

Auto-Encoding Scene Graphs for Image Captioning

Xu Yang, Kaihua Tang, Hanwang Zhang, Jianfei Cai
School of Computer Science and Engineering,
Nanyang Technological University,

{s170018@e,kaihua001@e}.ntu.edu.sg, {hanwangzhang@,ASJFCai@}ntu.edu.sg

Abstract

We propose Scene Graph Auto-Encoder (SGAE) that incorporates the language inductive bias into the encoder-decoder image captioning framework for more human-like captions. Intuitively, we humans use the inductive bias to compose collocations and contextual inference in discourse. For example, when we see the relation “person on bike”, it is natural to replace “on” with “ride” and infer “person riding bike on a road” even the “road” is not evident. Therefore, exploiting such bias as a language prior is expected to help the conventional encoder-decoder models less likely overfit to the dataset bias and focus on reasoning. Specifically, we use the scene graph — a directed graph (\mathcal{G}) where an object node is connected by adjective nodes and relationship nodes — to represent the complex structural layout of both image (\mathcal{I}) and sentence (\mathcal{S}). In the textual domain, we use SGAE to learn a dictionary (\mathcal{D}) that helps to reconstruct sentences in the $\mathcal{S} \rightarrow \mathcal{G} \rightarrow \mathcal{D} \rightarrow \mathcal{S}$ pipeline, where \mathcal{D} encodes the desired language prior; in the vision-language domain, we use the shared \mathcal{D} to guide the encoder-decoder in the $\mathcal{I} \rightarrow \mathcal{G} \rightarrow \mathcal{D} \rightarrow \mathcal{S}$ pipeline. Thanks to the scene graph representation and shared dictionary, the inductive bias is transferred across domains in principle. We validate the effectiveness of SGAE on the challenging MS-COCO image captioning benchmark, e.g., our SGAE-based single-model achieves a new state-of-the-art 127.8 CIDEr-D on the Karpathy split, and a competitive 125.5 CIDEr-D (c40) on the official server even compared to other ensemble models. Code has been made available at: <https://github.com/yangxuntu/SGAE>.

1. Introduction

Modern image captioning models employ an end-to-end encoder-decoder framework [10, 30, 29, 2, 27], i.e., the encoder encodes an image into vector representations and then the decoder decodes them into a language sequence. Since its invention inspired from neural machine translation [3], this framework has experienced several significant upgrades

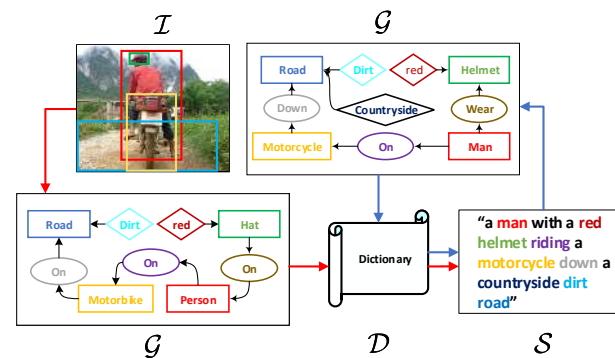


Figure 1. Illustration of auto-encoding scene graphs (blue arrows) into the conventional encoder-decoder framework for image captioning (red arrows), where the language inductive bias is encoded in the trainable shared dictionary. Word colors correspond to nodes in image and sentence scene graphs.

such as the top-bottom [50] and bottom-up [2] visual attentions for dynamic encoding, and the reinforced mechanism for sequence decoding [38, 9, 35]. However, a ubiquitous problem has never been substantially resolved: when we feed an unseen image scene into the framework, we usually get a simple and trivial caption about the salient objects such as “there is a dog on the floor”, which is no better than just a list of object detection [29]. This situation is particularly embarrassing in front of the booming “mid-level” vision techniques nowadays: we can already detect and segment almost everything in an image [12, 17, 36].

We humans are good at telling sentences about a visual scene. Not surprisingly, cognitive evidences [32] show that the visually grounded language generation is not end-to-end and largely attributed to the “high-level” symbolic reasoning, that is, once we abstract the scene into symbols, the generation will be almost *disentangled* from the visual perception. For example, as shown in Figure 1, from the scene abstraction “helmet-on-human” and “road dirty”, we can say “a man with a helmet in countryside” by using the common sense knowledge like “country road is dirty”. In fact, such collocations and contextual inference in human

language can be considered as the *inductive bias* that is apprehended by us from everyday practice, which makes us performing better than machines in high-level reasoning [22, 5]. However, the direct exploitation of the inductive bias, *e.g.*, early template/rule-based caption models [20, 8], is well-known ineffective compared to the encoder-decoder ones, due to the large gap between visual perception and language composition.

In this paper, we propose to incorporate the inductive bias of language generation into the encoder-decoder framework for image captioning, benefiting from the complementary strengths of both symbolic reasoning and end-to-end multi-modal feature mapping. In particular, we use scene graphs [13, 47] to bridge the gap between the two worlds. A scene graph (\mathcal{G}) is a unified representation that connects 1) the objects (or entities), 2) their attributes, and 3) their relationships in an image (\mathcal{I}) or a sentence (\mathcal{S}) by directed edges. Thanks to the recent advances in spatial Graph Convolutional Networks (GCNs) [31, 24], we can embed the graph structure into vector representations, which can be seamlessly integrated into the encoder-decoder. Our key insight is that the vector representations are expected to transfer the inductive bias from the pure language domain to the vision-language domain.

Specifically, to encode the language prior, we propose the Scene Graph Auto-Encoder (SGAE) that is a sentence self-reconstruction network in the $\mathcal{S} \rightarrow \mathcal{G} \rightarrow \mathcal{D} \rightarrow \mathcal{S}$ pipeline, where \mathcal{D} is a trainable dictionary for the re-encoding purpose of the node features, the $\mathcal{S} \rightarrow \mathcal{G}$ module is a fixed off-the-shelf scene graph language parser [1], and the $\mathcal{D} \rightarrow \mathcal{S}$ is a trainable RNN-based language decoder [2]. Note that \mathcal{D} is the “juice” — the language inductive bias — we extract from training SGAE. By sharing \mathcal{D} in the encoder-decoder training pipeline: $\mathcal{I} \rightarrow \mathcal{G} \rightarrow \mathcal{D} \rightarrow \mathcal{S}$, we can incorporate the language prior to guide the end-to-end image captioning. In particular, the $\mathcal{I} \rightarrow \mathcal{G}$ module is a visual scene graph detector [56] and we introduce a multi-modal GCN for the $\mathcal{G} \rightarrow \mathcal{D}$ module in the captioning pipeline, to complement necessary visual cues that are missing due to the imperfect visual detection. Interestingly, \mathcal{D} can be considered as a working memory [45] that helps to re-key the encoded nodes from \mathcal{I} or \mathcal{S} to a more generic representation with smaller domain gaps. More motivations and the incarnation of \mathcal{D} will be discussed in Section 4.3.

