

SGEITL: Scene Graph Enhanced Image-Text Learning for Visual Commonsense Reasoning

Zhecan Wang^{1*}, Haoxuan You^{1*}, Liunian Harold Li², Alireza Zareian¹, Suji Park¹, Yiqing Liang¹, Kai-Wei Chang², Shih-Fu Chang¹

¹ Columbia University

² University of California, Los Angeles

{olinzhecanwang, haoxuanyou}@gmail.com

Abstract

Answering complex questions about images is an ambitious goal for machine intelligence, which requires a joint understanding of images, text, and commonsense knowledge, as well as a strong reasoning ability. Recently, multimodal Transformers have made a great progress in the task of Visual Commonsense Reasoning (VCR), by jointly understanding visual objects and text tokens through layers of cross-modality attention. However, these approaches do not utilize the rich structure of the scene and the interactions between objects which are essential in answering complex commonsense questions. We propose a **Scene Graph Enhanced Image-Text Learning (SGEITL)** framework to incorporate visual scene graph in commonsense reasoning. In order to exploit the scene graph structure, at the model structure level, we propose a multihop graph transformer for regularizing attention interaction among hops. As for pre-training, a scene-graph-aware pre-training method is proposed to leverage structure knowledge extracted in visual scene graph. Moreover, we introduce a method to train and generate domain relevant visual scene graph using textual annotations in a weakly-supervised manner. Extensive experiments on VCR and other tasks show significant performance boost compared with the state-of-the-art methods, and prove the efficacy of each proposed component.

1 Introduction

Visual Commonsense Reasoning (Zellers et al. 2019) is a new addition to Vision-and-Language (VL) research, which has drawn significant attention in the past few years. Different from the conventional Visual Question Answering (VQA) task (Goyal et al. 2017a), VCR requires deeper understanding of the scene and commonsense knowledge. It also requires reasoning ability such as cause-effect and next-step prediction of the presented activity. The state-of-the-art (SOTA) performance on VCR has been improved by a series of recent works on Transformer-based VL models (Su et al. 2020; Chen et al. 2020; Li et al. 2019). In those models, visual object features from image and word embeddings from the image-question-answer pairs are jointly fed into the conventional Transformer model, which consists of several layers of multi-head attention both within each modality and across the two.

*Equal Contribution

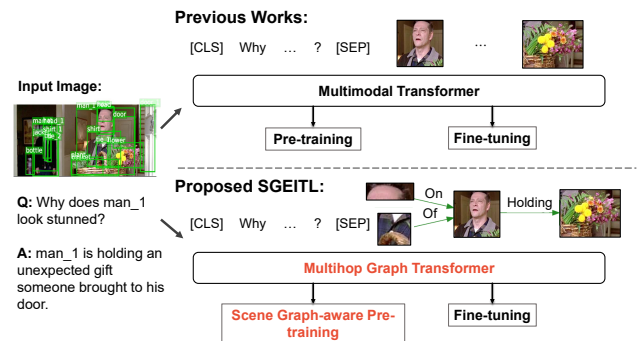


Figure 1: Overview. We propose to incorporate visual scene graph in VL model’s pre-training, fine-tuning and conventional Transformer’s structure for joint multimodal representations.

In spite of their great performance, most existing Transformer-based models reduce the image into a bag of object features extracted using a pre-trained object detector (Tan and Bansal 2019; Chen et al. 2020; Su et al. 2020). However, since the VCR task requires a comprehensive understanding of the visual scene and the commonsense reasoning beyond that, object information alone may not be sufficient to understand visual scenes. Therefore, a more comprehensive visual representation along with the paired text prompt is essential for progress in this field.

A scene graph represents an image as objects and their interactions, providing a structured understanding of the visual scene. Due to its compact, yet comprehensive representation, it has been used for several applications, such as image retrieval (Johnson et al. 2015), image captioning (Yang et al. 2019), image synthesis (Johnson, Gupta, and Fei-Fei 2018), and visual question answering (Antol et al. 2015). Nevertheless, key limitations have prevented us from utilizing scene graphs for VCR. Firstly, it is unclear how the top-performed Transformer-based methods that assume a set of visual and textual tokens as input can handle the information present in a graph structure. Secondly, existing pre-training tasks of VL models mainly use naive random masking when calculating pre-training loss. Such masking steps (Su et al. 2020; Li et al. 2020b) lack the ability to consider graph connections if tokens are from visual scene graphs. Thirdly, the only public dataset with sufficient annotation to train scene graph gen-

eration (SG) models is Visual Genome (VG) (Krishna et al. 2017) which is heavily biased in relationship classes (Zellers et al. 2018) and its classes also suffer from a severe semantic gap compared to the questions and answers in VCR.

