

# Transforming Visual Scene Graphs to Image Captions

Xu Yang<sup>1</sup> Jiawei Peng<sup>1</sup> Zihua Wang<sup>1</sup> Haiyang Xu<sup>2\*</sup> Qinghao Ye<sup>2</sup>  
Chenliang Li<sup>2</sup> Songfang Huang<sup>2</sup> Fei Huang<sup>2</sup> Zhangzikang Li<sup>1</sup> Yu Zhang<sup>1\*</sup>

<sup>1</sup> School of Computer Science & Engineering, Key Lab of Computer Network  
& Information Integration (Ministry of Education), Southeast Univ., Nanjing, China

<sup>2</sup>Alibaba Group

{101013120, pengjiawei, zihua, zhang\_yu}@seu.edu.cn, {shuofeng.xhy, yeqinghao.yqh,  
lcl193798, songfang.hsf}@alibaba-inc.com, {fairhuang, lizhangzikang}@gmail.com

## Abstract

We propose to TransForm Scene Graphs into more descriptive Captions (**TFSGC**). In TFSGC, we apply multi-head attention (MHA) to design the Graph Neural Network (GNN) for embedding scene graphs. After embedding, different graph embeddings contain diverse specific knowledge for generating the words with different part-of-speech, *e.g.*, object/attribute embedding is good for generating nouns/adjectives. Motivated by this, we design a Mixture-of-Expert (MOE)-based decoder, where each expert is built on MHA, for discriminating the graph embeddings to generate different kinds of words. Since both the encoder and decoder are built based on the MHA, as a result, we construct a **simple and homogeneous** encoder-decoder unlike the previous **heterogeneous** ones which usually apply Fully-Connected-based GNN and LSTM-based decoder. The homogeneous architecture enables us to unify the training configuration of the whole model instead of specifying different training strategies for diverse sub-networks as in the heterogeneous pipeline, which releases the training difficulty. Extensive experiments on the MS-COCO captioning benchmark validate the effectiveness of our TFSGC. The code is in: [https://anonymous.4open.science/r/ACL23\\_TFSGC](https://anonymous.4open.science/r/ACL23_TFSGC).

## 1 Introduction

Image captioning, which aims to generate one sentence for describing multi-aspects of an image, has made huge progress since the proposal of the encoder-decoder framework (Vinyals et al., 2015; Xu et al., 2015). Such a framework contains one visual encoder to extract a series of visual features from the image and one language decoder to generate captions from the extracted visual features. Since the visual encoder is usually well pre-trained by image classification and object detection, the

extracted features contain abundant knowledge of the object categories, which enables the captioning model to generate object-abundant captions.

However, object category is not the only visual pattern that matters for high-quality captions (Anderson et al., 2018; Jiang et al., 2020), object attributes and relations also play significant roles in generating descriptive captions, *i.e.*, the caption containing multi-aspects of an image. Motivated by this, researchers propose to incorporate additional semantic knowledge, *e.g.*, object categories, attributes, and relations, into the captioning models by using the scene graph as the mediator (Yao et al., 2018; Yang et al., 2020). Scene graphs assign each object node with certain attribute nodes and some pairwise objects with certain relation nodes. These nodes are represented by the corresponding semantic tags, *e.g.*, as shown in Fig. 1, the object “dog” is assigned with the attribute “black” and the pairwise objects “dog” and “fish” have the relation “bite” in between. To exploit the scene graph, Graph Neural Network (GNN) (Battaglia et al., 2018) is deployed to embed the graphs and the output embeddings are input to the decoder for captioning.

The top part of Fig. 1 shows the pipeline of the previous popular GNN-based captioning model (Yao et al., 2018; Yang et al., 2020), which implements GNN as a few Fully-Connected (FC) and non-linear activation layers. To update the node embedding, this GNN maps the concatenated neighbour embeddings into the new one (Xu et al., 2019). Then the updated graph embeddings are input into the language decoder that contains a few LSTM layers and an attention module. The LSTM layers are used to generate the context vector based on the partially generated captions. This context vector works as the query in the attention module for determining which graph embeddings should be used to generate the next word. Compared with the models without GNN, this GNN-LSTM pipeline usually gets better performance.

\*Corresponding authors.

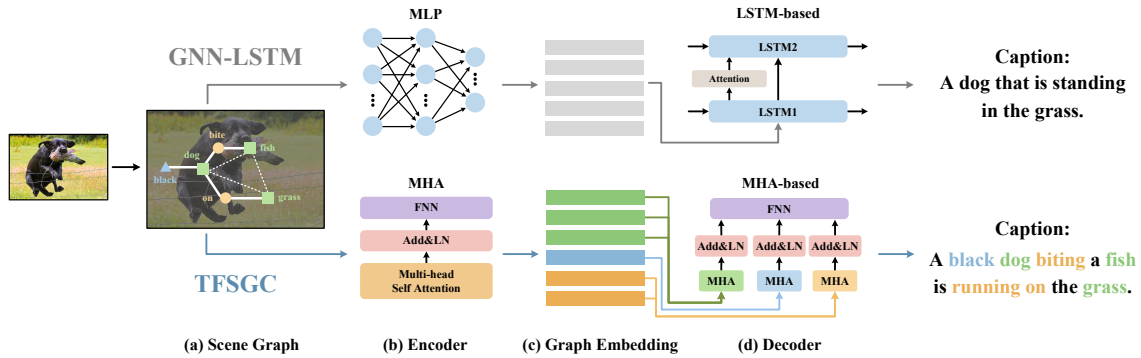


Figure 1: Comparison between traditional heterogeneous GNN-LSTM (top part) and our homogeneous TFSGC model (bottom part). In GNN-LSTM, they use MLP-based GNN and do not discriminate the graph embeddings (grey colour in (c) is used to strengthen such indiscrimination). In TFSGC, we use MHA to design the GNN and the decoder and discriminate diverse graph embeddings (different colours in (c) are used to strengthen such discrimination).

However, this GNN-LSTM framework implies two flaws which hinder the further improvement of applying scene graphs. First, FC-based GNN and LSTM do not share the same building blocks and thus the constructed model is a **heterogeneous** structure, which requires well-chosen training strategies, *e.g.*, choosing different learning rates or optimizers for different sub-networks, to achieve the best performance (Yang et al., 2020). Finding such training configurations is a labour-intensive process. Second, the graph embeddings are indiscriminately selected during captioning (the grey embeddings in Fig. 1 top (c) denote such indiscrimination), which causes less descriptive captions. While intuitively, different kinds of node embeddings should be used for generating the words with diverse part-of-speech (POS), *e.g.*, the object/attribute/relation embeddings should be more responsible for the nouns/adjectives/verbs, respectively (Yang et al., 2019b).

To alleviate the above-mentioned flaws, we propose a novel **homogeneous** captioning model to Transform Scene Graphs (TFSGC) into captions. Our TFSGC is built based on the Transformer (Vaswani et al., 2017) since it is more powerful than LSTM in image captioning (Herdade et al., 2019; Li et al., 2019; Cornia et al., 2020). TFSGC is homogeneous since we use multi-head attention (MHA) to design both the graph encoder to embed the scene graphs and the language decoder to generate the caption.

