepts into the LSTM-based caption generation process is constrained in these methods; *e.g.*, only through soft attention or initializa-

tion of the first step of the LSTM.

In this paper, we propose the Semantic Compositional Network (SCN) to more effectively assemble the meanings of individual tags to generate the caption that describes the overall meaning of the image, as illustrated in Figure 1a. Similar to the conventional CNN-LSTM-based image captioning framework, a CNN is used to extract the visual feature vector, which is then fed into a LSTM for generating the image caption (for simplicity, in this discussion we refer to images, but the method is also applicable to video). However, unlike the conventional LSTM, the SCN extends each weight matrix of the conventional LSTM to an *ensemble* of tag-dependent weight matrices, subject to the probabilities that the tags are present in the image. These tag-dependent weight matrices form a weight tensor with a large number of parameters. In order to make learning feasible, we factorize that tensor to be a three-way matrix product, which dramatically reduces the number of free parameters to be learned, while also yielding excellent performance.

The main contributions of this paper are as follows: (i) We propose the SCN to effectively compose individual semantic concepts for image captioning. (ii) We perform comprehensive evaluations on two image captioning benchmarks, demonstrating that the proposed method outperforms previous state-of-the-art approaches by a substantial margin. For example, as reported by the COCO official test server, we achieve a BLEU-4 of 33.1, an improvement of 1.5 points over the current published state-of-the-art [54]. (iii) We extend the proposed framework from image captioning to video captioning, demonstrating the versatility of the proposed model. (iv) We also perform a detailed analysis to study the SCN, showing that the model can adjust the caption smoothly by modifying the tags.

2. Related work

We focus on recent neural-network-based literature for caption generation, as these are most relevant to our work. Such models typically extract a visual feature vector via a CNN, and then send that vector to a language model for caption generation. Representative works include [7, 9, 10, 20, 23, 24, 29, 48] for image captioning and [10, 46, 47, 56, 3, 35, 11] for video captioning. The differences of the various methods mainly lie in the types of CNN architectures and language models. For example, the vanilla RNN [12] was used in [29, 20], while the LSTM [17] was used in [48, 46, 47]. The visual feature vector was only fed into the RNN once at the first time step in [48, 20], while it was used at each time step of the RNN in [29].

Most recently, [52] utilized an attention-based mechanism to learn where to focus in the image during caption generation. This work was followed by [53] which introduced a review module to improve the attention mechanism and [28] which proposed a method to improve the correct-

ness of visual attention. Moreover, a variational autoencoder was developed in [34] for image captioning. Other related work includes [32] for video captioning and [1] for composing sentences that describe novel objects.

Another class of models uses semantic information for caption generation. Specifically, [18] applied retrieved sentences as additional semantic information to guide the LSTM when generating captions, while [13, 49, 54] applied a semantic-concept-detection process [15] before generating sentences. In addition, [13] also proposes a deep multimodal similarity model to project visual features and captions into a joint embedding space. This line of methods represents the current state of the art for image captioning. Our proposed model also lies in this category; however, distinct from the aforementioned approaches, our model uses weight *tensors* in LSTM units. This allows learning an ensemble of semantic-concept-dependent weight matrices for generating the caption.

Related to but distinct from the hierarchical composition in a recursive neural network [37], our model carries out implicit composition of concepts, and there is no hierarchical relationship among these concepts. Figure 1b illustrates the semantic composition manifested in the SCN model. Specifically, a set of semantic concepts, such as “*baby, holding, toothbrush, mouth*”, are detected with high probabilities. If only one semantic concept is turned on, the model will generate a description covering only part of the input image, as shown in sentences 1-5 of Figure 1b; however, by assembling all these semantic concepts, the SCN is able to generate a comprehensive description “*a baby holding a toothbrush in its mouth*”. More interestingly, as shown in sentences 6-8 of Figure 1b, the SCN also has great flexibility to adjust the generation of the caption by changing certain semantic concepts.

The tensor factorization method is used to make the SCN compact and simplify learning. Similar ideas have been exploited in [25, 30, 38, 39, 41, 50, 14]. In [10, 19, 23] the authors also briefly discussed using the tensor factorization method for image captioning. Specifically, visual features extracted from CNNs are utilized in [10, 23], and an inferred scene vector is used in [19] for tensor factorization. In contrast to these works, we use the semantic-concept vector that is formed by the probabilities of all tags to weight the basis LSTM weight matrices in the ensemble. Our semantic-concept vector is more powerful than the visual-feature vector [10, 23] and the scene vector [19] in terms of providing explicit semantic information of an image, hence leading to significantly better performance, as shown in our quantitative evaluation. In addition, the usage of semantic concepts also makes the proposed SCN more interpretable than [10, 19, 23], as shown in our qualitative analysis, since each unit in the semantic-concept vector corresponds to an explicit tag.

3. Semantic compositional networks

3.1. Review of RNN for image captioning

Consider an image \mathbf{I} , with associated caption \mathbf{X} . We first extract feature vector $\mathbf{v}(\mathbf{I})$, which is often the top-layer features of a pretrained CNN. Henceforth, for simplicity, we omit the explicit dependence on \mathbf{I} , and represent the visual feature vector as \mathbf{v} . The length- T caption is represented as $\mathbf{X} = (x_1, \dots, x_T)$, with x_t a 1-of- V (“one hot”) encoding vector, with V the size of the vocabulary. The length T typically varies among different captions.

The t -th word in a caption, x_t , is linearly embedded into an n_x -dimensional real-valued vector $\mathbf{w}_t = \mathbf{W}_e \mathbf{x}_t$, where $\mathbf{W}_e \in \mathbb{R}^{n_x \times V}$ is a word embedding matrix (learned), *i.e.*, \mathbf{w}_t is a column of \mathbf{W}_e chosen by the one-hot \mathbf{x}_t . The probability of caption \mathbf{X} given image feature vector \mathbf{v} is defined as

$$p(\mathbf{X}|\mathbf{I}) = \prod_{t=1}^T p(x_t | x_0, \dots, x_{t-1}, \mathbf{v}), \quad (1)$$

where x_0 is defined as a special start-of-the-sentence token. All the words in the caption are sequentially generated using a RNN, until the end-of-the-sentence symbol is generated. Specifically, each conditional $p(x_t | x_{<t}, \mathbf{v})$ is specified as $\text{softmax}(\mathbf{V}\mathbf{h}_t)$, where \mathbf{h}_t is recursively updated through $\mathbf{h}_t = \mathcal{H}(\mathbf{w}_{t-1}, \mathbf{h}_{t-1}, \mathbf{v})$, and \mathbf{h}_0 is defined as a zero vector (\mathbf{h}_0 is *not* updated during training). \mathbf{V} is the weight matrix connecting the RNN’s hidden state, used for computing a distribution over words. Bias terms are omitted for simplicity throughout the paper.