In this work, we aim to address the three aforementioned challenges by incorporating visual scene graph into Transformer’s model structure, pre-training and fine-tuning. Those adjustments in VL models’ training pipelines are model-agnostic, and could be unified to a Scene Graph Enhanced Image-Text Learning (SGEITL) schema to assist popular VL models’ pre-training and fine-tuning, as shown in Fig. 1.

SGEITL takes text tokens and visual scene graph tokens as input. Then the proposed multihop graph structure is applied to existing Transformer’s attention to learn joint multimodal representations through a set of pre-training tasks. Modified on top of current VL Transformers, multihop graph Transformer can dynamically adjust the attention value between tokens within multiple hops based on mutual distance in a scene graph. In addition, for pre-training, we propose a scene-graph-aware pre-training method with triplet-based masking. Under this strategy, each triplet (subject/object/predicate) of visual scene graphs can have at most one component masked. Compared with random masking strategy in previous works, our method preserves the necessary local context for the masked node in each triplet unit and also emphasizes the semantic and structural difference among nodes in the scene graph with three prediction heads for subject, object and predicate predictions respectively. Lastly, for solving the problem of limited scene graph annotation and existing biased distribution of scene graph dataset, we further propose an innovative SGG method, independent from VL Transformer models, driven by the weak supervision of parsed text of each image-text pair.

The main contributions of our paper are twofolds.

1. To the best of our knowledge, this is the first work to demonstrate that the structure information in visual scene graph is helpful for complex semantic visual question answering task, such as the VCR task.
2. We systematically introduce multiple ways of enhancing current VL models with visual scene graphs, which enable a more structural understanding towards visual scene and hence facilitate multimodal learning. Our schema includes the following components and is generalizable to different VL models.
 - I Our multihop graph structure inherits from the conventional Transformer could explicitly model the multihop reasoning on visual scene graph.
 - II Our scene-graph-aware pre-training method takes into consideration of visual scene graph’s structure and could assist VL models for learning better visual representation.
 - III Our new SGG method, Text-VSPNet can learn to generate scene graphs even for datasets without scene graph annotation. With weakly-supervision from text, the model can generate semantically-rich and target domain-relevant visual scene graphs.

2 Related Work

2.1 Vision and Language Representations

Multimodal Transformers: Combining vision and language is essential for various tasks such as visual question answering (Antol et al. 2015; Zhang et al. 2016; Goyal et al. 2017b) and visual reasoning (Suhr et al. 2019; Zellers et al. 2019). The emerging trend of multimodal Transformers has shown promising progress in these fields, where general-purpose models are pre-trained on image-caption pairs, and then fine-tuned on downstream by transferring their rich representations from pre-training. These models are simpler, yet more effective and versatile (Lu et al. 2019; Tan and Bansal 2019; Li et al. 2019; Su et al. 2020; Chen et al. 2020; Li et al. 2020a). The idea behind is to extend the language Transformer model (*e.g.* BERT (Devlin et al. 2019)) to vision by adding visual tokens. The model in most cases is exactly the same as BERT. Some variants such as LXMERT (Tan and Bansal 2019) employ a two-stream architecture which applies Transformers separately on each modality, followed by a multimedia Transformer across modalities. In contrast to these approaches, our proposed multihop graph Transformer takes visual scene graphs as the visual input instead of a bag of object features, and incorporates the scene structure through a graph-based attention mechanism.

Multimodal Pre-training: Since (Devlin et al. 2019; ?), pre-training on general representations followed by transferring knowledge on downstream tasks has been a very popular strategy for many different tasks. Besides text-only pre-training, former works *e.g.* (Li et al. 2019; Tan and Bansal 2019; Lu et al. 2019) use a similar set of multimodal pre-training methods and demonstrate their efficacy in downstream multimodal tasks (Antol et al. 2015; Zhang et al. 2016; Goyal et al. 2017b; Suhr et al. 2019; Zellers et al. 2019). Most of them apply naive random masking and expect to learn strong multimodal representations. Following them, later methods like (Li et al. 2020b; Chen et al. 2020) focus and improve on semantic alignments between vision and text domains via variant pre-training methods related to using contrastive loss, word region alignment and object tags as anchors. However, these explicit alignments are limited to syntactic matching of tokens without considering structure information. A recent method (Yu et al. 2020) takes a step further to obtain scene-related graphs parsed from the paired texts during pre-training. However, the texts are mostly short, having very limited visual descriptive tokens related to the scene. This would surely bring an insufficient description of the visual scene by the solely text-parsed graph. Also, it still inherits from the former to use bags of tokens for the visual representation. This would cause inconsistency in representations between vision and language, hinder the learning of multimodal alignment and cannot ensure to have sufficient local context during reconstruction. Differently, in this work, we focus on incorporating visual scene graphs from images during pre-training to help models adapt to the structural scene domain.