Our design principle is quite simple where we do not need to revise the self-attention operation but only need to reformulate the input data structure. Specifically, to design GNN by MHA, we

first linearize the scene graph into a token sequence and introduce a binary mask to index which two nodes are connected in the graph. Then we use the masked MHA operation to deal with this linearized token sequence for graph embedding. In this process, each token embedding is added by a learnable type embedding to index the token type (*e.g.*, object/attribute/relation) and we will show that such type embedding can help distinguish the edge type during the attention calculation.

After graph operation, we get a series of object/attribute/relation embeddings, which will be used in the decoder for captioning. To make the decoder discriminate different embeddings for generating different words, we learn from MOE networks (Jacobs et al., 1991; Xue et al., 2022; Du et al., 2022) to revise the original Transformer decoder with two strategies. First, as Fig. 1 bottom (d) shows, we use three encoder-decoder attention layers, which are built on MHA, as three experts to address object/attribute/relation embeddings, respectively. Second, we incorporate an attention-based soft routing network to discriminate which kinds of embeddings should be more responsible for generating the next word. Both the MOE-decoder and the type embedding in the encoder help distinguish node embeddings for better captions. We carry exhaustive ablation studies and comparisons to validate the effectiveness of TFSGC and it achieves 132.3/138.6/139.5 CIDEr scores when using BUTD/Patch/VinVL features.

## 2 Related Work

**Image Captioning.** For a long time, the attention-based CNN-LSTM pipeline (Vinyals et al., 2015;

Xu et al., 2015) is the most popular backbone for captioning and various techniques have been added into it for better performance, including building stronger visual encoders (Lu et al., 2018; Jiang et al., 2020), designing more advanced attention mechanisms (Anderson et al., 2018; Wang et al., 2020), incorporating semantic knowledge (You et al., 2016; Gan et al., 2017; Yao et al., 2018; Yang et al., 2020), and exploiting language structure (Lu et al., 2017; Yang et al., 2019b).

Recently, Transformer (Vaswani et al., 2017) has gradually substituted LSTM as the mainstream language decoder in image captioning (Herdade et al., 2019; Li et al., 2019) since it achieves better performances than the LSTM-based models. Based on this new backbone, researchers develop more advanced strategies for further improving the effectiveness, including designing more sophisticated attention mechanisms (Huang et al., 2019; Pan et al., 2020), introducing additional memory blocks (Cornia et al., 2020; Yang et al., 2021b), distilling knowledge from the large-scale pre-training models (Radford et al., 2021; Li et al., 2021; Xu et al., 2021), and exploiting Transformer-based visual encoders (Wang et al., 2022; Fang et al., 2022), modularized design for large-scale multi-modal pretraining (Li et al., 2022; Xu et al., 2023; Ye et al., 2023). Since the recently proposed SOTA models use Transformer as the backbone, we also built TFSGC based on Transformer for fair comparison.

**Graph Neural Network (GNN).** Scene Graph abstracts the major visual patterns in a visual scene as a graph. It is usually used as the mediator to narrow the gap between the vision and the language domains. To incorporate scene graphs into deep networks, GNN (Battaglia et al., 2018) is used to embed the discrete node labels into dense embeddings. However, most of the previous GNNs are MLP-based (Yang et al., 2020; Yao et al., 2018; Xu et al., 2019; Milewski et al., 2020; Zhong et al., 2020), which may limit the effectiveness of embedding scene graphs in a Transformer architecture. In our research, we design an MHA-based GNN to remedy this limitation. Moreover, noisy scene graphs may damage the performances (Nguyen et al., 2021), so we use better scene graph parsers to minimize the impact of noise on our model.

**Mixture of Experts (MOE).** The major idea of MOE is to construct a network with lots of experts where different experts deal with diverse sam-

ples (Jacobs et al., 1991; Shazeer et al., 2017). When a sample is input to the MOE network, a routing network will decide which experts should be more responsible for this input. Thus MOE naturally fits our case where we hope diverse experts can discriminate graph embeddings for generating the words with different POS. Different from the existent MOE-based Transformer (Lepikhin et al., 2020; Xue et al., 2022; Du et al., 2022) which applies various feed-forward networks as different experts, we set three encoder-decoder attention layers as different experts where the query is set to the same context vector while the key and value are set to object/attribute/relation embeddings.

### 3 Revisiting of Transformer

We first revisit the Transformer-based captioning model and then introduce how to revise it to get our TFSGC in the next section. Transformer-based model contains a visual encoder to calculate the contexts of the extracted visual features and the output embeddings will be input into a language decoder for captioning. For both the encoder and the decoder, the most elemental building block is the multi-head attention (MHA). Given the query, key, and value matrices:  $Q \in \mathbb{R}^{N_Q \times d}$ ,  $K \in \mathbb{R}^{N_K \times d}$ ,  $V \in \mathbb{R}^{N_V \times d}$ , MHA calculates the output  $Y = \text{MHA}(Q, K, V)$  as:

$$\begin{aligned} \text{Input: } & Q, K, V \\ \text{Att: } & A_i = \text{Softmax}\left(\frac{QW_i^Q(KW_i^K)^T}{\sqrt{d}}\right) \\ \text{Head: } & H_i = A_i V W_i^V, \\ \text{Multi-Head: } & H = [H_1, H_2, \dots, H_h] W^H, \\ \text{Output: } & Y = \text{LN}(H + Q), \end{aligned} \quad (1)$$

where  $W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{d \times d_h}$ , and  $W_i^H \in \mathbb{R}^{d \times d}$  are all trainable matrices;  $h$  is the number of attention heads (set to 8 in the experiments) and  $d_h = d/h$ ;  $A_i$  denotes the  $i$ -th attention matrix used to calculate the  $i$ -th head matrix;  $[\cdot]$  means the concatenation operation; and LN is the Layer Normalization operation.

Besides MHA, another important module in Transformer is the Feed-Forward network (FFN):

$$\text{FFN}(Y) = \text{LN}(\text{FC}(\text{ReLU}(\text{FC}(Y)))) + Y, \quad (2)$$

where FC denotes the fully-connected layer and ReLU denotes the rectified linear function.

Given MHA and FFN, we can use them to build a Transformer-based captioning model. For the encoder, it stacks 6 identical blocks where each one contains an MHA and an FFN. Given the output of the former block as the input  $X$ , the next block

calculates its output as:

**Input:**  $X$ ,

**Self-ATT:**  $Y = \text{MHA}(Q = X, K = X, V = X)$ , (3)

**Output:**  $Z = \text{FFN}(Y)$ ,

Note that the variables  $X, Y, Z$  used here are “local variables” for conveniently introducing the work flow of Transformer architecture, whose values will be set to the specific values when introducing the concrete captioning model. In “Self-ATT”,  $Q, K, V$  are set to the same value and this operation is named as self-attention (Vaswani et al., 2017). After stacking 6 blocks defined in Eq. (3), a visual encoder is built. For the first block, its input is the extracted visual feature set of the given image. The output of the last block will be input into the language decoder.