We implement the proposed SGAE-based captioning model by using the recently released visual encoder [37] and language decoder [2] with RL-based training strategy [38]. Extensive experiments on MS-COCO [26] validates the superiority of using SGAE in image captioning. Particularly, in terms of the popular CIDEr-D metric [44], we achieve an absolute 7.2 points improvement over a strong baseline: an upgraded version of Up-Down [2]. Then, we advance to a new state-of-the-art *single-model*

achieving 127.8 on the Karpathy split and a competitive 125.5 on the official test server even compared to many ensemble models.

In summary, we would like to make the following technical contributions:

- A novel Scene Graph Auto-Encoder (SGAE) for learning the feature representation of the language inductive bias.
- A multi-modal graph convolutional network for modulating scene graphs into visual representations.
- A SGAE-based encoder-decoder image captioner with a shared dictionary guiding the language decoding.

2. Related Work

Image Captioning. There is a long history for researchers to develop automatic image captioning methods. Compared with early works which are rules/templates based [21, 33, 23], the modern captioning models have achieved striking advances by three techniques inspired from the NLP field, *i.e.*, encoder-decoder based pipeline [46], attention technique [50], and RL-based training objective [38]. Afterwards, researchers tried to discover more semantic information from images and incorporated them into captioning models for better descriptive abilities. For example, some methods exploit object [29], attribute [54], and relationship [53] knowledge into their captioning models. Compared with these approaches, we use the scene graph as the bridge to integrate object, attribute, and relationship knowledge together to discover more meaningful semantic contexts for better caption generations.

Scene Graphs. The scene graph contains the structured semantic information of an image, which includes the knowledge of present objects, their attributes, and pairwise relationships. Thus, the scene graph can provide a beneficial prior for other vision tasks like VQA [40, 43], image generation [13, 59], and visual grounding [28]. By observing the potential of exploiting scene graphs in vision tasks, a variety of approaches are proposed to improve the scene graph generation from images [57, 56, 52, 51, 49, 42, 6]. On the another hand, some researchers also tried to extract scene graphs only from textual data [1, 47]. In this research, we use [56] to parse scene graphs from images and [1] to parse scene graphs from captions.

Memory Networks. Recently, many researchers try to augment a working memory into network for preserving a dynamic knowledge base for facilitating subsequent inference [41, 48, 45]. Among these methods, differentiable attention mechanisms are usually applied to extract useful knowledge from memory for the tasks at hand. Inspired by these methods, we also implement a memory architecture to preserve humans’ inductive bias, guiding our image captioning model to generate more descriptive captions.

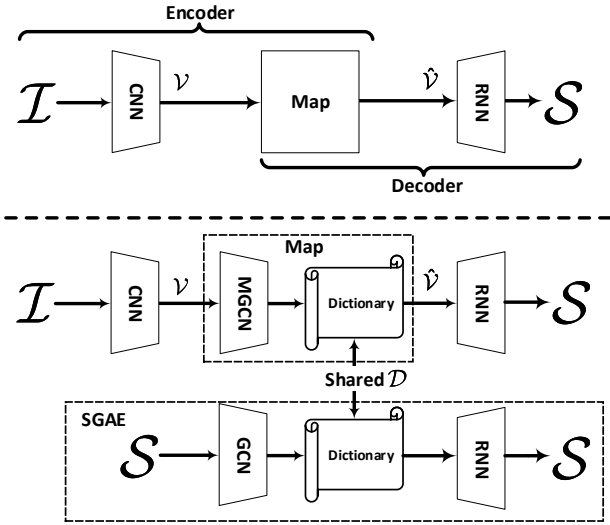


Figure 2. Top: the conventional encoder-decoder. Bottom: our proposed encoder-decoder, where the novel SGAE embeds the language inductive bias in the shared dictionary.

3. Encoder-Decoder Revisited

As illustrated in Figure 2, given an image \mathcal{I} , the target of image captioning is to generate a natural language sentence $\mathcal{S} = \{w_1, w_2, \dots, w_T\}$ describing the image. A state-of-the-art encoder-decoder image captioner can be formulated as:

$$\begin{aligned}
 \text{Encoder: } & \mathcal{V} \leftarrow \mathcal{I}, \\
 \text{Map: } & \hat{\mathcal{V}} \leftarrow \mathcal{V}, \\
 \text{Decoder: } & \mathcal{S} \leftarrow \hat{\mathcal{V}}.
 \end{aligned} \tag{1}$$

Usually, an encoder is a Convolutional Neural Network (CNN) [11, 37] that extracts the image feature \mathcal{V} ; the map is the widely used attention mechanism [50, 2] that re-encodes the visual features into more informative $\hat{\mathcal{V}}$ that is dynamic to language generation; a decoder is an RNN-based language decoder for the sequence prediction of \mathcal{S} . Given a ground truth caption \mathcal{S}^* for \mathcal{I} , we can train this encoder-decoder model by minimizing the cross-entropy loss:

$$L_{XE} = -\log P(\mathcal{S}^*), \tag{2}$$

or by maximizing a reinforcement learning (RL) based reward [38] as:

$$R_{RL} = \mathbb{E}_{\mathcal{S}^s \sim P(\mathcal{S})} [r(\mathcal{S}^s; \mathcal{S}^*)], \tag{3}$$

where r is a sentence-level metric for the sampled sentence \mathcal{S}^s and the ground-truth \mathcal{S}^* , e.g., the CIDEr-D [44] metric.