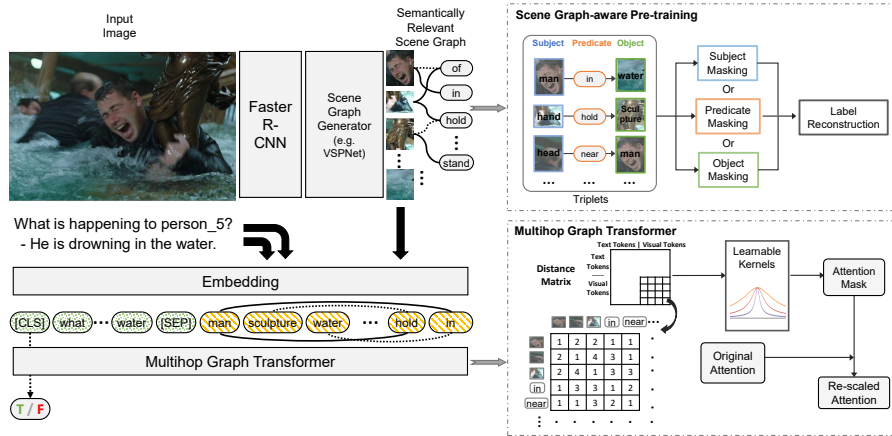


Figure 2: Pipeline. We first extract a scene graph from a given image, using either a conventional SGG method, or the proposed Text-VSPNet. Then we employ the multihop graph Transformer to tackle the scene graph input with regularized attention mechanism among multiple hops. For pre-training, a novel scene-graph-aware pre-training schema is utilized to conduct masking on triplet-level.

Scene Graph Generation: Scene Graph Generation (SGG) has attracted much attention since proposed in (Xu et al. 2017) and shown potential benefits for several downstream visual reasoning tasks (Krishna et al. 2018; Johnson, Gupta, and Fei-Fei 2018; Jiang et al. 2020; Jiaxin Shi 2019; Yang et al. 2019; Zhu et al. 2020). The goal of SGG is to take an image and extract a set of objects and pairwise interactions, which will form a graph where predicates are either edges (Yang et al. 2018) or nodes (Zareian, Karaman, and Chang 2020). Although most SGG methods need intensive supervision to be trained, recently VSPNet (Zareian, Karaman, and Chang 2020) is proposed as a more generalized form of SGG, which does not require bounding box supervision. Although VSPNet was originally trained on Visual Genome (VG), considering its capability of learning from weak supervision, we adopt VSPNet to the task of VCR to generate semantically rich and task-related scene graphs by pre-training on Visual Genome (Krishna et al. 2017) and then fine-tuning on the text annotations on VCR (Zellers et al. 2019).

3 Scene Graph Enhanced Image-Text Learning

In this section, we will first explain the high-level architecture of how we equip VL models with visual scene graphs in Sec. 3.1. Then we introduce the multihop graph Transformer in Sec. 3.3, scene-graph-aware pre-training strategy in Sec. 3.4, and explain how to generate visual scene graph semantically-relevant to text domain in Sec. 3.5.

3.1 Model Overview

The overall pipeline of our approach is illustrated in Fig. 2. For input, we take generated visual scene graph from an image and text tokens as input, and extract joint multimodal representations. For preserving the graph connections between visual nodes, we upgrade the conventional Transformer with the proposed multihop graph for learning multihop feature aggregation. To accommodate the visual scene

graph’s structure in input, we introduce a novel triplet-based masking strategy in the pre-training phase. Moreover, we introduce a new model-agnostic way to obtain visual scene graphs that are more relevant to the VCR domain by training an SGG model directly on VCR text annotations.