For the decoder, it also stacks 6 identical blocks where each one contains two MHAs and an FFN. Given the output of the former decoder block  $X_D$  and the output of the visual encoder  $X_E$ , the next decoder block calculates its output:

**Input:**  $X_D, X_E$

**Self-ATT:**  $Y_1 = \text{MHA}(Q = X_D, K = V = X_D)$ ,

**ED-ATT:**  $Y_2 = \text{MHA}(Q = Y_1, K = V = X_E)$ , (4)

**Output:**  $Z = \text{FFN}(Y_2)$ ,

Note that in “ED-ATT”,  $Q$  is set to the output of the former decoder block while  $K, V$  are set to the output of the visual encoder, and such operation is called encoder-decoder attention (Vaswani et al., 2017). After stacking 6 blocks defined in Eq. (4), a language decoder is built.

For the first block in the decoder,  $X_D$  in Eq. (4) is set to the word embedding set of the partially generated captions  $S = \{s_1, \dots, s_t\}$  at the  $t$ -th time step. For all the decoder blocks, the input  $X_E$  is set to the same value, which is the output of the visual encoder. The output of the last decoder block  $Z = \{z_1, \dots, z_t\}$  is used to calculate the word distribution of the next word:

$$P(s_{t+1}) = \text{Softmax}(z_t). \quad (5)$$

Given the ground-truth caption  $S^*$ , we can train this model by minimizing the cross-entropy loss:

$$L_{XE} = -\log P(S^*), \quad (6)$$

or by maximizing a reinforcement learning (RL) based reward (Rennie et al., 2017):

$$R_{RL} = \mathbb{E}_{S^s \sim P(S)}(r(S^s; S^*)), \quad (7)$$

where  $r$  is a sentence-level metric for the sampled sentence  $S^s$  and the ground-truth  $S^*$ , e.g., the CIDEr-D (Vedantam et al., 2015) metric.

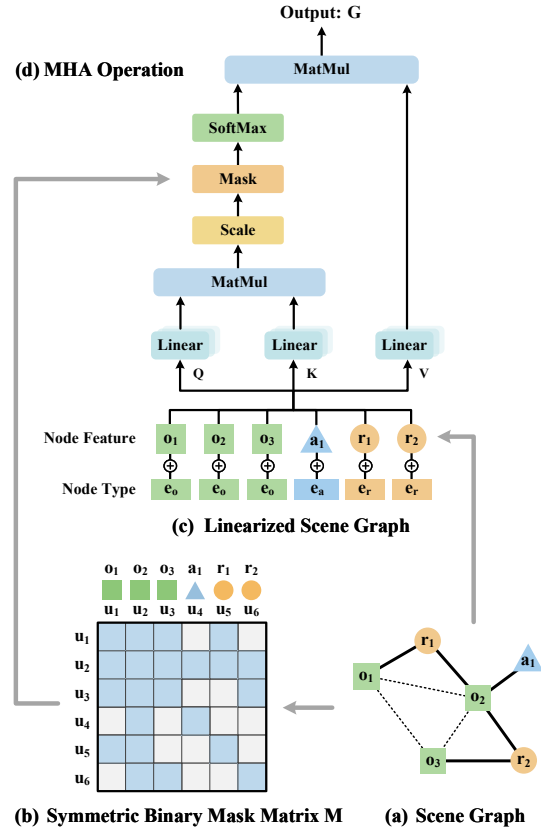


Figure 2: The sketch of the proposed MHA-GNN. In (a), square/triangle/circle demonstrate the object/attribute/relation embeddings, where the dash line means that all the object nodes should be connected for capturing visual contexts. (b) sketches the built binary mask matrix, where the top part shows the original graph embeddings ( $o/a/r$ ). For convenience, we also index the linearized annotation ( $u$ ) in the left and top. (c) shows the linearized scene graph, where the top and bottom parts are respectively node feature and type embeddings. (d) details how MHA achieves the graph operation.

## 4 Transforming Scene Graphs

In this section, we introduce how to revise the Transformer to get our TFSGC. We will first show how to get an MHA-based GNN and then introduce how to design an MOE-based decoder.

### 4.1 MHA-GNN

A visual scene graph (Krishna et al., 2017) contains three kinds of node embeddings: object/attribute/relationship embeddings  $o/a/r$ . These nodes are connected by the following rules: if an object  $o_i$  has an attribute  $a_k$ ,  $o_i$  and  $a_k$  are connected, e.g.,  $o_1$  connects  $a_1$  in Fig. 2 (a). If two objects  $o_i$  and  $o_j$  have the relation  $r_k$ , we connect  $r_k$  with  $o_i$  and  $o_j$ , e.g.,  $r_1$  connects  $o_1$  and  $o_2$ . Given an image, we extract a series of visual features from the im-

age as the object embeddings:  $\{\mathbf{o}_1, \dots, \mathbf{o}_{N_o}\}$ . To get the attribute/relation embeddings, we first use the attribute/relation annotations from VG (Krishna et al., 2017) to train the attribute/relation classifiers to predict the labels. Then we use two learnable embedding layers to respectively transform these labels into the dense attribute/relation embeddings  $\{\mathbf{a}_1, \dots, \mathbf{a}_{N_a}\} / \{\mathbf{r}_1, \dots, \mathbf{r}_{N_r}\}$ .

Given these original node embeddings, GNN will update each one by aggregating the neighbour embeddings. In the previous GNN-LSTM-based models, GNN is usually deployed by the FC layers, which aggregates the contexts by mapping the concatenated neighbour embeddings to the new one (Yao et al., 2018; Yang et al., 2020). However, in our TFSGC, since Transformer is applied as the backbone, we may more hope to use the basic building block of Transformer to deploy the GNN. Such a design principle has two advantages. First, it alleviates the coding implementation difficulty that we do not need to specify some additional GNN operations. Second, which is more important, when the GNN and the Transformer architecture are homogeneous, the whole model will be more easily trained. For example, we do not need to set different training strategies like the learning rate or optimizer to different sub-networks.

Since MHA (Eq. (1)) can learn the context knowledge between the embeddings (Vaswani et al., 2017), it can naturally be used to define the graph operation for aggregating the knowledge. We apply the following two steps to do this. Firstly, as shown in Fig. 2, we linearize the object, attribute, and relation embeddings into one sequence and add the learnable type embeddings as the linearized token set:  $\mathbf{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$ :

$$\begin{aligned} \text{Object: } & \mathbf{u}_i = \mathbf{o}_i + \mathbf{e}_o, & 1 \leq i \leq N_o, \\ \text{Attribute: } & \mathbf{u}_{N_o+i} = \mathbf{a}_i + \mathbf{e}_a, & 1 \leq i \leq N_a, \\ \text{Relation: } & \mathbf{u}_{N_o+N_a+i} = \mathbf{r}_i + \mathbf{e}_r, & 1 \leq i \leq N_r, \end{aligned} \quad (8)$$

where  $\mathbf{e}_o/\mathbf{e}_a/\mathbf{e}_r$  are respectively learnable type embeddings correspond to object/attribute/relation types and  $N = N_o + N_a + N_r$ . For example, in Fig. 2,  $N_o/N_a/N_r$  is 3/1/2, the objects  $\mathbf{o}_{1:3}$  become  $\mathbf{u}_{1:3}$ , the attributes  $\mathbf{a}_1$  becomes  $\mathbf{u}_4$ , and the relations  $\mathbf{r}_{1:2}$  become  $\mathbf{u}_{5:6}$ .