Without loss of generality, we begin by discussing an RNN with a simple transition function $\mathcal{H}(\cdot)$; this is generalized in Section 3.4 to the LSTM. Specifically, $\mathcal{H}(\cdot)$ is defined as

$$\mathbf{h}_t = \sigma(\mathbf{W}\mathbf{x}_{t-1} + \mathbf{U}\mathbf{h}_{t-1} + \mathbb{1}(t=1) \cdot \mathbf{C}\mathbf{v}), \quad (2)$$

where $\sigma(\cdot)$ is a logistic sigmoid function, and $\mathbb{1}(\cdot)$ represents an indicator function. Feature vector \mathbf{v} is fed into the RNN at the beginning, *i.e.*, at $t = 1$. \mathbf{W} is defined as the input matrix, and \mathbf{U} is termed the recurrent matrix. The model in (2) is illustrated in Figure 2(a).

3.2. Semantic concept detection

The SCN developed below is based on the detection of semantic concepts, *i.e.*, tags, in the image under test. In order to detect such from an image, we first select a set of tags from the caption text in the training set. Following [13], we use the K most common words in the training captions to determine the vocabulary of tags, which includes the most frequent nouns, verbs, or adjectives.

In order to predict semantic concepts given a test image, motivated by [49, 44], we treat this problem as a multi-label classification task. Suppose there are N training examples, and $\mathbf{y}_i = [y_{i1}, \dots, y_{iK}] \in \{0, 1\}^K$ is the label vector of the

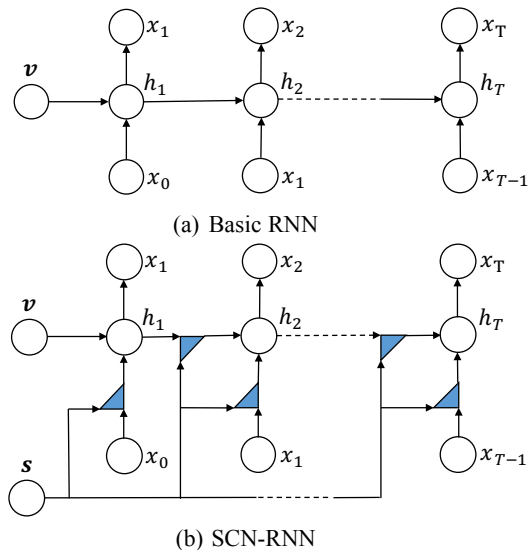


Figure 2: Comparison of our proposed model with a conventional recurrent neural network (RNN) for caption generation. \mathbf{v} and \mathbf{s} denote the visual feature and semantic feature, respectively. x_0 represents a special start-of-the-sentence token, (x_1, \dots, x_T) represents the caption, and $(\mathbf{h}_1, \dots, \mathbf{h}_T)$ denotes the RNN hidden states. Each triangle symbol represents an ensemble of tag-dependent weight matrices.

i -th image, where $y_{ik} = 1$ if the image is annotated with tag k , and $y_{ik} = 0$ otherwise. Let \mathbf{v}_i and \mathbf{s}_i represent the image feature vector and the semantic feature vector for the i -th image, the cost function to be minimized is

$$\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \left(y_{ik} \log s_{ik} + (1 - y_{ik}) \log(1 - s_{ik}) \right), \quad (3)$$

where $\mathbf{s}_i = \sigma(f(\mathbf{v}_i))$ is a K -dimensional vector with $s_i = [s_{i1}, \dots, s_{iK}]$, $\sigma(\cdot)$ is the logistic sigmoid function and $f(\cdot)$ is implemented as a multilayer perceptron (MLP).

In testing, for each input image, we compute a semantic-concept vector \mathbf{s} , formed by the probabilities of all tags, computed by the semantic-concept detection model.

3.3. SCN-RNN

The SCN extends each weight matrix of the conventional RNN to be an ensemble of a set of tag-dependent weight matrices, subjective to the probabilities that the tags are present in the image. Specifically, the SCN-RNN computes the hidden states as follows

$$\mathbf{h}_t = \sigma(\mathbf{W}(\mathbf{s})\mathbf{x}_{t-1} + \mathbf{U}(\mathbf{s})\mathbf{h}_{t-1} + \mathbf{z}), \quad (4)$$

where $\mathbf{z} = \mathbb{1}(t=1) \cdot \mathbf{C}\mathbf{v}$, and $\mathbf{W}(\mathbf{s})$ and $\mathbf{U}(\mathbf{s})$ are ensembles of tag-dependent weight matrices, subjective to the probabilities that the tags are present in the image, according to the semantic-concept vector \mathbf{s} .

Given $\mathbf{s} \in \mathbb{R}^K$, we define two weight tensors $\mathbf{W}_{\mathcal{T}} \in \mathbb{R}^{n_h \times n_x \times K}$ and $\mathbf{U}_{\mathcal{T}} \in \mathbb{R}^{n_h \times n_h \times K}$, where n_h is the number

of hidden units and n_x is the dimension of word embedding. $\mathbf{W}(\mathbf{s}) \in \mathbb{R}^{n_h \times n_x}$ and $\mathbf{U}(\mathbf{s}) \in \mathbb{R}^{n_h \times n_h}$ can be specified as

$$\mathbf{W}(\mathbf{s}) = \sum_{k=1}^K s_k \mathbf{W}_{\mathcal{T}}[k], \quad \mathbf{U}(\mathbf{s}) = \sum_{k=1}^K s_k \mathbf{U}_{\mathcal{T}}[k], \quad (5)$$

where s_k is the k -th element in \mathbf{s} ; $\mathbf{W}_{\mathcal{T}}[k]$ and $\mathbf{U}_{\mathcal{T}}[k]$ denote the k -th 2D ‘‘slice’’ of $\mathbf{W}_{\mathcal{T}}$ and $\mathbf{U}_{\mathcal{T}}$, respectively. The probability of the k -th semantic concept, s_k , is associated with a pair of RNN weight matrices $\mathbf{W}_{\mathcal{T}}[k]$ and $\mathbf{U}_{\mathcal{T}}[k]$, implicitly specifying K RNNs in total. Consequently, training such a model as defined in (4) and (5) can be interpreted as *jointly* training an ensemble of K RNNs.