This encoder-decoder framework is the core pillar underpinning almost all state-of-the-art image captioners since [46]. However, it is widely shown brittle to dataset

bias [14, 29]. We propose to exploit the language inductive bias, which is beneficial, to confront the dataset bias, which is pernicious, for more human-like image captioning. As shown in Figure 2, the proposed framework is formulated as:

$$\begin{aligned}
 \text{Encoder: } & \mathcal{V} \leftarrow \mathcal{I}, \\
 \text{Map: } & \hat{\mathcal{V}} \leftarrow R(\mathcal{V}, \mathcal{G}; \mathcal{D}), \mathcal{G} \leftarrow \mathcal{V}, \\
 \text{Decoder: } & \mathcal{S} \leftarrow \hat{\mathcal{V}}.
 \end{aligned} \tag{4}$$

As can be clearly seen that we focus on modifying the Map module by introducing the scene graph \mathcal{G} into a re-encoder R parameterized by a shared dictionary \mathcal{D} . As we will detail in the rest of the paper, we first propose a Scene Graph Auto-Encoder (SGAE) to learn the dictionary \mathcal{D} which embeds the language inductive bias from sentence to sentence self-reconstruction (cf. Section 4) with the help of scene graphs. Then, we equip the encoder-decoder with the proposed SGAE to be our overall image captioner (cf. Section 5). Specifically, we use a novel Multi-modal Graph Convolutional Network (MGCN) (cf. Section 5.1) to re-encode the image features by using \mathcal{D} , narrowing the gap between vision and language.

4. Auto-Encoding Scene Graphs

In this section, we will introduce how to learn \mathcal{D} through self-reconstructing sentence \mathcal{S} . As shown in Figure 2, the process of reconstructing \mathcal{S} is also an encoder-decoder pipeline. Thus, by slightly abusing the notations, we can formulate SGAE as:

$$\begin{aligned}
 \text{Encoder: } & \mathcal{X} \leftarrow \mathcal{G} \leftarrow \mathcal{S}, \\
 \text{Map: } & \hat{\mathcal{X}} \leftarrow R(\mathcal{X}; \mathcal{D}), \\
 \text{Decoder: } & \mathcal{S} \leftarrow \hat{\mathcal{X}}.
 \end{aligned} \tag{5}$$

Next, we will detail every component mentioned in Eq. (5).

4.1. Scene Graphs

We introduce how to implement the step $\mathcal{G} \leftarrow \mathcal{S}$, i.e., from sentence to scene graph. Formally, a scene graph is a tuple $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where \mathcal{N} and \mathcal{E} are the sets of nodes and edges, respectively. There are three kinds of nodes in \mathcal{N} : object node o , attribute node a , and relationship node r . We denote o_i as the i -th object, r_{ij} as the relationship between object o_i and o_j , and $a_{i,l}$ as the l -th attribute of object o_i . For each node in \mathcal{N} , it is represented by a d -dimensional vector, i.e., e_o , e_a , and e_r . In our implementation, d is set to 1,000. In particular, the node features are trainable label embeddings. The edges in \mathcal{E} are formulated as follows:

- if an object o_i owns an attribute $a_{i,l}$, assigning a directed edge from $a_{i,l}$ to o_i ;
- if there is one relationship triplet $\langle o_i - r_{ij} - o_j \rangle$ appearing, assigning two directed edges from o_i to r_{ij} and from r_{ij} to o_j , respectively.

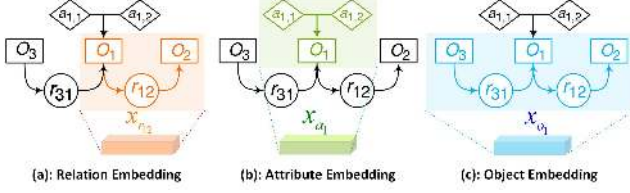


Figure 3. Graph Convolutional Network. In particular, it is spatial convolution, where the colored neighborhood is “convolved” for the resultant embedding.

Figure 3 shows one example of \mathcal{G} , which contains 7 nodes in \mathcal{N} and 6 directed edges in \mathcal{E} .

We use the scene graph parser provided by [1] for scene graphs \mathcal{G} from sentences, where a syntactic dependency tree is built by [18] and then a rule-based method [39] is applied for transforming the tree to a scene graph.

4.2. Graph Convolution Network

We present the implementation for the step $\mathcal{X} \leftarrow \mathcal{G}$ in Eq. (5), i.e., how to transform the original node embeddings e_o , e_a , and e_r into a new set of context-aware embeddings \mathcal{X} . Formally, \mathcal{X} contains three kinds of d -dimensional embeddings: relationship embedding $x_{r_{ij}}$ for relationship node r_{ij} , object embedding x_{o_i} for object node o_i , and attribute embedding x_{a_i} for object node o_i . In our implementation, d is set to 1,000. We use four *spatial graph convolutions*: g_r , g_a , g_s , and g_o for generating the above mentioned three kinds of embeddings. In our implementation, all these four functions have the same structure with independent parameters: a vector concatenation input to a fully-connected layer, followed by an ReLU.

Relationship Embedding $x_{r_{ij}}$: Given one relationship triplet $\langle o_i - r_{ij} - o_j \rangle$ in \mathcal{G} , we have:

$$x_{r_{ij}} = g_r(e_{o_i}, e_{r_{ij}}, e_{o_j}), \quad (6)$$

where the context of a relationship triplet is incorporated together. Figure 3 (a) shows such an example.

Attribute Embedding x_{a_i} : Given one object node o_i with all its attributes $a_{i,1:N_{a_i}}$ in \mathcal{G} , where N_{a_i} is the number of attributes that the object o_i has, then x_{a_i} for o_i is:

$$x_{a_i} = \frac{1}{N_{a_i}} \sum_{l=1}^{N_{a_i}} g_a(e_{o_i}, e_{a_{i,l}}), \quad (7)$$

where the context of this object and all its attributes are incorporated. Figure 3 (b) shows such an example.