After linearizing, the topological knowledge of the graph is lost, *i.e.*, this token sequence does not show which two nodes are connected or not. To remedy such knowledge, we use a symmetric binary mask matrix  $\mathbf{M} \in \mathbb{R}^{N \times N}$  to all Transformer blocks to mask the attention weights of the un-

connected nodes to control whether two nodes are connected or not. If two nodes  $\mathbf{u}_i$  and  $\mathbf{u}_j$  are connected in the original scene graph, we set  $M_{i,j} = 1$  and  $M_{i,j} = 0$  otherwise. Specifically, the values of  $\mathbf{M}$  are set as follows: 1) If  $\mathbf{o}_i$  has one attribute  $\mathbf{a}_j$ , we set  $M_{i,j+N_o} = 1$ , *e.g.*,  $\mathbf{o}_2$  ( $\mathbf{u}_2$ ) and  $\mathbf{a}_1$  ( $\mathbf{u}_4$ ) in Fig. 2 are connected and  $M_{2,4} = 1$ . 2) If  $\mathbf{r}_k$  connects with  $\mathbf{o}_i$  and  $\mathbf{o}_j$ , we set  $M_{i,k+N_o+N_a} = M_{j,k+N_o+N_a} = 1$ , *e.g.*,  $\mathbf{r}_1$  ( $\mathbf{u}_5$ ) connects  $\mathbf{o}_1$  ( $\mathbf{u}_1$ ) and  $\mathbf{o}_2$  ( $\mathbf{u}_2$ ) and  $M_{1,5} = M_{2,5} = 1$ . 3) All the object nodes are connected with each other since they are visual features that their contexts play a key role in captioning (Herdade et al., 2019). Thus  $\forall i, j \leq N_o, M_{i,j} = 1$ . 4) Since the scene graph is an undirected graph,  $\mathbf{M}$  should be symmetric:  $M_{i,j} = M_{j,i}$ .

After getting  $\mathbf{U}$  and  $\mathbf{M}$ , we can revise the Transformer encoder to get our MHA-GNN. Specifically, we use  $\mathbf{U}$  as the input of the encoder defined in Eq.(3) and revise the **Att** operation in Eq. (1) as the following **Masked Att** operation:

$$\mathbf{A}_i = \text{Softmax}(\mathbf{M} \odot \frac{\mathbf{Q}\mathbf{W}_i^Q(\mathbf{K}\mathbf{W}_i^K)^T}{\sqrt{d}}), \quad (9)$$

where  $\odot$  denotes the element-wise product. In this way, the graph operation is defined by MHA. Specifically, for each node embedding, it is updated by weighted summing its neighbour embeddings, where the weights are from the attention heads  $\mathbf{A}_i$  calculated by the **Att** operation in Eq. (1). During weighted summing, the binary matrix control whether two nodes are connected or not. Note that the edge type is implicitly embedded in **Att** operation due to the added node type embedding. For example, after adding node type embeddings  $\mathbf{e}_o$  and  $\mathbf{e}_r$  to the object and relation embeddings  $\mathbf{o}$  and  $\mathbf{r}$ , respectively, the inner-product becomes\*:

$$(\mathbf{o} + \mathbf{e}_o)^T(\mathbf{r} + \mathbf{e}_r) = \mathbf{o}^T\mathbf{r} + \mathbf{e}_o^T\mathbf{r} + \mathbf{e}_r^T\mathbf{o} + \mathbf{e}_o^T\mathbf{e}_r, \quad (10)$$

where the right three terms are affected by the node type embedding. Thus, when the edge type changes (*e.g.*, the object-relation edge changes to object-attribute edge), the corresponding node type embeddings also change (*e.g.*,  $\mathbf{e}_r$  changes to  $\mathbf{e}_a$ ), which means Eq. (10) encodes the knowledge of edge types into the embeddings.

By stacking more such layers, the receptive field is increasing and thus each node can be updated by aggregating more neighbour embeddings, which naturally follows the design principle of GNN (Battaglia et al., 2018). The output graph

\*For convenience, we omit the trainable matrices  $\mathbf{W}^Q, \mathbf{W}^K$  in Eq. (1) in this inner-product operation.

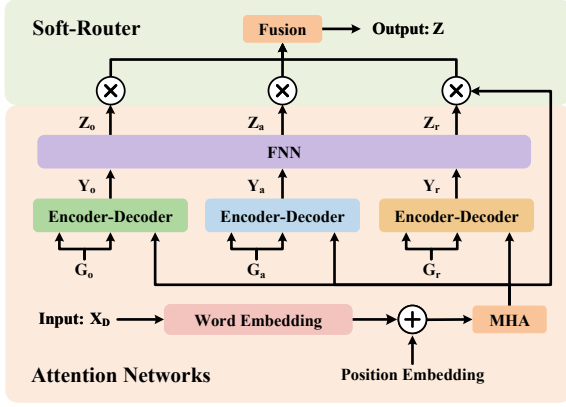


Figure 3: The sketch of the MOE-decoder, we use different colours to denote different experts: green/blue/yellow correspond to object/attribute/relation experts.

embedding set  $G$  are input to the decoder for captioning.

## 4.2 MOE-decoder

As mentioned before, a caption contains different kinds of words for describing diverse visual patterns, *e.g.*, nouns/adjectives/verbs for object/s/attributes/relations (Yang et al., 2019b), which suggests that different experts should be used to address diverse visual knowledge for generating the corresponding words. Motivated by this idea, we design an MOE-based (Jacobs et al., 1991; Du et al., 2022) language decoder to discriminate diverse graph embeddings by setting three encoder-decoder attention layers as different experts. As shown in Fig. 1 (c), the graph embeddings  $G = \{g_1, \dots, g_N\}$  output from the MHA-GNN can be naturally divided according to the original token types in the scene graph: object/attribute/relation sets  $G_o = \{g_1, \dots, g_{N_o}\}$ ,  $G_a = \{g_{N_o+1}, \dots, g_{N_o+N_a}\}$ ,  $G_r = \{g_{N_o+N_a+1}, \dots, g_N\}$ . Then we only need to input them into the corresponding experts for discriminating them. Fig. 3 sketches the designed MOE-based decoder, which is got by revising the decoder defined in Eq. (4) as:

$$\begin{aligned}
 \text{Input: } & X_D, G_o, G_a, G_r \\
 \text{SA: } & X = \text{MHA}(Q = K = V = X_D), \\
 \text{EXP}_O: & Y_o = \text{MHA}(Q = X, K = V = G_o), \\
 \text{EXP}_A: & Y_a = \text{MHA}(Q = X, K = V = G_a), \\
 \text{EXP}_R: & Y_r = \text{MHA}(Q = X, K = V = G_r), \\
 \text{FFN: } & Z_o, Z_a, Z_r = \text{FFN}(Y_o, Y_a, Y_r), \\
 \text{SR: } & Z = \text{SR}(Z_o, Z_a, Z_r, X)
 \end{aligned} \tag{11}$$

where  $\text{EXP}_O$ ,  $\text{EXP}_A$ , and  $\text{EXP}_R$  denote three different experts (encoder-decoder attentions) used to address object, attribute, and relation embeddings, respectively. They have the same structure while

with different parameters.