Though appealing, the number of parameters is proportional to K , which is prohibitive for large K (e.g., $K = 1000$ for COCO). In order to remedy this problem, we adopt ideas from [30] to factorize $\mathbf{W}(\mathbf{s})$ and $\mathbf{U}(\mathbf{s})$ defined in (5) as

$$\mathbf{W}(\mathbf{s}) = \mathbf{W}_a \cdot \text{diag}(\mathbf{W}_b \mathbf{s}) \cdot \mathbf{W}_c, \quad (6)$$

$$\mathbf{U}(\mathbf{s}) = \mathbf{U}_a \cdot \text{diag}(\mathbf{U}_b \mathbf{s}) \cdot \mathbf{U}_c, \quad (7)$$

where $\mathbf{W}_a \in \mathbb{R}^{n_h \times n_f}$, $\mathbf{W}_b \in \mathbb{R}^{n_f \times K}$ and $\mathbf{W}_c \in \mathbb{R}^{n_f \times n_x}$. Similarly, $\mathbf{U}_a \in \mathbb{R}^{n_h \times n_f}$, $\mathbf{U}_b \in \mathbb{R}^{n_f \times K}$ and $\mathbf{U}_c \in \mathbb{R}^{n_f \times n_h}$. n_f is the number of factors. Substituting (6) and (7) into (4), we obtain our SCN with an RNN as

$$\tilde{\mathbf{x}}_{t-1} = \mathbf{W}_b \mathbf{s} \odot \mathbf{W}_c \mathbf{x}_{t-1}, \quad (8)$$

$$\tilde{\mathbf{h}}_{t-1} = \mathbf{U}_b \mathbf{s} \odot \mathbf{U}_c \mathbf{h}_{t-1}, \quad (9)$$

$$\mathbf{z} = \mathbb{1}(t=1) \cdot \mathbf{C} \mathbf{v}, \quad (10)$$

$$\mathbf{h}_t = \sigma(\mathbf{W}_a \tilde{\mathbf{x}}_{t-1} + \mathbf{U}_a \tilde{\mathbf{h}}_{t-1} + \mathbf{z}). \quad (11)$$

where \odot denotes the element-wise multiply (Hadamard) operator.

\mathbf{W}_a and \mathbf{W}_c are shared among all the captions, effectively capturing common linguistic patterns; while the diagonal term, $\text{diag}(\mathbf{W}_b \mathbf{s})$, accounts for semantic aspects of the image under test, captured by \mathbf{s} . The same analysis also holds true for $\mathbf{U}_{a,b,c}$. In this factorized model, the RNN weight matrices that correspond to each semantic concept share ‘‘structure.’’ This factorized model (termed SCN-RNN) is illustrated in Figure 2(b).

To provide further motivation for and insight into the decompositions in (6) and (7), let \mathbf{w}_{bk} represent the k th column of \mathbf{W}_b , then

$$\mathbf{W}(\mathbf{s}) = \sum_{k=1}^K s_k [\mathbf{W}_a \cdot \text{diag}(\mathbf{w}_{bk}) \cdot \mathbf{W}_c]. \quad (12)$$

A similar decomposition is manifested for $\mathbf{U}(\mathbf{s})$. The matrix $\mathbf{W}_a \cdot \text{diag}(\mathbf{w}_{bk}) \cdot \mathbf{W}_c$ may be interpreted as the k -th ‘‘slice’’ of a weight tensor, with each slice corresponding to one of the K semantic concepts (K total tensor ‘‘slices,’’ each of size $n_h \times n_x$). Hence, via the decomposition in (6) and (7), we effectively learn an ensemble of K sets of RNN parameters,

one for each semantic concept. This is efficiently done by sharing \mathbf{W}_a and \mathbf{W}_c when composing each member of the ensemble. The weight with which the k -th slice of this tensor contributes to the RNN parameters for a given image is dependent on the respective probability s_k with which the k -th semantic concept is inferred to be associated with image \mathbf{I} .

The number of parameters in the basic RNN model (see Figure 2(a)) is $n_h \cdot (n_x + n_h)$, while the number of parameters in the SCN-RNN model (see Figure 2(b)) is $n_f \cdot (n_x + 2K + 3n_h)$. In experiments, we set $n_f = n_h$. Therefore, the additional number of parameters is $2 \cdot n_h \cdot (n_h + K)$. This increased model complexity also indicates increased training/testing time.

3.4. SCN-LSTM

RNNs with LSTM units [17] have emerged as a popular architecture, due to their representational power and effectiveness at capturing long-term dependencies. We generalize the SCN-RNN model by using LSTM units. Specifically, we define $\mathbf{h}_t = \mathcal{H}(\mathbf{x}_{t-1}, \mathbf{h}_{t-1}, \mathbf{v}, \mathbf{s})$ as

$$\mathbf{i}_t = \sigma(\mathbf{W}_{ia} \tilde{\mathbf{x}}_{i,t-1} + \mathbf{U}_{ia} \tilde{\mathbf{h}}_{i,t-1} + \mathbf{z}), \quad (13)$$

$$\mathbf{f}_t = \sigma(\mathbf{W}_{fa} \tilde{\mathbf{x}}_{f,t-1} + \mathbf{U}_{fa} \tilde{\mathbf{h}}_{f,t-1} + \mathbf{z}), \quad (14)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_{oa} \tilde{\mathbf{x}}_{o,t-1} + \mathbf{U}_{oa} \tilde{\mathbf{h}}_{o,t-1} + \mathbf{z}), \quad (15)$$

$$\tilde{\mathbf{c}}_t = \sigma(\mathbf{W}_{ca} \tilde{\mathbf{x}}_{c,t-1} + \mathbf{U}_{ca} \tilde{\mathbf{h}}_{c,t-1} + \mathbf{z}), \quad (16)$$

$$\mathbf{c}_t = \mathbf{i}_t \odot \tilde{\mathbf{c}}_t + \mathbf{f}_t \odot \mathbf{c}_{t-1}, \quad (17)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t), \quad (18)$$

where $\mathbf{z} = \mathbb{1}(t=1) \cdot \mathbf{C} \mathbf{v}$. For $\star = i, f, o, c$, we define

$$\tilde{\mathbf{x}}_{\star,t-1} = \mathbf{W}_{\star b} \mathbf{s} \odot \mathbf{W}_{\star c} \mathbf{x}_{t-1}, \quad (19)$$

$$\tilde{\mathbf{h}}_{\star,t-1} = \mathbf{U}_{\star b} \mathbf{s} \odot \mathbf{U}_{\star c} \mathbf{h}_{t-1}. \quad (20)$$

Since we implement the SCN with LSTM units, we name this model SCN-LSTM. In experiments, since LSTM is more powerful than classical RNN, we only report results using SCN-LSTM.

In summary, distinct from previous image-captioning methods, our model has a unique way to utilize and combine the visual feature \mathbf{v} and semantic-concept vector \mathbf{s} extracted from an image \mathbf{I} . \mathbf{v} is fed into the LSTM to initialize the first step, which is expected to provide the LSTM an overview of the image content. While the LSTM state is initialized with the overall visual context \mathbf{v} , an ensemble of K sets of LSTM parameters is utilized when decoding, weighted by the semantic-concept vector \mathbf{s} , to generate the caption.