Object Embedding x_{o_i} : In \mathcal{G} , o_i can act as “subject” or “object” in relationships, which means o_i will play different roles due to different edge directions. Then, different functions should be used to incorporate such knowledge. For avoiding ambiguous meaning of the same “predicate” in different context, knowledge of the whole relationship triplets

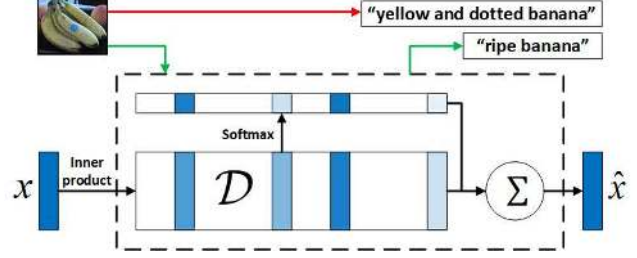


Figure 4. The visualization of the re-encoder function R . The black dashed block shows the operation of re-encoding. The top part demonstrates how “imagination” is achieved by re-encoding: green line shows the generated phrase by re-encoding, while the red line shows the one without re-encoding.

where o_i appears should be incorporated into x_{o_i} . One simple example for ambiguity is that, in $\langle \text{hand-with-cup} \rangle$, the predicate “with” may mean “hold”, while in $\langle \text{head-with-hat} \rangle$, “with” may mean “wear”. Therefore, x_{o_i} can be calculated as:

$$x_{o_i} = \frac{1}{Nr_i} \left[\sum_{o_j \in \text{subj}(o_i)} g_s(e_{o_i}, e_{o_j}, e_{r_{ij}}) + \sum_{o_k \in \text{obj}(o_i)} g_o(e_{o_k}, e_{o_i}, e_{r_{ki}}) \right]. \quad (8)$$

For each node $o_j \in \text{subj}(o_i)$, it acts as “object” while o_i acts as “subject”, e.g., $\text{subj}(o_1) = \{o_2\}$ in Figure 3 (c). $Nr_i = |\text{subj}(i)| + |\text{obj}(i)|$ is the number of relationship triplets where o_i is present. Figure 3 (c) shows this example.

4.3. Dictionary

Now we introduce how to learn the dictionary \mathcal{D} and then use it to re-encode $\hat{\mathcal{X}} \leftarrow R(\mathcal{X}; \mathcal{D})$ in Eq. (5). Our key idea is inspired by using the working memory to preserve a dynamic knowledge base for run-time inference, which is widely used in textual QA [41], VQA [48], and one-shot classification [45]. Our \mathcal{D} aims to embed language inductive bias in language composition. Therefore, we propose to place the dictionary learning into the sentence self-reconstruction framework. Formally, we denote \mathcal{D} as a $d \times K$ matrix $\mathbf{D} = \{\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_K\}$. The K is set as 10,000 in implementation. Given an embedding vector $x \in \mathcal{X}$, the re-encoder function $R_{\mathcal{D}}$ can be formulated as:

$$\hat{x} = R(x; \mathcal{D}) = \mathbf{D}\alpha = \sum_{k=1}^K \alpha_k \mathbf{d}_k, \quad (9)$$

where $\alpha = \text{softmax}(\mathbf{D}^T x)$ can be viewed as the “key” operation in memory network [41]. As shown in Figure 4, this re-encoding offers some interesting “imagination” in human common sense reasoning. For example, from “yellow and dotted banana”, after re-encoding, the feature will be more likely to generate “ripe banana”.

We deploy the attention structure in [2] for reconstructing \mathcal{S} . Given a reconstructed \mathcal{S} , we can use the training objective in Eq. (2) or (3) to train SGAE parameterized by \mathcal{D} in an end-to-end fashion. Note that training SGAE is unsupervised, that is, SGAE offers a potential never-ending learning from large-scale unsupervised inductive bias learning for \mathcal{D} . Some preliminary studies are reported in Section 6.2.2.

5. Overall Model: SGAE-based Encoder-Decoder

In this section, we will introduce the overall model: SGAE-based Encoder-Decoder as sketched in Figure 2 and Eq. (4).

5.1. Multi-modal Graph Convolution Network

The original image features extracted by CNN are not ready for use for the dictionary re-encoding as in Eq. (9), due to the large gap between vision and language. To this end, we propose a Multi-modal Graph Convolution Network (MGCN) to first map the visual features \mathcal{V} into a set of scene graph-modulated features \mathcal{V}' .

Here, the scene graph \mathcal{G} is extracted by an image scene graph parser that contains an object proposal detector, an attribute classifier, and a relationship classifier. In our implementation, we use Faster-RCNN as the object detector [37], MOTIFS relationship detector [56] as the relationship classifier, and we use our own attribute classifier: a small fc-ReLU-fc-Softmax network head. The key representation difference between the sentence-parsed \mathcal{G} and the image-parsed \mathcal{G} is that the node o_i is not only the label embedding. In particular, we use the RoI features pre-trained from Faster RCNN and then fuse the detected label embedding e_{o_i} with the visual feature v_{o_i} , into a new node feature u_{o_i} :

$$u_{o_i} = \text{ReLU}(\mathbf{W}_1 e_{o_i} + \mathbf{W}_2 v_{o_i}) - (\mathbf{W}_1 e_{o_i} - \mathbf{W}_2 v_{o_i})^2. \quad (10)$$

where \mathbf{W}_1 and \mathbf{W}_2 are the fusion parameters following [58]. Compared to the popular bi-linear fusion [58], Eq (10) is empirically shown a faster convergence of training the label embeddings in our experiments. The rest node embeddings: $u_{r_{ij}}$ and u_{a_i} are obtained in a similar way. The differences between two scene graphs generated from \mathcal{I} and \mathcal{S} are visualized in Figure 1, where the image \mathcal{G} is usually more simpler and nosier than the sentence \mathcal{G} .

Similar to the GCN used in Section 4.2, MGCN also has an ensemble of four functions f_r , f_a , f_s and f_o , each of which is a two-layer structure: fc-ReLU with independent parameters. And the computation of relationship, attribute and object embeddings are similar to Eq. (6), Eq. (7), and Eq. (8), respectively. After computing \mathcal{V}' by using MGCN, we can adopt Eq. (9) to re-encode \mathcal{V}' as $\hat{\mathcal{V}}$ and feed $\hat{\mathcal{V}}$ to the decoder for generating language \mathcal{S} . In particular, we deploy the attention structure in [2] for the generation.

5.2. Training and Inference

Following the common practice in deep-learning feature transfer [7, 55], we use the SGAE pre-trained \mathcal{D} as the initialization for the \mathcal{D} in our overall encoder-decoder for image captioning. In particular, we intentionally use a very small learning rate (*e.g.*, 10^{-5}) for fine-tuning \mathcal{D} to impose the sharing purpose. The overall training loss is hybrid: we use the cross-entropy loss in Eq. (2) for 20 epochs and then use the RL-based reward in Eq. (3) for another 40 epochs.

For inference in language generation, we adopt the beam search strategy [38] with a beam size of 5.

6. Experiments

6.1. Datasets, Settings, and Metrics

MS-COCO [26]. There are two standard splits of MS-COCO: the official online test split and the 3rd-party Karpathy split [15] for offline test. The first split has 82, 783/40, 504/40, 775 train/val/test images, each of which has 5 human labeled captions. The second split has 113, 287/5, 000/5, 000 train/val/test images, each of which has 5 captions.