The key difference between scene graph enhanced VL learning and previous multimodal VL learning is twofold. (1) Each image is modeled as a scene graph, consisting of object features, relation features and connectivity information. (2) With bringing visual scene graph into previous VL learning frameworks and our proposed adjustments in three perspectives, we not only transform the generated graph to be domain-relevant via weak supervision but also effectively exploit the relation features and graph structures through scene-graph-aware pre-training and modeling.

3.2 Scene Graph as Visual Input

Each input sample consists of a text segment and a scene graph represented as a set of objects and relations and their connections. For the text embedding layer, following BERT (Devlin et al. 2019), we tokenize the input sentence into a sequence of WordPieces (Wu et al. 2016) and obtain its BERT Embedding. For the scene graph embedding layer, we represent each object and relation by two types of features: (1) position features, (2) visual features. Position features of objects and relations are respectively transformed from bounding box coordinates and union box coordinates. For visual features of objects, we take the region of interest (ROI) features from the object detector; for visual features of relations, we take the relation features before inputting into the final prediction layer in SGG models. Besides, since ROI features and relation features are in different semantic space, two fully-connected layers are applied individually to object visual features and relation visual features.

The obtained text embeddings and visual embeddings are fed into the Transformer-based model.

3.3 Multihop Graph Transformer

In previous VL Transformer models, since the visual input is either a sequence of objects (Lu et al. 2019; Li et al. 2019) or pixels (Huang et al. 2020), they directly use the conventional Transformer (Vaswani et al. 2017), in which every token (both visual and text) can freely attend to each others’ belittling local connections between relevant objects. However, when the input includes a scene graph, the attention should favor the local interactions between connected nodes (objects and predicates) in graph structure since they are essential to the entire visual scene understanding. Toward such motivation, we modified the conventional Transformer with multihop graph attention mechanisms. So the improved Transformer can dynamically adjust the attention weight between visual tokens within multiple hops.

We denote a generated visual scene graph as $G(V, E)$, where V denotes vertices of objects and predicates and E denotes edges connecting them. Following (Yang et al. 2018), besides edges in $G(V, E)$, we also add skip edges between all objects to allow direct information flow among objects and get an enhanced scene graph $G'(V, E')$. Given a sequence of input tokens $\{x_i\}_{i=1}^{i=n}$ (including text tokens, objects tokens and predicates tokens), we first pre-compute a distance matrix D based on $G'(V, E')$. The distance between i -th token and j -th token is defined as the number of edges (hops) in the shortest path. It’s noted that the distance between visual tokens and text tokens is always set to 1 because we want the two modalities to be fully connected. The distance between text tokens is also set to 1 to prevent from disrupting the knowledge learned during BERT pre-training.

$$D_{ij} = \begin{cases} \text{Distance}(i, j) & x_i \text{ and } x_j \text{ in } G'(V, E'), \\ 1 & x_i \text{ and } x_j \text{ in different modalities,} \\ 1 & x_i \text{ and } x_j \text{ in text modality.} \end{cases}$$

Following the input sequence, a L -layer multihop graph attention mechanism is applied. Inside each layer, multiple attention heads are included. At each attention head, the output of previous layer is treated as input to Key W^K and Query W^Q and Value W^V to be projected to hidden dimension.

$$Q = H_{in}W^Q, K = H_{in}W^K, V = H_{in}W^V.$$

Then the attention matrix is computed by a scaled dot-product between Q and K . Similar to (Zareian et al. 2020; Ahmad, Peng, and Chang 2020), we applied a binary mask to zero out the attention values beyond h hops.

$$A = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} + M \right), \quad (1)$$

$$\text{where } M_{ij} = \begin{cases} 0 & D_{ij} \leq h, \\ -\infty & D_{ij} > h. \end{cases} \quad (2)$$

After getting the attention matrix A , the next step of the conventional Transformer is to multiply it with value V to aggregate the features as output. However, in visual scene graph, it’s broadly proved that the predicates largely depend on the objects they link to (Zellers et al. 2018; Chen et al. 2019). Based on above finding, we further hypothesize that

the attention between closer nodes in scene graph should be emphasized and vice versa. That also coheres with the conclusion drawn by (Ahmad, Peng, and Chang 2020) when they utilize dependency graph in text Transformer to help cross-lingual relation and event extraction tasks. Therefore, in our multihop graph attention, built on top of the conventional attention mechanism, a monotonically decreasing function $F : D \rightarrow \mathbb{R}^1$ is introduced to generate an additional attention mask for re-scaling the original attention values in A . And one more normalization function σ is cascaded to make sure the sum to be 1.