Model learning Given the image \mathbf{I} and associated caption \mathbf{X} , the objective function is the sum of the log-likelihood of the caption conditioned on the image representation:

$$\log p(\mathbf{X}|\mathbf{I}) = \sum_{t=1}^T p(\mathbf{x}_t | \mathbf{x}_0, \dots, \mathbf{x}_{t-1}, \mathbf{v}, \mathbf{s}). \quad (21)$$

Methods	COCO						Flickr30k				
	B-1	B-2	B-3	B-4	M	C	B-1	B-2	B-3	B-4	M
NIC [48]	0.666	0.451	0.304	0.203	–	–	0.663	0.423	0.277	0.183	–
m-RNN [29]	0.67	0.49	0.35	0.25	–	–	0.60	0.41	0.28	0.19	–
Hard-Attention [52]	0.718	0.504	0.357	0.250	0.230	–	0.669	0.439	0.296	0.199	0.185
ATT [54]	0.709	0.537	0.402	0.304	0.243	–	0.647	0.460	0.324	0.230	0.189
Att-CNN+LSTM [49]	0.74	0.56	0.42	0.31	0.26	0.94	0.73	0.55	0.40	0.28	–
LSTM-R	0.698	0.525	0.390	0.292	0.238	0.889	0.657	0.437	0.296	0.201	0.186
LSTM-T	0.716	0.546	0.411	0.312	0.250	0.952	0.691	0.483	0.336	0.232	0.202
LSTM-RT	0.724	0.555	0.419	0.316	0.252	0.970	0.706	0.486	0.339	0.235	0.204
LSTM-RT ₂	0.730	0.568	0.430	0.322	0.249	0.977	0.724	0.523	0.370	0.257	0.210
SCN-LSTM	0.728	0.566	0.433	0.330	0.257	1.012	0.735	0.530	0.377	0.265	0.218
SCN-LSTM Ensemble of 5	0.741	0.578	0.444	0.341	0.261	1.041	0.747	0.552	0.403	0.288	0.223

Table 1: Performance of the proposed model (SCN-LSTM) and other state-of-the-art methods on the COCO and Flickr30k datasets, where B-*N*, M and C are short for BLEU-*N*, METEOR and CIDEr-D scores, respectively.

The above objective corresponds to a single image-caption pair. In training, we average over all training pairs.

3.5. Extension to video captioning

The above framework can be readily extended to the task of video captioning [10, 46, 47, 56, 3, 51]. In order to effectively represent the spatiotemporal visual content of a video, we use a two-dimensional (2D) and a three-dimensional (3D) CNN to extract visual features of video frames/clips. We then perform a mean pooling process [47] over all 2D CNN features and 3D CNN features, to generate two feature vectors (one from 2D CNN features and the other from 3D CNN features). The representation of each video, v , is produced by concatenating these two features. Similarly, we also obtain the semantic-concept vector s by running the semantic-concept detector based on the video representation v . After v and s are obtained, we employ the same model proposed above directly for video-caption generation, as described in Figure 2(b).

4. Experiments

4.1. Datasets

We present results on three benchmark datasets: COCO [27], Flickr30k [55] and Youtube2Text [5]. COCO and Flickr30k are for image captioning, containing 123287 and 31783 images, respectively. Each image is annotated with at least 5 captions. We use the same pre-defined splits as [20] for all the datasets: on Flickr30k, 1000 images for validation, 1000 for test, and the rest for training; and for COCO, 5000 images are used for both validation and testing. We further tested our model on the official COCO test set consisting of 40775 images (human-generated captions for this split are not publicly available), and evaluated our model on the COCO evaluation server. We also follow the publicly available code [20] to preprocess the captions, yielding vocabulary sizes of 8791 and 7414 for COCO and Flickr30k,

respectively.

Youtube2Text is used for video captioning, which contains 1970 Youtube clips, and each video is annotated with around 40 sentences. We use the same splits as provided in [47], with 1200 videos for training, 100 videos for validation, and 670 videos for testing. We convert all captions to lower case and remove the punctuation, yielding vocabulary size of 12594 for Youtube2Text.

4.2. Training procedure

For image representation, we take the output of the 2048-way *pool5* layer from ResNet-152 [16], pretrained on the ImageNet dataset [36]. For video representation, in addition to using the 2D ResNet-152 to extract features on each video frame, we also utilize a 3D CNN (C3D) [43] to extract features on each video. The C3D is pretrained on Sports-1M video dataset [21], and we take the output of the 4096-way *fc7* layer from C3D as the video representation. We consider the RGB frames of videos as input, with 2 frames per second. Each video frame is resized as 112×112 and 224×224 for the C3D and ResNet-152 feature extractor, respectively. The C3D feature extractor is applied on video clips of length 16 frames (as in [21]) with an overlap of 8 frames.

We use the procedure described in Section 3.2 for semantic concept detection. The semantic-concept vocabulary size is determined to reflect the complexity of the dataset, which is set to 1000, 200 and 300 for COCO, Flickr30k and Youtube2Text, respectively. Since Youtube2Text is a relatively small dataset, we found that it is very difficult to train a reliable semantic-concept detector using the Youtube2Text dataset alone, due to its limited amount of data. In experiments, we utilize additional training data from COCO.

For model training, all the parameters in the SCN-LSTM are initialized from a uniform distribution in $[-0.01, 0.01]$. All bias terms are initialized to zero. Word embedding vectors are initialized with the publicly available *word2vec* vectors [31]. The embedding vectors of words not present in the

Model	BLEU-1		BLEU-2		BLEU-3		BLEU-4		METEOR		ROUGE-L		CIDEr-D	
	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40	c5	c40
SCN-LSTM	0.740	0.917	0.575	0.839	0.436	0.739	0.331	0.631	0.257	0.348	0.543	0.696	1.003	1.013
ATT	0.731	0.900	0.565	0.815	0.424	0.709	0.316	0.599	0.250	0.335	0.535	0.682	0.943	0.958
OV	0.713	0.895	0.542	0.802	0.407	0.694	0.309	0.587	0.254	0.346	0.530	0.682	0.943	0.946
MSR Cap	0.715	0.907	0.543	0.819	0.407	0.710	0.308	0.601	0.248	0.339	0.526	0.680	0.931	0.937

Table 2: Comparison to published state-of-the-art image captioning models on the blind test set as reported by the COCO test server. SCN-LSTM is our model. ATT refers to ATT VC [54], OV refers to OriolVinyals [48], and MSR Cap refers to MSR Captivator [9].

Model	B-4	M	C
S2VT [46]	–	0.292	–
LSTM-E [32]	0.453	0.310	–
GRU-RCN [3]	0.479	0.311	0.678
h-RNN [56]	0.499	0.326	0.658
LSTM-R	0.448	0.310	0.640
LSTM-C	0.445	0.309	0.644
LSTM-CR	0.469	0.317	0.688
LSTM-T	0.473	0.324	0.699
LSTM-CRT	0.475	0.316	0.647
LSTM-CRT ₂	0.469	0.326	0.706
SCN-LSTM	0.502	0.334	0.770
SCN-LSTM Ensemble of 5	0.511	0.335	0.777

Table 3: Results on BLEU-4 (B-4), METEOR (M) and CIDEr-D (C) metrics compared to other state-of-the-art results and baselines on Youtube2Text.

pretrained set are initialized randomly. The number of hidden units and the number of factors in SCN-LSTM are both set to 512 and we use mini-batches of size 64. The maximum number of epochs we run for all the three datasets is 20. Gradients are clipped if the norm of the parameter vector exceeds 5 [40]. We do not perform any dataset-specific tuning and regularization other than dropout [57] and early stopping on validation sets. The Adam algorithm [22] with learning rate 2×10^{-4} is utilized for optimization. All experiments are implemented in Theano [42]¹.