Visual Genome [19] (VG). This dataset has abundant scene graph annotations, *e.g.*, objects' categories, objects' attributes, and pairwise relationships, which can be exploited to train the object proposal detector, attribute classifier, and relationship classifier [56] as our image scene graph parser.

Settings. For captions, we used the following steps to pre-process the captions: we first tokenized the texts on white space; then we changed all the words to lowercase; we also deleted the words which appear less than 5 times; at last, we trimmed each caption to a maximum of 16 words. This results in a vocabulary of 10, 369 words. This pre-processing was also applied in VG. It is noteworthy that except for ablation studies, these additional text descriptions from VG were not used for training the captioner. Since the object, attribute, and relationship annotations are very noisy in VG dataset, we filter them by keeping the objects, attributes, and relationships which appear more than 2, 000 times in the training set. After filtering, the remained 305 objects, 103 attributes, and 64 relationships are used to train our object detector, attribute classifier and relationship classifier.

We chose the language decoder proposed in [2]. The number of hidden units of both LSTMs used in this decoder is set to 1000. For training SGAE in Eq. (5), the decoder is firstly set as $\mathcal{S} \leftarrow \mathcal{X}$ and \mathcal{D} is not trained to learn a rudiment encoder and decoder. We used the corss-entropy loss in Eq. (2) to train them for 20 epochs. Then the decoder was set as $\mathcal{S} \leftarrow \hat{\mathcal{X}}$ to train \mathcal{D} by cross-entropy loss for another 20 epochs. The learning rate was initialized to $5e^{-4}$ for all parameters and we decayed them by 0.8 for every 5 epochs. For training our SGAE-based encoder-decoder, we followed Eq. (4) to generate \mathcal{S} with shared \mathcal{D} pre-trained

Table 1. The performances of various methods on MS-COCO Karpathy split. The metrics: B@N, M, R, C and S denote BLEU@N, METEOR, ROUGE-L, CIDEr-D and SPICE. Note that the *fuse* subscript indicates fused models while the rest methods are all single models. The best results for each metric on fused models and single models are marked in boldface separately.

Models	B@1	B@4	M	R	C	S
SCST [38]	–	34.2	26.7	55.7	114.0	–
LSTM-A [54]	78.6	35.5	27.3	56.8	118.3	20.8
StackCap [9]	78.6	36.1	27.4	–	120.4	–
Up-Down [2]	79.8	36.3	27.7	56.9	120.1	21.4
CAVP [27]	–	38.6	28.3	58.5	126.3	21.6
GCN-LSTM [†] [53]	80.0	37.1	28.0	57.3	122.8	21.1
GCN-LSTM [53]	80.5	38.2	28.5	58.3	127.6	22.0
Base	79.9	36.8	27.7	57.0	120.6	20.9
Base+MGCN	80.2	37.2	27.9	57.5	123.4	21.2
Base+ <i>D</i> w/o GCN	80.2	37.3	27.8	58.0	124.2	21.4
Base+ <i>D</i>	80.4	37.7	28.1	58.2	125.7	21.4
SGAE	80.8	38.4	28.4	58.6	127.8	22.1
SGAE _{<i>fuse</i>}	81.0	39.0	28.4	58.9	129.1	22.2
GCN-LSTM _{<i>fuse</i>} [53]	80.9	38.3	28.6	58.5	128.7	22.1

from SGAE. The decoder was set as $\mathcal{S} \leftarrow \{\hat{\mathcal{V}}, \mathcal{V}'\}$, where \mathcal{V}' and $\hat{\mathcal{V}}$ can provide visual clues and high-level semantic contexts respectively. In this process, cross-entropy loss was first used to train the network for 20 epochs and then the RL-based reward was used to train for another 80 epochs. The learning rate for \mathcal{D} was initialized to $5e^{-5}$ and for other parameters it was $5e^{-4}$, and all these learning rates were decayed by 0.8 for every 5 epochs. Adam optimizer [16] was used for batch size 100.

Metrics. We used five standard automatic evaluation metrics: CIDEr-D [44], BLEU [34], METEOR [4], ROUGE [25] and SPICE [1].

6.2. Ablative Studies

We conducted extensive ablations for architecture (Section 6.2.1), language corpus (Section 6.2.2), and sentence reconstruction quality (Section 6.2.3). For simplicity, we use **SGAE** to denote our SGAE-based encoder-decoder captioning model.

6.2.1 Architecture

Comparing Methods. For quantifying the importance of the proposed GCN, MGCN, and dictionary \mathcal{D} , we ablated our SGAE with the following baselines: **Base:** We followed the pipeline given in Eq (1) without using GCN, MGCN, and \mathcal{D} . This baseline is the benchmark for other ablative baselines. **Base+MGCN:** We added MGCN to compute the multi-modal embedding set $\hat{\mathcal{V}}$. This baseline is designed for validating the importance of MGCN. **Base+*D* w/o GCN:** We learned \mathcal{D} by using Eq. (5), while GCN is not used and only word embeddings of \mathcal{S} were input to the decoder. Also, MGCN in Eq. (4) is not used. This baseline is designed for validating the importance of GCN. **Base+*D*:** Compared to

Table 2. The performances of using different language corpora

Models	B@1	B@4	M	R	C	S
Base	79.9	36.8	27.7	57.0	120.6	20.9
VG	80.2	37.8	28.0	58.2	123.2	21.3
COCO	80.8	38.4	28.4	58.6	127.8	22.1
VG+COCO	81.1	38.9	28.6	58.8	128.6	22.2

Table 3. The performances of using different scene graphs

Models	B@1	B@4	M	R	C	S
$\tilde{\mathcal{X}}$	90.3	53.8	34.3	66.5	153.2	30.6
\mathcal{X}	93.9	65.2	38.5	71.8	177.0	34.3
SGAE	80.8	38.4	28.4	58.6	127.8	22.1

Table 4. The performances of various methods on MS-COCO Karpathy split trained by cross-entropy loss only.