$$H_{out} = \sigma(F(D) \cdot A)V, \quad (3)$$

$$\sigma(e_{ij}) = \frac{e_{ij}}{\sum_{j=1}^n e_{ij}}. \quad (4)$$

Considering that different heads should have different functionalities, instead of a fixed handcrafted function as in (Ahmad, Peng, and Chang 2020), we propose to use individual parametric kernel in each head with learnable parameters. Based on experiments, we choose the rational quadratic (RQ) kernel, which is similar to Gaussian kernel but with a smoother prior.

Furthermore, one important property of any scene graph is that it’s always a bipartite graph where the edges are only between objects and predicates. Even though we add the skip connections between objects for information flow, the neighbor distributions of objects and predicates still vary a lot. For example, within scene graph, the second-hop neighbors of predicates and objects are always predicates and objects respectively. With the motivation to disentangle them, we employ two learnable kernels separately for objects and predicates.

$$F(D)_{ij} = \begin{cases} \left(1 + \frac{(D_{ij} - 1)^2}{2 \cdot \alpha_o \cdot l_o^2}\right)^{-\alpha_o} & x_i \in \text{objects} \\ \left(1 + \frac{(D_{ij} - 1)^2}{2 \cdot \alpha_p \cdot l_p^2}\right)^{-\alpha_p} & x_i \in \text{predicates} \\ 1 & \text{else} \end{cases} \quad (5)$$

where α_o (α_p) and l_o (l_p) are 1×1 learnable scale-mixture parameter and 1×1 length-scale parameter for objects (predicates). It’s noted that in order to upgrade the conventional attention to multihop graph attention, we only need to add four parameters per attention head which is quite efficient.

3.4 Scene-Graph-Aware Pre-training

Pre-training has become a conventional practice for multimodal frameworks to fuse different modalities. Former VL models (Chen et al. 2020; Li et al. 2020b, 2019; Gan et al. 2020) have proved that large-scale pre-training on image-text data would significantly benefit the downstream tasks. However, most of the current pre-training methods rely on naive random masking ignoring the semantic difference of tokens, structure information and local context. For instance, in a triplet, $[man, on, floor](subject, relationship, object)$, with random masking, it is possible that both the subject and the object are masked out. Under this situation, it is challenging to find sufficient local context to reconstruct the missing token. Without constraint of context, relationship tokens like $[on]$

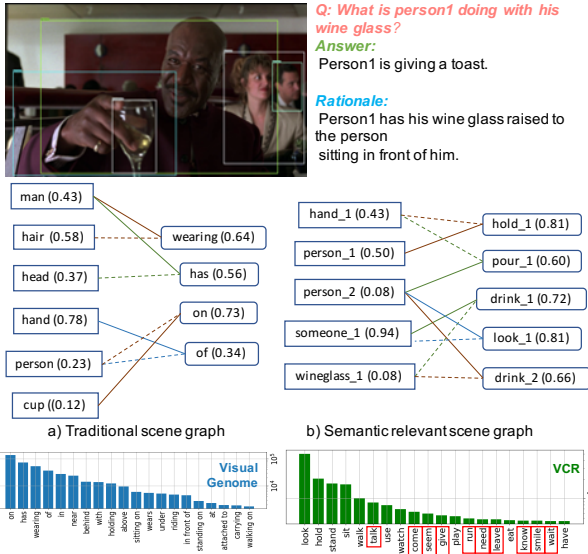


Figure 3: Comparison between traditional and semantically-relevant scene graphs (Best viewed by zoom-in).

would have too many possible combinations of subject and object. To solve this problem, we incorporate visual scene graphs in pre-training steps to preserve the structural context during masking. Visual scene graphs are consisted by triplets. Each triplet has two entities *i.e.* subject as well as object and also one relationship connecting those two. In scene-graph-aware pre-training, we use visual scene graphs as the guidance to selectively mask tokens conditioning on triplet to ensure that sufficient local context would be preserved for reconstructing the masked nodes. This facilitates the model to learn the correlation in neighbor.