In testing, we use beam search for caption generation, which selects the top- k best sentences at each time step and considers them as the candidates to generate new top- k best sentences at the next time step. We set the beam size to $k = 5$ in experiments.

4.3. Evaluation

The widely used BLEU [33], METEOR [4], ROUGE-L [26], and CIDEr-D [45] metrics are reported in our quantitative evaluation of the performance of the proposed model and baselines in the literature. All the metrics are computed by using the code released by the COCO evaluation server [6]. For COCO and Flickr30k datasets, besides comparing to results reported in previous work, we also re-implemented strong baselines for comparison. The results of image captioning are presented in Table 1. The models we implemented are as follows.

¹Code is publicly available at https://github.com/zhegan27/Semantic_Compositional_Nets.

1. *LSTM-R / LSTM-T / LSTM-RT*: R , T , RT denotes using different features. Specifically, R denotes ResNet visual feature vector, T denotes Tags (*i.e.*, the semantic-concept vector), and RT denotes the concatenation of R and T . The features are fed into a standard LSTM decoder only at the initial time step. In particular, *LSTM-T* is the model proposed in [49].
2. *LSTM-RT₂*: The ResNet feature vector is sent to a standard LSTM decoder at the first time step, while the tag vector is sent to the LSTM decoder at every time step in addition to the input word. This model is similar to [54] without using semantic attention. This is the model closest to ours, which provides a direct comparison to our proposed model.
3. *SCN-LSTM*: This is the model presented in Section 3.4.

For video captioning experiments, we use the same notation. For example, *LSTM-C* means we leverage the C3D feature for caption generation.

4.4. Quantitative results

Performance on COCO and Flickr30k We first present results on the task of image captioning, summarized in Table 1. The use of tags (*LSTM-T*) provides better performance than leveraging visual features alone (*LSTM-R*). Combining both tags and visual features further enhances performance, as expected. Compared with only feeding the tags into the LSTM at the initial time step (*LSTM-RT*), *LSTM-RT₂* yields better results, since it takes as input the tag feature at each time step. Further, the direct comparison between *LSTM-RT₂* and *SCN-LSTM* demonstrates the advantage of our proposed model, indicating that our approach is a better method to fuse semantic concepts into the LSTM.

We also report results averaging an ensemble of 5 identical *SCN-LSTM* models trained with different initializations, which is a common strategy adopted widely [54] (note that now we employ ensembles in two ways: an ensemble of LSTM parameters linked to tags, and an overarching ensemble atop the entire model). We obtain state-of-the-art results on both COCO and Flickr30k datasets. Remarkably, we improve the state-of-the-art BLEU-4 score by 3.1 points on COCO.

Performance on COCO test server We also evaluate the proposed *SCN-LSTM* model by uploading results to the on-




	Tags: dog (1), grass (0.996), laying (0.97), outdoor (0.943), next (0.788), sitting (0.651), lying (0.542), white (0.507)		Tags: road (1), double (0.999), bus (0.996), red (0.996), street (0.926), building (0.859), driving (0.796)		Tags: zebra (1), animal (0.985), mammal (0.948), dirt (0.937), grass (0.902), standing (0.878), group (0.848), field (0.709)
Caption generated by our model: <i>a dog laying on the ground next to a frisbee</i> Semantic composition: 1. Replace “dog” with “cat”: <i>a white cat laying on the ground</i> 2. Replace “grass” with “bed”: <i>a white dog laying on top of a bed</i> 3. Replace “grass” with “laptop”: <i>a dog laying on the ground next to a laptop</i>	Caption generated by our model: <i>a red double decker bus driving down a street</i> Semantic composition: 1. Replace “red” with “blue”: <i>a blue double decker bus driving down a street</i> 2. Replace “bus” with “train”: <i>a red train traveling down a city street</i> 3. Replace “road” and “street” with “ocean”: <i>a red bus is driving in the ocean</i>	Caption generated by our model: <i>a herd of zebra standing on top of a dirt field</i> Semantic composition: 1. Replace “zebra” with “horse”: <i>a group of horses standing in a field</i> 2. Replace “standing” with “running”: <i>a herd of zebra running across a dirt field</i> 3. Replace “field” with “snow”: <i>a group of zebras standing in the snow</i>			

Figure 3: Illustration of semantic composition. Our model can adjust the caption smoothly as the semantic concepts are modified.




	Tags: indoor (0.952), dog (0.828), sitting (0.647), stuffed (0.602), white (0.544), next (0.527), laying (0.509), cat (0.402)		Tags: snow(1), outdoor (0.992), covered (0.847), nature (0.812), skiing (0.61), man (0.451), pile (0.421), building (0.369)		Tags: person (1), cabinet (0.931), man (0.906), shelf (0.771), table (0.707), front (0.683), holding (0.662), food (0.587)
Generated captions: SCN-LSTM-T: a dog laying on top of a stuffed animal SCN-LSTM: a teddy bear laying on top of a stuffed animal	Generated captions: SCN-LSTM-T: a person that is standing in the snow SCN-LSTM: a stop sign is covered in the snow	Generated captions: SCN-LSTM-T: a man sitting at a table with a plate of food SCN-LSTM: a man is holding a glass of wine			

Figure 4: Detected tags and sentence generation results on COCO. The output captions are generated by: 1) SCN-LSTM, and 2) SCN-LSTM-T, a SCN-LSTM model without the visual feature inputs, *i.e.*, with only tag inputs.

line COCO test server. Table 2 shows the comparison to the published state-of-the-art image captioning models on the blind test set as reported by the COCO test server. We include the models that have been published and perform at top-3 in the table. Compared to these methods, our proposed SCN-LSTM model achieves the best performance across all the evaluation metrics on both c5 and c40 testing sets.²

Performance on Youtube2Text Results on video captioning are presented in Table 3. The SCN-LSTM achieves significantly better results over all competing methods in all metrics, especially in CIDEr-D. For self-comparison, it is also worth noting that our model improves over LSTM-CRT₂ by a substantial margin. Again, using an overarching ensemble further enhances performance.

4.5. Qualitative analysis

Figure 3 shows three examples to illustrate the semantic composition on caption generation. Our model properly describes the image content by using the correctly detected tags. By manually replacing specific tags, our model can adjust the caption smoothly. For example, in the left image, by replacing the tag “grass” with “bed”, our model imagines “a dog laying on top of a bed”. Our model is also able to

generate novel captions that are highly unlikely to occur in real life. For instance, in the middle image, by replacing the tag “road” and “street” with “ocean”, our model imagines “a bus driving in the ocean”; in the right image, by replacing the tag “field” with “snow”, our model dreams “a group of zebras standing in the snow”.