Note that the input  $X_D$  is the word embeddings of the partially generated captions and at  $t$ -th step,  $X_D = \{x_D^1, \dots, x_D^t\}$ . Then all the  $X/Z_o/Z_a/Z_r$  also contain  $t$  elements, *e.g.*,  $Z_o = \{z_o^1, \dots, z_o^t\}$ . Soft Router (SR) calculates an ensemble embedding  $z$  at each time step to construct the embedding set  $Z = \{z^1, \dots, z^t\}$ . Specifically, for each element  $x/z_o/z_a/z_r$  in  $X/Z_o/Z_a/Z_r$ , a corresponding output  $z$  can be got<sup>†</sup>:

$$\begin{aligned}
 \text{Input: } & x, z_o, z_a, z_r, \\
 \text{ATT: } & \alpha = \{\alpha_o, \alpha_a, \alpha_r\} \\
 & = \text{Softmax}(\{x^T z_o, x^T z_a, x^T z_r\})
 \end{aligned} \tag{12}$$

$$\text{Output: } z = \alpha_o z_o + \alpha_a z_a + \alpha_r z_r,$$

where ATT operation calculates the soft routing weights, since  $x$  accumulates the context knowledge of the partially generated caption, it can help judge which kind of word should be generated at the next step. For example, if the last word of this partially generated caption is an adjective “black”, the next word is more like to be a noun and thus  $\alpha_o$  should be a large value for using more object embeddings instead of the other embeddings.

## 5 Experiments

### 5.1 Datasets, Metrics, and Implementation Details

**Datasets. MSCOCO.** We use MSCOCO (Lin et al., 2014) to validate our TFSGC. This dataset has 123,287 images and each one is labeled with 5 captions. We use two splits in the experiments: the offline Karpathy split (113,287/5,000/5,000 train/val/test images) and the Official online split (82,783/40,504/40,775 train/val/test images).

**Visual Genome (Krishna et al., 2017)** provides scene graph annotations for training the scene graph parser. We follow (Yang et al., 2020) to filter the noisy dataset (*e.g.*, lots of labels only appear a few times in the dataset) by removing the attribute/relation labels appearing less than 2000 times and use the remained 103/64 attribute/relation labels to train the attribute/relation classifiers.

**Implementation Details.** In the experiments, we use three kinds of visual features to exhaustively compare to the other SOTA models, which are BUTD (Anderson et al., 2018), Patch (Liu et al., 2021), and VinVL (Zhang et al., 2021). During training and inference, for BUTD/Patch/VinVL, we respectively follow (Yang et al., 2020) and

<sup>†</sup>For convenience, we remove the superscript representing for different time steps of each embedding.

VinVL’s official parser<sup>‡</sup> to parse SGs, where the latter is more powerful. For all the visual features, we set the batch size to 20 and use Adam (Kingma and Ba, 2014) as the optimizer. For BUTD/Patch/VinVL features, We sequentially use cross-entropy loss (Eq. (6)) and RL-based reward (Eq. (7)) to train the models 20/20/30 and 30/30/30 epochs. For BUTD/Patch/VinVL features, the learning rate used in the cross-entropy stage is initialized as  $5e^{-4}/5e^{-4}/2e^{-5}$  and both decayed by 0.8 every 5 epochs, the learning rate used in the RL-reward stage is reset to  $5e^{-5}/2e^{-5}/5e^{-6}$  and both decayed by 0.8 every 5 epochs. During inference, we use beam search where the beam size is 5. We evaluate the captions by CIDEr-D (Vedantam et al., 2015), BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE (Lin, 2004) and SPICE (Anderson et al., 2016).

## 5.2 Ablation Studies

To confirm the effectiveness of the proposed MHA-GNN and MOE-decoder, we deploy exhaustive ablations as follows. Note that we use BUTD feature in this section. **BASE**: We directly use the classic Transformer architecture. **SG**: We incorporate the scene graphs into the Transformer by using the node embeddings without any graph operations. **MLP-GNN**: We apply MLP-based Graph Neural Network (Xu et al., 2019) for embedding the scene graphs. **MHA-GNN w/o e**: We apply the proposed MHA-GNN while do not use node type embedding. **MHA-GNN**: We apply the proposed MHA-GNN and keep the decoder unchanged as BASE. **GNN-FC**: We remove the binary mask matrix  $M$  from TFSGC. **MOE**: We use the proposed MOE-decoder and do not use GNN but input the original node embeddings into the decoder. **TFSGC**: We apply the integral TFSGC.

Table 1 compares the similarity metrics of the ablation models. Firstly, we can find that the integral TFSGC achieves the highest scores, which confirms its effectiveness. Next, we respectively compare the ablation models to validate the effectiveness of the proposed MHA-GNN and MOE-decoder. By comparing MLP-GNN, SG, and BASE, it is easy to see that using GNN can gain more profits than only using node embeddings. Furthermore, we can find that MHA-GNN has higher CIDEr than MLP-GNN, which suggests that designing GNN by MHA is more powerful than by

| Models        | B@4         | M           | R           | C            | S           |
|---------------|-------------|-------------|-------------|--------------|-------------|
| BASE          | 38.4        | 28.5        | 58.1        | 128.7        | 22.0        |
| SG            | 38.5        | 28.6        | 58.1        | 129.0        | 22.2        |
| MLP-GNN       | 38.9        | 28.8        | 58.4        | 129.5        | 22.4        |
| MHA-GNN w/o e | 39.1        | 28.9        | 58.5        | 130.1        | 22.4        |
| MHA-GNN       | 39.5        | 29.2        | 58.9        | 130.9        | 22.8        |
| GNN-FC        | 39.2        | 29.0        | 58.3        | 130.5        | 22.3        |
| MOE           | 39.2        | 28.8        | 58.5        | 130.1        | 22.5        |
| TFSGC         | <b>39.8</b> | <b>29.6</b> | <b>59.3</b> | <b>132.3</b> | <b>23.4</b> |

Table 1: The performances of various ablation models. The metrics: B@N, M, R, C, and S denote BLEU@N, METEOR, ROUGE-L, CIDEr-D, and SPICE.

| Models        | nouns       | adjectives  | verbs       | prepositions |
|---------------|-------------|-------------|-------------|--------------|
| BASE          | 43.8        | 12.7        | 20.3        | 40.1         |
| SG            | 44.2        | 13.5        | 20.9        | 40.8         |
| MLP-GNN       | 45.4        | 14.8        | 21.8        | 41.6         |
| MHA-GNN w/o e | 48.8        | 16.4        | 24.3        | 44.3         |
| MHA-GNN       | 49.6        | 17.0        | 24.7        | 44.6         |
| MOE           | 48.4        | 16.6        | 24.4        | 44.0         |
| TFSGC         | <b>52.8</b> | <b>18.2</b> | <b>25.8</b> | <b>45.8</b>  |

Table 2: The recalls (%) of five part-of-speech words.

MLP in the Transformer architecture. Next, to see whether discriminating the graph embeddings is beneficial or not, we can compare MHA-GNN with MHA-GNN w/o e and find that using node type embedding performs better. From the comparison between GNN-FC and TFSGC, it can be seen that when removing  $M$ , the graph becomes a fully-connected graph, which introduces more noises. Also, it can be seen that MOE and TFSGC respectively achieve better performances than SG and MHA-GNN, which validates the effectiveness of the MOE-decoder.