The recent work (Yu et al. 2020) claims to obtain "scene graph" parsed from the text and utilize it during pre-training. However, its parsed graph should be more accurately regarded as scene-related graph since it does not closely reflect the object interaction and spatial relationship in the image. It is also not consisted of the conventional triplet (*subject, relationship, object*) thus it reflects more of loose text dependency instead of close visual object interactions. Differently, our visual scene graph is directly extracted from the image. It can closely represent the interaction between objects and spatial relationships. Furthermore, following the triplet partition, our masking strategy is relationship-centric focusing on interactions between objects (Shi and Lin 2019). For each image-text pair, (Yu et al. 2020)'s graph information is never inputted into the VL model and only used for masking text tokens for calculating pre-training loss. Due to the loose text dependency connection of its graph, the pre-training would not allow model to better learn the visual structure of the scene. However, differently, we do feed the visual scene graph into the VL model and even directly mask to modify the visual object input. We have three prediction tasks corresponding to three masked semantic roles (subject, object and relationship). Each visual sample would be equally possible to be assigned for only one of the three types of prediction tasks.

After assigned, for each sample, all the visual objects of the semantic role corresponding to the assigned prediction task, we would randomly select 30% of them and assign a special *[MASK]* token. Therefore, for any triplet, only one of the three nodes would be masked leaving sufficient context information for modeling. Furthermore, we also have three different prediction heads corresponding to the three prediction tasks. This explicit separation allows the model to learn the semantic difference among nodes and better understand the structural knowledge. Under this mechanism, we introduce the pre-training loss of **Masked Node Modeling (MNM)**.

$$\mathcal{L}_{MNM}(\theta) = \mathcal{L}_{sbj}(\theta) + \mathcal{L}_{obj}(\theta) + \mathcal{L}_{rel}(\theta), \quad (6)$$

where θ is the trainable parameters. The model would be supervised to predict the masked node of each triplet, $\langle \mathbf{n}_{s_i}, \mathbf{n}_{o_i}, \mathbf{n}_{r_i} \rangle$ based on the observation of the other two unmasked nodes in the triplet, other surrounding nodes \mathbf{N} and all tokens from the other modality \mathbf{V} , by minimizing the negative log-likelihood:

$$\mathcal{L}_{sbj}(\theta) = -\mathbb{E}_{(\mathbf{n}, \mathbf{v}) \sim D} \log P_{\theta}(\mathbf{n}_{s_i} | \mathbf{n}_{\setminus s_i}, \mathbf{n}_{o_i}, \mathbf{n}_{r_i}, \mathbf{N}, \mathbf{V}),$$

$$\mathcal{L}_{obj}(\theta) = -\mathbb{E}_{(\mathbf{n}, \mathbf{v}) \sim D} \log P_{\theta}(\mathbf{n}_{o_i} | \mathbf{n}_{\setminus o_i}, \mathbf{n}_{s_i}, \mathbf{n}_{r_i}, \mathbf{N}, \mathbf{V}),$$

$$\mathcal{L}_{rel}(\theta) = -\mathbb{E}_{(\mathbf{n}, \mathbf{v}) \sim D} \log P_{\theta}(\mathbf{n}_{r_i} | \mathbf{n}_{\setminus r_i}, \mathbf{n}_{s_i}, \mathbf{n}_{o_i}, \mathbf{N}, \mathbf{V}),$$

where i is a sampled location for masking nodes, s represents subject, o represents object and r represents their relationship.

3.5 Weakly Supervised Scene Graph Generation

Besides modifying current VL models with scene graph related pre-training and structure modification, we also explore to improve visual scene graph at input level. SGG methods are typically trained on the Visual Genome (VG) dataset, which is the only large-scale source for SGG supervision. However, due to the limited and biased distribution of VG, the conventionally produced scene graphs sometimes may not be ideally semantically-relevant for VCR questions. In order to produce more useful scene graphs, inspired by (Zareian, Karaman, and Chang 2020), we innovatively make the first attempt to train SGG with weak supervision from text data.