SCN not only picks up the tags well (and imagines the corresponding scenes), but also selects the right functional words for different concepts to form syntactically correct caption. As illustrated in sentence 6 of Figure 1b, by replacing the tag “baby” with “girl”, the generated captions not only changes “a baby” to “a little girl”, but more importantly, changes “in its mouth” to “in her mouth”. In addition, the SCN also infers the underlying semantic relatedness between different tags. As illustrated in sentence 4 of Figure 1b, when only switching on the tag “mouth”, the generated caption becomes “a man with a toothbrush”, indicating the semantic closeness between “mouth”, “man” and “toothbrush”. By further switching on “baby”, we generate a more detailed description “a baby brushing its teeth”.

The above analysis shows the importance of tags in generating captions. However, SCN generates captions using both semantic concepts and the global visual feature vector. The language model learns to assemble semantic concepts (weighted by their likelihood), in consideration of the global

²Please check <https://competitions.codalab.org/competitions/3221#results> for the most recent results.




	Tags: book (1), shelf (1), table (0.965), sitting (0.955), person (0.955), library (0.908), room (0.829), front (0.464)		Tags: person (1), table (0.822), wine (0.672), people (0.657), man (0.62), woman (0.601), sitting (0.502), holding (0.494)		Tags: grass (1), red (0.982), fire (0.953), hydrant (0.852), dog (0.723), standing (0.598), next (0.476), field (0.341)
Generated captions: LSTM-R: a young girl is playing a video game LSTM-RT₂: a group of people sitting at a table SCN-LSTM: two women sitting at a table in a library		Generated captions: LSTM-R: a group of people standing around a table eating food LSTM-RT₂: a group of people sitting at a table SCN-LSTM: a man pouring wine into a wine glass		Generated captions: LSTM-R: a dog that is sitting on the ground LSTM-RT₂: a dog standing next to a fire hydrant SCN-LSTM: a dog standing next to a red fire hydrant	

Figure 5: Detected tags and sentences generation results on COCO. The output captions are generated by: 1) LSTM-R, 2) LSTM-RT₂, and 3) our SCN-LSTM.




		
Tags: man (0.806), game (0.629), playing (0.577), ball (0.555), football (0.522), men (0.435), running (0.386), soccer (0.252)	Tags: man (0.976), person (0.881), guy (0.603), boy (0.456), gun (0.41), shooting (0.269), movie (0.232), standing (0.209)	Tags: man (0.808), person (0.603), street (0.522), road (0.512), doing (0.424), riding (0.405), running (0.397), walking (0.296)
Generated captions: LSTM-CR: a man is running LSTM-CRT₂: a man is hitting a goal SCN-LSTM: the men are playing soccer	Generated captions: LSTM-CR: a man is playing a guitar LSTM-CRT₂: a man is playing with a machine SCN-LSTM: a man is shooting a gun	Generated captions: LSTM-CR: a man is walking LSTM-CRT₂: a man is dancing SCN-LSTM: a man is running

Figure 6: Detected tags and sentence generation results on Youtube2Text. The output captions are generated by: 1) LSTM-CR, 2) LSTM-CRT₂, and 3) our SCN-LSTM.

visual information, into a coherent meaningful sentence that captures the overall meaning of the image. In order to demonstrate the importance of visual feature vectors, we train another SCN-LSTM-T model, which is a SCN-LSTM model without the visual feature inputs, *i.e.*, with only tag inputs. As shown in the first example of Figure 4, the image tagger detects “dog” with high probability. Using only tag inputs, SCN-LSTM-T can only generate the wrong caption “a dog laying on top of a stuffed animal”. With additional visual feature inputs, our SCN-LSTM model correctly replaces “dog” with “teddy bear”.

We further present examples of generated captions on COCO with various other methods in Figure 5, along with the detected tags. As can be seen, our model often generates more reasonable captions than *LSTM-R*, due to the use of high-level semantic concepts. For example, in the first image, *LSTM-R* outputs an irrelevant caption to the image, while the detection of “table” and “library” helps our model to generate more sensible caption. Further, although both our model and *LSTM-RT₂* utilize detected tags for caption generation, our model often depicts the image content more comprehensively; *LSTM-RT₂* has a larger potential to miss important details in the image. For instance, in the 3rd image, the tag “red” appears in the caption generated by our

model, which is missed by *LSTM-RT₂*. This observation might be due to the fact that the SCN provides a better approach to fuse tag information into the process of caption generation. Similiar observations can also be found in the video captioning experiments, as demonstrated in Figure 6.

5. Conclusion

We have presented Semantic Compositional Network (SCN), a new framework to effectively compose the individual semantic meaning of tags for visual captioning. The SCN extends each weight matrix of the conventional LSTM to be a three-way matrix product, with one of these matrices dependent on the inferred tags. Consequently, the SCN can be viewed an ensemble of tag-dependent LSTM bases, with the contribution of each LSTM basis unit proportional to the likelihood that the tag is present in the image. Experiments conducted on three visual captioning datasets validate the superiority of the proposed approach.

Acknowledgements Most of this work was done when the first author was an intern at Microsoft Research. This work was also supported in part by ARO, DARPA, DOE, NGA, ONR and NSF.