Besides evaluating these ablation models by similarities, in Table 2, we calculate the recalls of the words with different POS to evaluate the descriptiveness. Table 2 shows that the captions generated from TFSGC have the highest recalls, suggesting that TFSGC generates the most descriptive captions. Also, we can find that both the proposed MHA-GNN (MHA-GNN vs. MLP-GNN) and MOE-based decoder (MOE vs. SG) can boost the recalls, suggesting both of them improve the descriptiveness. Moreover, We use Stanford Parser to get the POS to train the route weights by the cross-entropy loss, then the CIDEr of TFSGC boosts from 132.3 to 132.9, suggesting the advantage of using POS knowledge. Fig. 4 shows 4 examples of the captions generated from diverse models, where we can see that TFSGC generates more descriptive captions. BASE generates less descriptive captions since it does not use scene graphs and thus loses se-

<sup>‡</sup>[https://github.com/microsoft/scene\\_graph\\_benchmark](https://github.com/microsoft/scene_graph_benchmark)

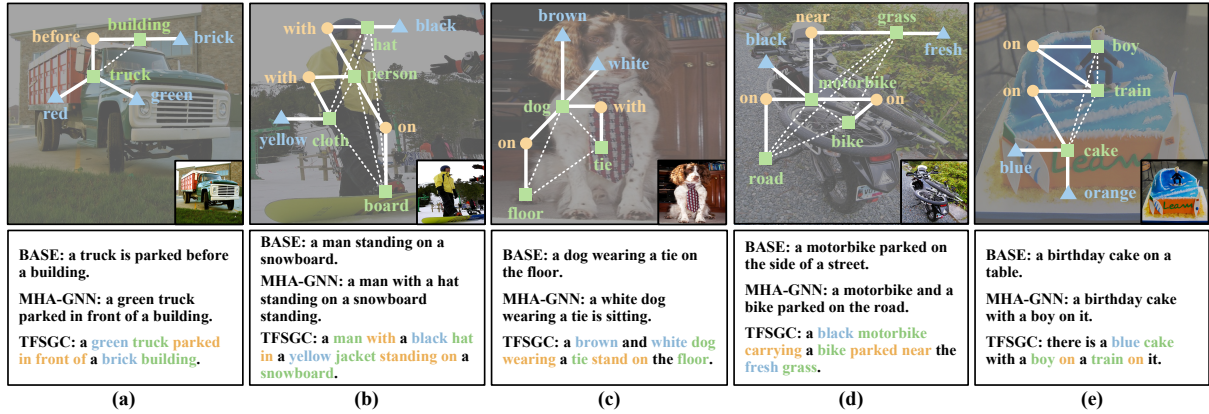


Figure 4: The captions generated by **BASE**, **MHA-GNN**, and **TFSGC**. It can be found that TFSGC generates more descriptive captions, *e.g.*, it describes more attributes like “green truck” and “brick building” in (a), or using more fine-grained nouns like “jacket” in (b). Diverse colors in TFSGC denote these words use more knowledge from different experts: green/blue/orange corresponds to object/attribute/relation expert, which is got by checking which one of  $\alpha_o/\alpha_a/\alpha_r$  in Eq (12) is the largest.

mantic knowledge compared with SG. MHA-GNN does not use the soft router and thus can not select the most suitable experts for generating corresponding words, which may lose certain details. Also, we show that which expert is more responsible for the words with different POS, *e.g.*, in Fig. 4 (b), the adjective “yellow” is generated by using more knowledge from the attribute expert.

### 5.3 Comparisons with SOTA

Recently, various SOTA captioning models with diverse settings have been proposed, including using different language decoders (LSTM, GRU, and Transformer), different visual features, and whether distilling knowledge from large-scale pre-training models (CLIP (Radford et al., 2021) or VinVL (Zhang et al., 2021)). In fairness, we compare our TFSGC with the models that also use Transformer-based decoder by three features: BUTD (Anderson et al., 2018), Patch (Liu et al., 2021), and VinVL (Zhang et al., 2021).

We compare with the following SOTA models: **ETA** (Li et al., 2019), **ORT** (Herdade et al., 2019), **AoANet** (Huang et al., 2019),  $\mathcal{M}^2$  **Transformer** (Cornia et al., 2020), **CATT** (Yang et al., 2021b), **APN** (Yang et al., 2021a), **DLCT** (Luo et al., 2021), **X-Transformer** (Pan et al., 2020), **Up-Down** (Anderson et al., 2018), **SGAE** (Yang et al., 2019a), **NG-SAN** (Guo et al., 2020), **PureT** (Wang et al., 2022), and **ViTCAP** (Fang et al., 2022). Specifically, ETA and ORT are preliminary Transformer-based models; AoANet and X-Transformer design stronger attention mech-

anisms; CATT and  $\mathcal{M}^2$  Transformer introduce additional memory networks; and APN exploits the hidden tree structures. We set the training batch size of TFSGC\* to 50 as X-Transformer for fairly comparing. These methods use BUTD features, while ViTCAP and PureT use Patch features. RSTNet and ViTCAP distill knowledge from pre-trained vision-language BERTs. Note that in VinVL (Zhang et al., 2021), they use OSCAR (Li et al., 2020) as the captioner, while OSCAR is trained on 6.5 million image-text pairs. To fairly compare, we input VinVL feature into the classic Transformer to generate the captions, which is denoted as **VinVL(Transformer)**. Note that we do not compare with extra large scale models trained by millions of image-text pairs or with excessively large parameter sizes, such as **RSTNet** (Zhou et al., 2020) and **LEMON** (Hu et al., 2022). All the results are shown in Table 3.

From Table 3 we can find that TFSGC almostly achieves the highest scores in different settings, *i.e.*, achieving 132.3, 138.6, 139.5 CIDEr scores when using BUTD, Patch, and VinVL features, respectively. DLCT’s visual features extractor (Jiang et al., 2020) is stronger than TFSGC(BUTD) (Anderson et al., 2018) and thus DLCT is a little better than TFSGC(BUTD). Among these compared methods, although other SOTAs do not use scene graphs, they usually have some other training burdens. For example, APN and X-Linear apply more complex attention mechanisms and it requires more computation resource for well-training them, while our TFSGC only apply the simplest attention opera-