Models	B@1	B@4	M	R	C	S
SCST[36]	–	30.0	25.9	53.4	99.4	–
LSTM-A[50]	73.4	32.6	25.4	54.0	100.2	18.6
StackCap[8]	76.2	35.2	26.5	–	109.1	–
Up-Down[2]	77.2	36.2	27.0	56.4	113.5	20.3
GCN-LSTM[49]	77.3	36.8	27.9	57.0	116.3	20.9
Base	76.8	36.1	27.1	56.3	113.1	20.3
Base+MGCN	77.1	36.2	27.2	56.5	114.2	20.5
SGAE	77.6	36.9	27.7	57.2	116.7	20.9

Base, we learned \mathcal{D} by using GCN. And MGCN in Eq. (4) was not used. This baseline is designed for validating the importance of the shared \mathcal{D} .

Results. The middle section of Table 1 shows the performances of the ablative baselines on MS-COCO Karpathy split. Compared with Base, our SGAE can boost the CIDEr-D by absolute 7.2. By comparing Base+MGCN, Base+*D* w/o GCN, and Base+*D* with Base, we can find that all the performances are improved, which demonstrate that the proposed MGCN, GCN, and \mathcal{D} are all indispensable for advancing the performances. We can also observe that the performances of Base+*D* or Base+*D* w/o GCN are better than Base+MGCN, which suggests that the language inductive bias plays an important role in generating better captions.

Qualitative Examples. Figure 5 shows 6 examples of the generated captions using different baselines. We can see that compared with captions generated by Base, Base+MGCN’s descriptions usually contain more descriptions about objects’ attributes and pairwise relationships. For captions generated by SGAE, they are more complex and descriptive. For example, in Figure 5 (a), the word “busy” will be used to describe the heavy traffic; in (b) the scene “forest” can be deduced from “trees”; and in (d), the weather “rain” will be inferred from “umbrella”.

6.2.2 Language Corpus

Comparing Methods. To test the potential of using large-scale corpus for learning a better \mathcal{D} , we used the texts provided by VG instead of MS-COCO to learn \mathcal{D} , and then share the learned \mathcal{D} in the encoder-decoder pipeline. The results are demonstrated in Table 2, where **VG** and

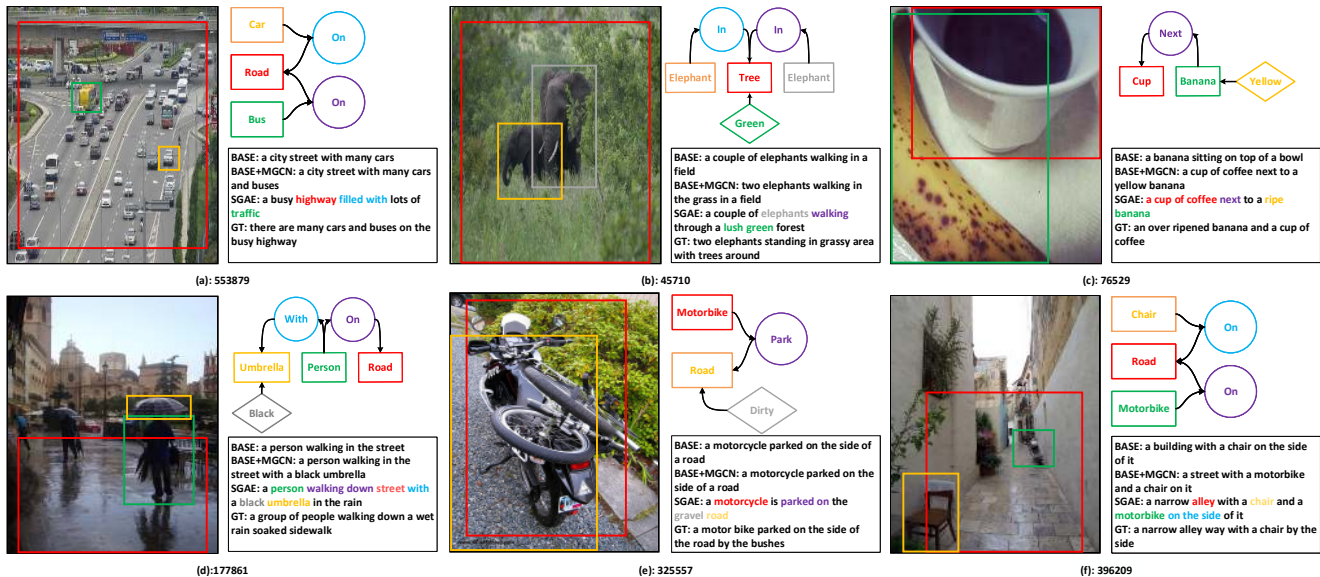


Figure 5. Qualitative examples of different baselines. For each figure, the image scene graph is pruned to avoid clutter. The id refers to the image id in MS-COCO. Word colors correspond to nodes in the detected scene graphs.



Figure 6. Captions generated by using different language corpora.

Table 5. The performances of various methods on the online MS-COCO test server.

Model	B@4		M		R-L		C-D	
	c5	c40	c5	c40	c5	c40	c5	c40
SCST [38]	35.2	64.5	27.0	35.5	56.3	70.7	114.7	116.0
LSTM-A [54]	35.6	65.2	27.0	35.4	56.4	70.5	116.0	118.0
StackCap [9]	34.9	64.6	27.0	35.6	56.2	70.6	114.8	118.3
Up-Down [2]	36.9	68.5	27.6	36.7	57.1	72.4	117.9	120.5
CAVP [27]	37.9	69.0	28.1	37.0	58.2	73.1	121.6	123.8
SGAE _{single}	37.8	68.7	28.1	37.0	58.2	73.1	122.7	125.5
SGAE _{fuse}	38.5	69.7	28.2	37.2	58.6	73.6	123.8	126.5

VG+COCO mean the results obtained by using texts from VG and VG+COCO, respectively.

Results. We can observe that by using the web description texts, the performances are boosted, which validates the potential of our proposed model in exploiting additional Web texts. We can also see that by using texts provided by MS-COCO itself (SGAE), the generated captions have better scores than using Web texts. This is intuitively reasonable since \mathcal{D} can preserve more useful clues when a matched language corpus is given. All of these observations validate the

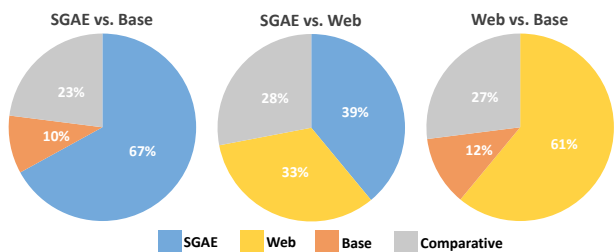


Figure 7. The pie charts each comparing the two methods in human evaluation. Each color indicates the percentage of users who consider that the corresponding method generates more descriptive captions. In particular, the gray color indicates that the two methods are comparative.

effectiveness of \mathcal{D} in two aspects: \mathcal{D} can memorize common inductive bias from the additional unmatched Web texts or specific inductive bias from a matched language corpus. In addition, when language corpora of both VG and COCO are exploited, a consistent improvement can be achieved compared with the baselines where only one corpus is used.