References

- [1] L. Anne Hendricks, S. Venugopalan, M. Rohrbach, R. Mooney, K. Saenko, and T. Darrell. Deep compositional captioning: Describing novel object categories without paired training data. In *CVPR*, 2016. 2
- [2] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015. 1
- [3] N. Ballas, L. Yao, C. Pal, and A. Courville. Delving deeper into convolutional networks for learning video representations. In *ICLR*, 2016. 2, 5, 6
- [4] S. Banerjee and A. Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop*, 2005. 6
- [5] D. L. Chen and W. B. Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL*, 2011. 5
- [6] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollár, and C. L. Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv:1504.00325*, 2015. 6
- [7] X. Chen and C. Lawrence Zitnick. Mind’s eye: A recurrent visual representation for image caption generation. In *CVPR*, 2015. 2
- [8] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*, 2014. 1
- [9] J. Devlin, H. Cheng, H. Fang, S. Gupta, L. Deng, X. He, G. Zweig, and M. Mitchell. Language models for image captioning: The quirks and what works. In *ACL*, 2015. 2, 6
- [10] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 1, 2, 5
- [11] J. Dong, X. Li, W. Lan, Y. Huo, and C. G. Snoek. Early embedding and late reranking for video captioning. In *ACMMM*, 2016. 2
- [12] J. L. Elman. Finding structure in time. *Cognitive science*, 1990. 2
- [13] H. Fang, S. Gupta, F. Iandola, R. K. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt, et al. From captions to visual concepts and back. In *CVPR*, 2015. 1, 2, 3
- [14] C. Gan, Z. Gan, X. He, J. Gao, and L. Deng. Stylenet: Generating attractive visual captions with styles. In *CVPR*, 2017. 2
- [15] C. Gan, T. Yang, and B. Gong. Learning attributes equals multi-source domain generalization. In *CVPR*, 2016. 2
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [17] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 1997. 1, 2, 4
- [18] X. Jia, E. Gavves, B. Fernando, and T. Tuytelaars. Guiding long-short term memory for image caption generation. In *ICCV*, 2015. 2
- [19] J. Jin, K. Fu, R. Cui, F. Sha, and C. Zhang. Aligning where to see and what to tell: image caption with region-based attention and scene factorization. *arXiv:1506.06272*, 2015. 2
- [20] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. 1, 2, 5
- [21] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 5
- [22] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [23] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Multimodal neural language models. In *ICML*, 2014. 2
- [24] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv:1411.2539*, 2014. 2
- [25] R. Kiros, R. Zemel, and R. R. Salakhutdinov. A multiplicative model for learning distributed text-based attribute representations. In *NIPS*, 2014. 2
- [26] C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. In *ACL workshop*, 2004. 6
- [27] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5
- [28] C. Liu, J. Mao, F. Sha, and A. Yuille. Attention correctness in neural image captioning. *arXiv:1605.09553*, 2016. 2
- [29] J. Mao, W. Xu, Y. Yang, J. Wang, Z. Huang, and A. Yuille. Deep captioning with multimodal recurrent neural networks (m-rnn). In *ICLR*, 2015. 1, 2, 5
- [30] R. Memisevic and G. Hinton. Unsupervised learning of image transformations. In *CVPR*, 2007. 2, 4
- [31] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013. 5
- [32] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui. Jointly modeling embedding and translation to bridge video and language. In *CVPR*, 2016. 1, 2, 6
- [33] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002. 6
- [34] Y. Pu, Z. Gan, R. Henao, X. Yuan, C. Li, A. Stevens, and L. Carin. Variational autoencoder for deep learning of images, labels and captions. In *NIPS*, 2016. 2
- [35] Y. Pu, M. R. Min, Z. Gan, and L. Carin. Adaptive feature abstraction for translating video to language. *arXiv:1611.07837*, 2016. 2
- [36] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 5
- [37] R. Socher, A. Karpathy, Q. V. Le, C. D. Manning, and A. Y. Ng. Grounded compositional semantics for finding and describing images with sentences. *TACL*, 2014. 2
- [38] J. Song, Z. Gan, and L. Carin. Factored temporal sigmoid belief networks for sequence learning. In *ICML*, 2016. 2
- [39] I. Sutskever, J. Martens, and G. E. Hinton. Generating text with recurrent neural networks. In *ICML*, 2011. 2
- [40] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014. 1, 6
- [41] G. W. Taylor and G. E. Hinton. Factored conditional restricted boltzmann machines for modeling motion style. In *ICML*, 2009. 2

- [42] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv:1605.02688*, 2016. 6
- [43] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 5
- [44] K. Tran, X. He, L. Zhang, J. Sun, C. Carapcea, C. Thrasher, C. Buehler, and C. Sienkiewicz. Rich image captioning in the wild. In *CVPR Workshops*, 2016. 3
- [45] R. Vedantam, C. Lawrence Zitnick, and D. Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015. 6
- [46] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko. Sequence to sequence-video to text. In *ICCV*, 2015. 1, 2, 5, 6
- [47] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. In *NAACL*, 2015. 1, 2, 5
- [48] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015. 1, 2, 5, 6
- [49] Q. Wu, C. Shen, L. Liu, A. Dick, and A. v. d. Hengel. What value do explicit high level concepts have in vision to language problems? In *CVPR*, 2016. 1, 2, 3, 5, 6
- [50] Y. Wu, S. Zhang, Y. Zhang, Y. Bengio, and R. Salakhutdinov. On multiplicative integration with recurrent neural networks. In *NIPS*, 2016. 2
- [51] J. Xu, T. Mei, T. Yao, and Y. Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, 2016. 5
- [52] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 1, 2, 5
- [53] Z. Yang, Y. Yuan, Y. Wu, R. Salakhutdinov, and W. W. Cohen. Review networks for caption generation. In *NIPS*, 2016. 2
- [54] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo. Image captioning with semantic attention. In *CVPR*, 2016. 1, 2, 5, 6
- [55] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *TACL*, 2014. 5
- [56] H. Yu, J. Wang, Z. Huang, Y. Yang, and W. Xu. Video paragraph captioning using hierarchical recurrent neural networks. In *CVPR*, 2016. 1, 2, 5, 6
- [57] W. Zaremba, I. Sutskever, and O. Vinyals. Recurrent neural network regularization. *arXiv:1409.2329*, 2014. 6

A. More results on image captioning













	Tags: polar (0.999), rock (0.998), animal (0.997), bear (0.993), mammal (0.988), zoo (0.881), white (0.779), large (0.748)		Tags: refrigerator (0.992), food (0.976), open (0.97), cabinet (0.953), shelf (0.582), filled (0.45), door (0.426), lots (0.329)		Tags: person (0.925), building (0.839), people (0.787), umbrella (0.779), group (0.704), child (0.519), standing (0.369), holding (0.272)
Generated captions: LSTM-R: a polar bear standing on top of a rock LSTM-RT₂: a polar bear standing on a rock SCN-LSTM: a large white polar bear standing on a rock		Generated captions: LSTM-R: a display case filled with lots of food LSTM-RT₂: a shelf full of different kinds of food SCN-LSTM: a refrigerator filled with lots of food and drinks		Generated captions: LSTM-R: a group of people standing next to each other LSTM-RT₂: a group of people standing in front of an umbrella SCN-LSTM: a group of people standing in the rain with umbrellas	
	Tags: bicycle (1), parked (0.923), next (0.889), group (0.829), sidewalk (0.783), many (0.698), lot (0.611), rack (0.596)		Tags: food (0.939), oranges (0.839), fruit (0.836), slice (0.792), sliced (0.783), orange (0.764), plate (0.759), table (0.704)		Tags: clock (1), building (0.999), large (0.902), station (0.876), mounted (0.644), sitting (0.621), tower (0.574), building (0.418)
Generated captions: LSTM-R: a group of motorcycles parked next to each other LSTM-RT₂: a row of bikes parked in a row SCN-LSTM: a bunch of bikes parked in a park lot		Generated captions: LSTM-R: a bowl of fruit sitting on top of a table LSTM-RT₂: a bunch of oranges sitting on a table SCN-LSTM: a bunch of oranges sitting on a plate		Generated captions: LSTM-R: a clock on the wall of a building LSTM-RT₂: a clock on the side of a building SCN-LSTM: a large clock mounted to the side of a building	
	Tags: dog (1), water (0.998), beach (0.805), standing (0.666), walking (0.451), next (0.435), ocean (0.301), white (0.225)		Tags: water (0.985), beach (0.975), ocean (0.655), next (0.493), shore (0.324), sand (0.288), sandy (0.209), bench(0.204)		Tags: bench (0.997), fence (0.98), park (0.974), grass (0.877), sitting (0.771), wooden (0.582), next (0.511), green (0.377)
Generated captions: LSTM-R: a dog that is playing with a frisbee LSTM-RT₂: a couple of dogs standing on a beach SCN-LSTM: a white dog walking on a beach		Generated captions: LSTM-R: a bench that is sitting on the beach LSTM-RT₂: a person sitting on a bench on a beach SCN-LSTM: a wooden bench sitting on top of a sandy beach		Generated captions: LSTM-R: a park bench sitting in the middle of a forest LSTM-RT₂: a park bench sitting on a park bench SCN-LSTM: a wooden bench sitting in the middle of a park	
	Tags: road (0.958), street (0.911), green (0.856), sign (0.601), traffic (0.549), car (0.401), truck (0.382), city (0.374)		Tags: person (0.958), woman (0.728), sitting (0.708), bench (0.394), people (0.381), next (0.371), group (0.361), front (0.311)		Tags: person (0.932), man (0.787), young (0.458), black (0.439), white (0.43), jumping (0.342), riding (0.242), trick (0.156)
Generated captions: LSTM-R: a bus that is driving down the road LSTM-RT₂: a bus parked on the side of a road SCN-LSTM: a green bus driving down a city street		Generated captions: LSTM-R: a couple of women standing next to each other LSTM-RT₂: a couple of people sitting on a toilet SCN-LSTM: a group of people sitting on a bench		Generated captions: LSTM-R: a man sitting on top of a wooden bench LSTM-RT₂: a man riding a skateboard down a street SCN-LSTM: a black and white photo of a skateboarder doing a trick	