| Models   | Cross-Entropy Loss |             |             |              |             | CIDEr optimization |             |             |              |             |
|--|--------------------|-------------|-------------|--------------|-------------|--------------------|-------------|-------------|--------------|-------------|
|  | B@4                | M           | R           | C            | S           | B@4                | M           | R           | C            | S           |
| <b>BUTD feature</b>                            |                    |             |             |              |             |                    |             |             |              |             |
| ETA  | 37.1               | 28.2        | 57.1        | 117.9        | 21.4        | 39.3               | 28.8        | 58.9        | 126.6        | 22.7        |
| ORT  | 35.5               | 28.0        | 56.6        | 115.4        | 21.2        | 38.6               | 28.7        | 58.4        | 128.3        | 22.6        |
| AoANet   | 37.2               | 28.4        | 57.5        | 119.8        | 21.4        | 38.9               | 29.2        | 58.8        | 129.8        | 22.4        |
| $\mathcal{M}^2$ Transformer                    | -                  | -           | -           | -            | -           | 39.1               | 29.2        | 58.6        | 131.2        | 22.6        |
| CATT   | 37.3               | 28.5        | 57.4        | 119.0        | 21.5        | 39.4               | 29.3        | 58.9        | 131.7        | 22.8        |
| APN  | -                  | -           | -           | -            | -           | 39.6               | 29.2        | 59.1        | 131.8        | 23.0        |
| DLCT   | -                  | -           | -           | -            | -           | <b>39.8</b>        | 29.5        | 59.1        | <b>133.8</b> | 23.0        |
| TFSGC  | <b>38.1</b>        | <b>28.6</b> | <b>57.7</b> | <b>120.2</b> | <b>21.9</b> | <b>39.8</b>        | <b>29.6</b> | <b>59.3</b> | 132.3        | <b>23.4</b> |
| <b>BUTD feature &amp; Larger Batch Size</b>    |                    |             |             |              |             |                    |             |             |              |             |
| X-Transformer                                  | 38.2               | <b>28.8</b> | <b>58.0</b> | 122.0        | 21.9        | 39.7               | 29.5        | 59.2        | 132.8        | 23.2        |
| TFSGC*   | <b>38.4</b>        | <b>28.8</b> | 57.8        | <b>122.3</b> | <b>22.1</b> | <b>39.9</b>        | <b>29.8</b> | <b>59.4</b> | <b>133.0</b> | <b>23.4</b> |
| <b>Large Visual-language model pretraining</b> |                    |             |             |              |             |                    |             |             |              |             |
| RSTNet   | -                  | -           | -           | -            | -           | 40.1               | 28.9        | 59.5        | 135.6        | 23.3        |
| LEMON <sub>base</sub>                          | 40.3               | 30.2        | -           | 133.3        | 23.3        | 41.6               | 30.1        | -           | 142.7        | 25.1        |
| LEMON <sub>huge</sub>                          | 41.5               | 30.8        | -           | 139.1        | 24.1        | 42.6               | 31.4        | -           | 145.5        | 25.5        |
| <b>Patch feature</b>                           |                    |             |             |              |             |                    |             |             |              |             |
| PureT  | -                  | -           | -           | -            | -           | 40.9               | <b>30.2</b> | <b>60.1</b> | 138.2        | 24.2        |
| ViTCAP <sub>small</sub>                        | 35.7               | 28.8        | 57.6        | 121.8        | 22.1        | 40.1               | 29.4        | 59.4        | 133.1        | 23.0        |
| TFSGC  | <b>38.8</b>        | <b>29.4</b> | <b>58.2</b> | <b>122.2</b> | <b>22.3</b> | <b>41.4</b>        | 30.1        | <b>60.1</b> | <b>138.6</b> | <b>24.4</b> |
| <b>VinVL feature</b>                           |                    |             |             |              |             |                    |             |             |              |             |
| VinVL(Transformer)                             | 35.4               | 28.6        | 57.5        | 121.5        | 21.3        | 40.6               | 30.0        | 59.8        | 137.3        | 23.7        |
| TFSGC  | <b>38.5</b>        | <b>29.2</b> | <b>58.8</b> | <b>122.7</b> | <b>22.4</b> | <b>41.7</b>        | <b>30.5</b> | <b>60.4</b> | <b>139.5</b> | <b>24.6</b> |

Table 3: The performances of SOTA methods on MS-COCO Karpathy split. All models used are single models.

| Models                      | B@4         |             | M           |             | R           |             | C            |              |
|-----------------------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------|--------------|
|                             | c5          | c40         | c5          | c40         | c5          | c40         | c5           | c40          |
| Up-Down                     | 36.9        | 68.5        | 27.6        | 36.7        | 57.1        | 72.4        | 117.9        | 120.5        |
| SGAE                        | 37.8        | 68.7        | 28.1        | 37.0        | 58.2        | 73.1        | 122.7        | 125.5        |
| ETA                         | 38.9        | 70.2        | 28.6        | 38.0        | 58.6        | 73.9        | 122.1        | 124.4        |
| APN                         | 38.9        | 70.2        | 28.8        | 38.0        | 58.7        | 73.7        | 126.3        | 127.6        |
| NG-SAN                      | 38.8        | 70.2        | 29.0        | 38.4        | 58.7        | 74.0        | 126.3        | 128.6        |
| TFSGC <sup>S</sup>          | <b>39.0</b> | <b>70.9</b> | <b>29.1</b> | <b>38.4</b> | <b>58.9</b> | <b>74.4</b> | <b>127.2</b> | <b>129.8</b> |
| AoANet                      | 39.4        | 71.2        | 29.1        | 38.5        | 58.9        | 74.5        | 126.9        | 129.6        |
| X-Transformer               | <b>39.9</b> | 71.8        | 29.5        | 39.0        | 59.3        | 74.9        | 129.3        | 131.4        |
| $\mathcal{M}^2$ Transformer | 39.7        | <b>72.8</b> | 29.4        | 39.0        | 59.2        | 74.8        | 129.3        | 132.1        |
| TFSGC <sup>E</sup>          | 39.7        | 72.6        | <b>29.6</b> | <b>39.2</b> | <b>59.5</b> | <b>75.1</b> | <b>129.6</b> | <b>133.1</b> |

Table 4: The scores on the MS-COCO online test server. <sup>S</sup> indicates single model, <sup>E</sup> indicates ensemble model.

tion. Moreover, as detailed in Sec 4.1 of ViTCAP, it applies much more training data (9.9M image-text pairs from 4 datasets including VG) to pre-train a concept network to get more powerful discrete tags, while we only use one dataset VG to get scene graphs and achieve better performance, which suggests that connecting discrete tags into a graph is a useful strategy if the discrete tags is not very powerful. Moreover, TFSGC achieves comparable performances with LEMON (Hu et al., 2022) which uses 200 million image-text pairs and uses 12/24-layer, 768/1024-hidden units Transformers(base/huge). To sum up, the advantage of TFSGC is that it can effectively embed and discriminate the semantic knowledge from scene graphs to balance the (usu-

ally more) burden of using more training data or of training more complex networks.

We also submit the single model TFSGC<sup>S</sup> and 4-ensembled model TFSGC<sup>E</sup> trained by regional-based features into the online server for testing, where the results are shown in Table 4. From this table, we can discover that both TFSGC<sup>S</sup> and TFSGC<sup>E</sup> have the highest CIDEr-D scores, which further confirm the effectiveness of the proposed TFSGC.

## 6 Conclusion

We proposed to transform the scene graph into captions (TFSGC) by a simple and homogeneous network. Specifically, we use MHA to design the GNN by linearizing the scene graph and remedying the lost topological knowledge with a binary mask matrix. Furthermore, we add learnable type embedding and design an MOE-based decoder to distinguish node embeddings for more descriptive captions. At last, we compared TFSGC with various SOTA models and demonstrated that our model can achieve comparable performances to some strong benchmarks.

## Limitations

There are two major limitations of the proposed TFSGC. The first one is that the effectiveness of TFSGC depends on the quality of the scene graph. Since MSCOCO does not have SG annotations, we evaluate the parsers in Visual Genome: for BUTD/Patch and VinVL, the recall@50 of relation/attribute are respectively 65.2/68.4 and 73.4/76.6. We use VinVL’s SGs in TFSGC(BUTD) and CIDEr improves from 132.3 to 133.1, suggesting better SGs are beneficial. If the scene graph quality is poor, then TFSGC will not achieve good performance. When an incorrect node in the scene graph, it also affects the output of the caption. *e.g.*, in Fig. 4 (e), the correct object label should be "surfboard" instead of "train". In this paper, we use Visual Genome, which contains abundant and useful scene graph annotations for parsing effective scene graphs, but current performance is not the best, and we will improve the scene graph parser based on the latest scene graph parsing methods in the future.