Qualitative Examples. Figure 6 shows 6 examples of generated captions by using different language corpora. Generally, compared with captions generated by Base, the captions of Web and SGAE are more descriptive. Specifically, the captions generated by using the matched language corpus can usually describe a scene by some specific expressions in the dataset, while more general expressions will appear in captions generated by using Web texts. For example, in Figure 6 (b), SGAE uses “lush green field” as GT captions while Web uses “grass”; or in (e), SGAE prefers “dirt” while Web prefers “sand”.

Human Evaluation. For better evaluating the qualities of the generated captions by using different language corpora, we conducted human evaluation with 30 workers. We showed them two captions generated by different methods and asked them which one is more descriptive. For each pairwise comparison, 100 images are randomly extracted from the Karpathy split for them to compare. The results of the comparisons are shown in Figure 7. From these pie charts, we can observe that when a \mathcal{D} is used, the generated captions are evaluated to be more descriptive.

6.2.3 Sentence Reconstruction

Comparing Methods. We investigated how well the sentences are reconstructed in training SGAE in Eq. (5), with or without using the re-encoding by \mathcal{D} , that is, we denote $\hat{\mathcal{X}}$ as the pipeline using \mathcal{D} and \mathcal{X} as the pipeline directly reconstructing sentences from their scene graph node features. Such results are given in Table 3.

Analysis. As we can see, the performances of using direct scene graph features $\hat{\mathcal{X}}$ are much better than those (\mathcal{X}) imposed with \mathcal{D} for re-encoding. This is reasonable since \mathcal{D} will regularize the reconstruction and thus encourages the learning of language inductive bias. Interestingly, the gap between $\hat{\mathcal{X}}$ and SGAE suggest that we should develop a more powerful image scene graph parser for improving the quality of \mathcal{G} in Eq. (4), and a stronger re-encoder should be designed for extracting more preserved inductive bias when only low-quality visual scene graphs are available.

6.3. Comparisons with State-of-The-Arts

Comparing Methods. Though there are various captioning models developed in recent years, for fair comparison, we only compared SGAE with some encoder-decoder methods trained by the RL-based reward (Eq. (3)), due to their superior performances. Specifically, we compared our methods with SCST [38], StackCap [9], Up-Down [2], LSTM-A [54], GCN-LSTM [53], and CAVP [27]. Among these methods, SCST and Up-Down are two baselines where the more advanced self-critic reward and visual features are used. Compared with SCST, StackCap proposes a more complex RL-based reward for learning captions with more

details. All of LSTM-A, GCN-LSTM, and CAVP try to exploit information of visual scene graphs, *e.g.*, LSTM-A and GCN-LSTM exploit attributes and relationships information respectively, while CAVP tries to learn pairwise relationships in the decoder. Noteworthy, in GCN-LSTM, they set the batch size as 1,024 and the training epoch as 250, which is quite large compared with some other methods like Up-Down or CAVP, and is beyond our computation resources. For fair comparison, we also re-implemented a version of their work (since they do not publish the code), and set the batch size and training epoch both as 100, such result is denoted as GCN-LSTM[†] in Table 1. In addition, the best result reported by GCN-LSTM is obtained by fusing two probabilities computed from two different kinds of relationships, which is denoted as GCN-LSTM_{fuse}, and our counterpart is denoted as SGAE_{fuse}.

Analysis. From Table 1 and 4, we can see that our single model can outperform the other image captioners. In particular, when both cross-entropy loss and RL reward are used, our single model can achieve a new state-of-the-art score among all the compared methods in terms of CIDEr-D, 127.8. And compared with GCN-LSTM_{fuse}, our fusion model SGAE_{fuse} also achieves better performances. By exploiting the inductive bias in \mathcal{D} , even when our decoder or RL-reward is not as sophisticated as CVAP or StackCap, our method still has better performances. Moreover, our small batch size and fewer training epochs still lead to higher performances than GCN-LSTM, whose batch size and training epochs are much larger. Table 5 reports the performances of different methods test on the official server. Compared with the published captioning methods (by the date of 16/11/2018), our single model has competitive performances and can achieve the highest CIDEr-D score.

7. Conclusions

We proposed to incorporate the language inductive bias — a prior for more human-like language generation — into the prevailing encoder-decoder framework for image captioning. In particular, we presented a novel unsupervised learning method: Scene Graph Auto-Encoder (SGAE), for embedding the inductive bias into a dictionary, which can be shared as a re-encoder for language generation and significantly improves the performance of the encoder-decoder. We validated the SGAE-based framework by extensive ablations and comparisons with state-of-the-art performances on MS-COCO. As we believe that SGAE is a general solution for capturing the language inductive bias, we are going to apply it in other vision-language tasks.

Acknowledgements. This work is partially supported by NTU Data Science and Artificial Intelligence Research Center (DSAIR) and NTU-Alibaba Lab.