Figure 7: Detected tags and sentences generation results on COCO. The output captions are generated by: 1) LSTM-R, 2) LSTM-RT₂, and 3) our SCN-LSTM.

B. More results on video captioning










		
<p>Tags: playing (0.694), animal (0.673), baby (0.63), person (0.471), eating (0.419), something (0.333), food (0.329), hand (0.311)</p>	<p>Tags: man (0.807), person (0.733), car (0.442), driving (0.39), playing (0.382), road (0.365), moving (0.189), pushing (0.129)</p>	<p>Tags: woman (0.88), girl (0.732), lady (0.699), making (0.516), something (0.501), water (0.267), glass (0.244), drinking (0.204)</p>
<p>Generated captions: LSTM-CR: a person is eating LSTM-CRT₂: a person is holding a small animal SCN-LSTM: a small animal is eating</p>	<p>Generated captions: LSTM-CR: a man is doing a wheelie LSTM-CRT₂: a man is riding a bike SCN-LSTM: a man is pushing a car</p>	<p>Generated captions: LSTM-CR: a woman is pouring sugar in a glass LSTM-CRT₂: a woman is pouring water SCN-LSTM: a woman is drinking something</p>
		
<p>Tags: man (0.635), woman (0.545), riding (0.541), person (0.465), water (0.465), girl (0.4), doing (0.387), horse (0.132)</p>	<p>Tags: man (0.843), person (0.774), doing (0.393), playing (0.385), open (0.298), gun (0.283), shooting (0.276), field (0.259)</p>	<p>Tags: man (0.958), song (0.869), stage (0.866), singing (0.859), men (0.845), music (0.826), playing (0.762), guitar (0.759)</p>
<p>Generated captions: LSTM-CR: a girl is riding a horse LSTM-CRT₂: a woman is riding a horse SCN-LSTM: a man is riding a horse</p>	<p>Generated captions: LSTM-CR: a girl is firing a gun LSTM-CRT₂: a girl is shooting SCN-LSTM: a man is shooting a gun</p>	<p>Generated captions: LSTM-CR: a group of people are dancing on stage LSTM-CRT₂: a man is dancing on stage SCN-LSTM: a band is performing on stage</p>
		
<p>Tags: doing (0.616), boy (0.557), room (0.554), playing (0.51), floor (0.493), dancing (0.491), dance (0.361), kid (0.281)</p>	<p>Tags: woman (0.829), girl (0.743), doing (0.593), using (0.408), makeup (0.211), applying (0.2), face (0.171), hand (0.139)</p>	<p>Tags: playing (0.776), dog (0.625), floor (0.423), trying (0.399), woman (0.356), running (0.293), puppy (0.202), toy (0.182)</p>
<p>Generated captions: LSTM-CR: a baby is walking LSTM-CRT₂: a baby is dancing SCN-LSTM: a boy is dancing</p>	<p>Generated captions: LSTM-CR: a girl is singing LSTM-CRT₂: a woman is playing SCN-LSTM: a woman is plucking her eyebrow</p>	<p>Generated captions: LSTM-CR: a group of girls are playing with a toy LSTM-CRT₂: the children are playing SCN-LSTM: a dog is playing with a toy</p>

Figure 8: Detected tags and sentence generation results on Youtube2Text. The output captions are generated by: 1) LSTM-CR, 2) LSTM-CRT₂, and 3) our SCN-LSTM.

C. More results for Figure 4







	<p>Tags: outdoor (1), elephant (0.995), animal (0.988), grass (0.962), standing (0.89), rock (0.781), zoo (0.682), enclosure (0.619)</p>		<p>Tags: indoor (0.966), table (0.919), food (0.849), kitchen (0.714), sitting (0.545), counter (0.436), top (0.285), doughnut (0.251)</p>		<p>Tags: outdoor (0.998), building (0.996), man (0.613), front (0.434), standing (0.333), woman (0.255), walking (0.249), next (0.247)</p>
<p>Generated captions: SCN-LSTM-T: a couple of elephants standing next to each other SCN-LSTM: a large elephant standing next to a tree</p>		<p>Generated captions: SCN-LSTM-T: a kitchen with a lot of food on it SCN-LSTM: a bunch of doughnuts sitting on top of a counter</p>		<p>Generated captions: SCN-LSTM-T: a man standing in front of a building SCN-LSTM: a statue of a man standing next to a building</p>	
	<p>Tags: table (0.997), indoor (0.893), chair (0.876), room (0.692), sitting (0.583), window (0.58), wooden (0.542), small (0.344)</p>		<p>Tags: fence (1), giraffe (0.994), animal (0.921), wooden (0.677), fenced (0.66), standing (0.592), next (0.493), zoo (0.442)</p>		<p>Tags: grass (1), outdoor (0.992), giraffe (0.985), mammal (0.98), animal (0.978), field (0.93), eating (0.558), standing (0.508)</p>
<p>Generated captions: SCN-LSTM-T: a dining room with a table and chairs SCN-LSTM: a wooden table with a laptop on it</p>		<p>Generated captions: SCN-LSTM-T: a giraffe standing next to a wooden fence SCN-LSTM: a couple of giraffe standing next to each other</p>		<p>Generated captions: SCN-LSTM-T: a giraffe standing on top of a lush green field SCN-LSTM: a giraffe is eating grass in a field</p>	

Figure 9: Detected tags and sentence generation results on COCO. The output captions are generated by: 1) SCN-LSTM, and 2) SCN-LSTM-T, a SCN-LSTM model without the visual feature inputs, *i.e.*, with only tag inputs.