The second limitation of TFSGC is that if the visual features contain abundant attribute or relation knowledge, then the improvement of TFSGC compared with the classic Transformer will be weakened. For example, compared with the BUTD feature case where the relative improvement of CIDEr-D is 3.6 (TFSGC-BASE in Table 1), the VinVL feature is more powerful since it is trained by much more data samples with more semantic labels, thus the relative improvement is lower, which is 2.2 (TFSGC-VinVL(Transformer) in Table 3).

## Acknowledgements

This work is supported by National Key R&D Program of China (2018AAA0100104, 2018AAA0100100), National Science Foundation of China (62206048), Natural Science Foundation of Jiangsu Province (BK20220819, BK20211164), and Young Elite Scientists Sponsorship Program of Jiangsu Association for Science and Technology Tj-2022-027.

## References

Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, pages 382–398. Springer.

Peter Anderson, Xiaodong He, Chris Buehler, Damien

Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 5, page 6.

Satanjeev Banerjee and Alon Lavie. 2005. 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.

Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.

Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. 2020. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10578–10587.

Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pages 5547–5569. PMLR.

Zhiyuan Fang, Jianfeng Wang, Xiaowei Hu, Lin Liang, Zhe Gan, Lijuan Wang, Yezhou Yang, and Zicheng Liu. 2022. Injecting semantic concepts into end-to-end image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18009–18019.

Zhe Gan, Chuang Gan, Xiaodong He, Yunchen Pu, Kenneth Tran, Jianfeng Gao, Lawrence Carin, and Li Deng. 2017. Semantic compositional networks for visual captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5630–5639.

Longteng Guo, Jing Liu, Xinxin Zhu, Peng Yao, Shichen Lu, and Hanqing Lu. 2020. Normalized and geometry-aware self-attention network for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10327–10336.

Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. 2019. Image captioning: Transforming objects into words. In *Advances in Neural Information Processing Systems*, pages 11137–11147.

Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. 2022. Scaling up vision-language pre-training for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17980–17989.

- Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. 2019. Attention on attention for image captioning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4634–4643.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, and Xinlei Chen. 2020. In defense of grid features for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10267–10276.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*.
- Chenliang Li, Haiyang Xu, Junfeng Tian, Wei Wang, Ming Yan, Bin Bi, Jiabo Ye, Hehong Chen, Guohai Xu, Zheng Cao, et al. 2022. mplug: Effective and efficient vision-language learning by cross-modal skip-connections. *arXiv preprint arXiv:2205.12005*.
- Guang Li, Linchao Zhu, Ping Liu, and Yi Yang. 2019. Entangled transformer for image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Xiujun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020. Oscar: Object-semantics aligned pre-training for vision-language tasks. *ECCV 2020*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*.
- Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. 2017. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *CVPR*, volume 6, page 2.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2018. Neural baby talk. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7219–7228.
- Yunpeng Luo, Jiayi Ji, Xiaoshuai Sun, Liujuan Cao, Yongjian Wu, Feiyue Huang, Chia-Wen Lin, and Rongrong Ji. 2021. Dual-level collaborative transformer for image captioning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2286–2293.
- Victor Milewski, Marie-Francine Moens, and Iacer Calixto. 2020. Are scene graphs good enough to improve image captioning? *arXiv preprint arXiv:2009.12313*.
- Kien Nguyen, Subarna Tripathi, Bang Du, Tanaya Guha, and Truong Q Nguyen. 2021. In defense of scene graphs for image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1407–1416.
- Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. 2020. X-linear attention networks for image captioning. In *CVPR*, pages 10971–10980.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. 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.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *CVPR*, volume 1, page 3.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*, pages 5998–6008.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *CVPR*.
- Li Wang, Zechen Bai, Yonghua Zhang, and Hongtao Lu. 2020. Show, recall, and tell: image captioning with recall mechanism. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12176–12183.
- Yiyu Wang, Jungang Xu, and Yingfei Sun. 2022. End-to-end transformer based model for image captioning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8053–8072.
- Haiyang Xu, Ming Yan, Chenliang Li, Bin Bi, Songfang Huang, Wenming Xiao, and Fei Huang. 2021. E2e-vlp: end-to-end vision-language pre-training enhanced by visual learning. *arXiv preprint arXiv:2106.01804*.
- Haiyang Xu, Qinghao Ye, Ming Yan, Yaya Shi, Jiabo Ye, Yuanhong Xu, Chenliang Li, Bin Bi, Qi Qian, Wei Wang, et al. 2023. mplug-2: A modularized multi-modal foundation model across text, image and video. *arXiv preprint arXiv:2302.00402*.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How powerful are graph neural networks? In *International Conference on Learning Representations*.
- Fuzhao Xue, Ziji Shi, Futao Wei, Yuxuan Lou, Yong Liu, and Yang You. 2022. Go wider instead of deeper. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8779–8787.
- Xu Yang, Chongyang Gao, Hanwang Zhang, and Jianfei Cai. 2021a. Auto-parsing network for image captioning and visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2197–2207.
- Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. 2019a. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10685–10694.
- Xu Yang, Hanwang Zhang, and Jianfei Cai. 2019b. Learning to collocate neural modules for image captioning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4250–4260.
- Xu Yang, Hanwang Zhang, and Jianfei Cai. 2020. Auto-encoding and distilling scene graphs for image captioning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. 2021b. Causal attention for vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9847–9857.
- Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. 2018. Exploring visual relationship for image captioning. In *Computer Vision—ECCV 2018*, pages 711–727. Springer.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. 2023. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*.
- Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image captioning with semantic attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4651–4659.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588.
- Yiwu Zhong, Liwei Wang, Jianshu Chen, Dong Yu, and Yin Li. 2020. Comprehensive image captioning via scene graph decomposition. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 211–229. Springer.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13041–13049.

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*The Limitation section is given in page 10.*
- A2. Did you discuss any potential risks of your work?  
*This work deal with image captioning, which is a task that does not have potential moral risks.*
- A3. Do the abstract and introduction summarize the paper’s main claims?  
*The Abstract and Introduction sections are given in the page 1.*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*Left blank.*

- B1. Did you cite the creators of artifacts you used?  
*The creators of artifacts are cited in the “Datasets, Metrics, and Implementation Details” section in the page 6.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*When we use the artifacts, we cite them in the paper, for example, the metrics are cited in the “Datasets, Metrics, and Implementation Details” section in the page 6.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*We use the evaluation metrics like BLEU or CIDEr-D to measure the similarities between the generated and ground-truth captions, which is consistent with their intended use.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*Such information is introduced in the Dataset section in the page 6.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*Not applicable. Left blank.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*Not applicable. Left blank.*

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a [question on AI writing assistance](#).*

**C  Did you run computational experiments?**

*Left blank.*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

*No response.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*No response.*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*No response.*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*No response.*

**D  Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*No response.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*No response.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*No response.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*No response.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*No response.*