References

- [1] P. Anderson, B. Fernando, M. Johnson, and S. Gould. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, pages 382–398. Springer, 2016. 2, 4, 6
- [2] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, volume 3, page 6, 2018. 1, 2, 3, 5, 6, 7, 8
- [3] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *ICLR*, 2015. 1
- [4] S. Banerjee and A. Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. 6
- [5] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. 2
- [6] L. Chen, H. Zhang, J. Xiao, X. He, S. Pu, and S.-F. Chang. Counterfactual critic multi-agent training for scene graph generation. *arXiv preprint arXiv:1812.02347*, 2018. 2
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 5
- [8] 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. 2
- [9] J. Gu, J. Cai, G. Wang, and T. Chen. Stack-captioning: Coarse-to-fine learning for image captioning. *AAAI*, 2017. 1, 6, 7, 8
- [10] J. Gu, G. Wang, J. Cai, and T. Chen. An empirical study of language cnn for image captioning. In *ICCV*, 2017. 1
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [12] R. Hu, P. Dollr, K. He, T. Darrell, and R. Girshick. Learning to segment every thing. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [13] J. Johnson, A. Gupta, and L. Fei-Fei. Image generation from scene graphs. *arXiv preprint*, 2018. 2
- [14] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 1988–1997. IEEE, 2017. 3
- [15] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015. 5
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [17] A. Kirillov, K. He, R. Girshick, C. Rother, and P. Dollár. Panoptic segmentation. *arXiv preprint arXiv:1801.00868*, 2018. 1
- [18] D. Klein and C. D. Manning. Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pages 423–430. Association for Computational Linguistics, 2003. 4
- [19] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al. Visual genome: Connecting language and vision using crowd-sourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017. 5
- [20] G. Kulkarni, V. Premraj, V. Ordonez, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg. Babytalk: Understanding and generating simple image descriptions. In *CVPR*, 2011. 2
- [21] P. Kuznetsova, V. Ordonez, A. C. Berg, T. L. Berg, and Y. Choi. Collective generation of natural image descriptions. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 359–368. Association for Computational Linguistics, 2012. 2
- [22] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017. 2
- [23] S. Li, G. Kulkarni, T. L. Berg, A. C. Berg, and Y. Choi. Composing simple image descriptions using web-scale n-grams. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 220–228. Association for Computational Linguistics, 2011. 2
- [24] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015. 2
- [25] C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*, 2004. 6
- [26] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 5
- [27] D. Liu, Z.-J. Zha, H. Zhang, Y. Zhang, and F. Wu. Context-aware visual policy network for sequence-level image captioning. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 1416–1424. ACM, 2018. 1, 6, 7, 8
- [28] D. Liu, H. Zhang, Z.-J. Zha, and F. Wu. Learning to assemble neural module tree networks for visual grounding. *arXiv preprint arXiv:1812.03299*, 2018. 2
- [29] J. Lu, J. Yang, D. Batra, and D. Parikh. Neural baby talk. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7219–7228, 2018. 1, 2, 3
- [30] R. Luo, B. Price, S. Cohen, and G. Shakhnarovich. Discriminability objective for training descriptive captions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6964–6974, 2018. 1
- [31] D. Marcheggiani and I. Titov. Encoding sentences with graph convolutional networks for semantic role labeling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1506–1515, 2017. 2

- [32] D. Marr. Vision: A computational investigation into the human representation and processing of visual information. mit press. Cambridge, Massachusetts, 1982. 1
- [33] M. Mitchell, X. Han, J. Dodge, A. Mensch, A. Goyal, A. Berg, K. Yamaguchi, T. Berg, K. Stratos, and H. Daumé III. Midge: Generating image descriptions from computer vision detections. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 747–756. Association for Computational Linguistics, 2012. 2
- [34] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics, 2002. 6
- [35] M. Ranzato, S. Chopra, M. Auli, and W. Zaremba. Sequence level training with recurrent neural networks. 2015. 1
- [36] J. Redmon and A. Farhadi. Yolo9000: Better, faster, stronger. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6517–6525. IEEE, 2017. 1
- [37] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 2, 3, 5
- [38] S. J. Rennie, E. Marcheret, Y. Mroueh, J. Ross, and V. Goel. Self-critical sequence training for image captioning. In *CVPR*, volume 1, page 3, 2017. 1, 2, 3, 5, 6, 7, 8
- [39] S. Schuster, R. Krishna, A. Chang, L. Fei-Fei, and C. D. Manning. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *Proceedings of the fourth workshop on vision and language*, pages 70–80, 2015. 4
- [40] J. Shi, H. Zhang, and J. Li. Explainable and explicit visual reasoning over scene graphs. *arXiv preprint arXiv:1812.01855*, 2018. 2
- [41] S. Sukhbaatar, J. Weston, R. Fergus, et al. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448, 2015. 2, 4
- [42] K. Tang, H. Zhang, B. Wu, W. Luo, and W. Liu. Learning to compose dynamic tree structures for visual contexts. *arXiv preprint arXiv:1812.01880*, 2018. 2
- [43] D. Teney, L. Liu, and A. van den Hengel. Graph-structured representations for visual question answering. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3233–3241. IEEE, 2017. 2
- [44] R. Vedantam, C. Lawrence Zitnick, and D. Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. 2, 3, 6
- [45] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pages 3630–3638, 2016. 2, 4
- [46] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015. 2, 3
- [47] Y.-S. Wang, C. Liu, X. Zeng, and A. Yuille. Scene graph parsing as dependency parsing. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 397–407. Association for Computational Linguistics, 2018. 2
- [48] C. Xiong, S. Merity, and R. Socher. Dynamic memory networks for visual and textual question answering. In *International conference on machine learning*, pages 2397–2406, 2016. 2, 4
- [49] D. Xu, Y. Zhu, C. B. Choy, and L. Fei-Fei. Scene graph generation by iterative message passing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, 2017. 2
- [50] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015. 1, 2, 3
- [51] J. Yang, J. Lu, S. Lee, D. Batra, and D. Parikh. Graph r-cnn for scene graph generation. In *European Conference on Computer Vision*, pages 690–706. Springer, 2018. 2
- [52] X. Yang, H. Zhang, and J. Cai. Shuffle-then-assemble: Learning object-agnostic visual relationship features. In *European Conference on Computer Vision*, pages 38–54. Springer, 2018. 2
- [53] T. Yao, Y. Pan, Y. Li, and T. Mei. Exploring visual relationship for image captioning. In *Computer Vision—ECCV 2018*, pages 711–727. Springer, 2018. 2, 6, 8
- [54] T. Yao, Y. Pan, Y. Li, Z. Qiu, and T. Mei. Boosting image captioning with attributes. In *IEEE International Conference on Computer Vision, ICCV*, pages 22–29, 2017. 2, 6, 7, 8
- [55] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, pages 3320–3328, 2014. 5
- [56] R. Zellers, M. Yatskar, S. Thomson, and Y. Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018. 2, 5
- [57] H. Zhang, Z. Kyaw, S.-F. Chang, and T.-S. Chua. Visual translation embedding network for visual relation detection. In *CVPR*, volume 1, page 5, 2017. 2
- [58] Y. Zhang, J. Hare, and A. Prügel-Bennett. Learning to count objects in natural images for visual question answering. In *ICLR*, 2018. 5
- [59] B. Zhao, L. Meng, W. Yin, and L. Sigal. Image generation from layout. *arXiv preprint arXiv:1811.11389*, 2018. 2