VSPNet (Zareian, Karaman, and Chang 2020) is a graph-based neural network architecture that takes object proposal features as input and creates a bipartite graph consisting of entities (objects) and predicates (visual relations). The original VSPNet was originally tested on VG dataset to obtain its supervision signal via alignment against labeled ground truth graph. However, most image-text datasets do not provide scene graph annotation, generating scene graphs for those unlabeled datasets is a difficult problem but also very important. It is even more challenging for highly semantic dataset like VCR. Considering this, we make the first attempt to utilize the paired text prompts in image-text datasets to generate weak supervision signal to train our scene graph generator, Text-VSPNet. We use a semantic parser (Shi and Lin 2019) to extract verbs, nouns, and semantic roles from the text prompts and then use a coreference resolution model (Lee et al. 2017) to merge corresponding noun nodes of such graphs. This is followed by a set of

rule-based post-processing steps to create clean graphs as the "ground truth" for training Text-VSPNet. With our flexible weakly-supervised training, the generated scene graph can obtain richer classes of objects and predicates. More importantly, when using the target domain's text prompts as the supervision signal, the generated scene graph is also more semantically-relevant to the target domain's text and images, as shown in the bottom right of Fig. 3. This especially facilitates the utilization of scene graph on highly semantic dataset like VCR. In VCR, we use both questions and answers from the train dataset only to extract pseudo-ground-truth graphs. This process results in 100K images with ground truth semantic graphs that include 672 object classes and 521 predicate classes, which comprise the most frequent words in VCR QAs, and hence are much more likely to be relevant to VCR questions than VG annotations, shown in Fig. 3.

In practice, we further notice that Text-VSPNet is prone to data frequency bias, and it does not learn infrequent classes well. To mitigate this problem, we augment the cross-entropy entity and predicate classification loss terms in VSPNet by focal loss (Lin et al. 2017) and class-balanced loss (Cui et al. 2019). This results in much more diverse predictions that we empirically found essential for VCR.

4 Experiments

In this section, we analyze different components of our framework and compare the performance with the SOTA methods. Additionally, visualization is shown to illustrate the intuition behind our model.

4.1 Ablation Study

We show the effectiveness of the proposed methods on the validation set of VCR. In Tab. 1, we show the experimental results of proposed three components: multihop graph Transformer **HopTrans**, scene-graph-aware pre-training **Pretrain-V** and semantically-relevant scene graphs generated by Text-VSPNet trained by proposed strategy, **SceneGraph+**. Besides, **SceneGraph** means relation features generated by Neural Motif (Zellers et al. 2018) trained on Visual Genome dataset. The baseline of our comparison, the 1st row in table, means that only 36 object features from object detector together with text are inputted into a vanilla Transformer model. The baseline model here is a pre-trained VL-BERT (Su et al. 2020).

Adding Scene Graphs We first investigate the effect of directly using scene graph and scene graph+ for visual representation. We take the top 18 predicates from both the conventional SGG and our Text-VSPNet according to the predicted confidence score¹. Comparing between the 1st, 2nd and 3rd row, we find that solely adding predicate features from both SGGs would already bring around 0.6% improvement on Q→A task and 0.9% on QA→R task.

¹It's noted that our Text-VSPNet and the conventional SSG would share the same backbone, object detector thus they would produce two identical set of visual objects but with different predicates. When merging the two types of generated graphs, we simply merge the same type of visual objects and preserve their connected predicates from both graphs.

Scene Graph-aware Pre-training Given the merged visual scene graph from Text-VSPNet and the conventional SGG as input, We find it beneficial to replace random masking with triplet-based masking for visual tokens. Compared with the 2nd row, the 3rd experiment with triplet-based masking on visual scene graph boosts the performance of Q→A and QA→R by 0.9% and 0.7% respectively.

Multihop Graph Transformer Furthermore, the proposed multihop graph Transformer is utilized to replace vanilla Transformer model. Through 4th row vs. 3rd row and 6th row vs. 5th row, we demonstrate that multihop graph Transformer can benefit scene graph input with 0.4% and 0.3% improvement respectively on Q→A and QA→R by incorporating graph structure in Transformer.

To further study the behavior of our framework with multihop graph Transformer, we give a more comprehensive ablation on two parts: multihop and kernel function $F()$. In Tab. 2, the baseline (1st row) we use for comparison is the model with ordinary scene graph and triplet-based masking on visual scene graph (Pretrain-V), where the model structure is vanilla Transformer. We find if we only consider one hop neighbor with the identity mapping as $F()$, then the performance drops compared with baseline. If we include multiple hops but still keep identity mapping, as in the 3rd row, the drop of performance gets mitigated a little bit. Then we further substitute identity mapping with learnable Gaussian kernels. This gives us positive result with very marginal improvement. After realizing that the learned Gaussian kernel tends to be very sharp, we replace it with a relatively smoother one - Rational Quadratic kernels and finally obtain satisfactory boost. We also find larger number of hops does not always mean better performance, which might be because that predicates are mostly decided by their close neighbors. Above experiments demonstrate the importance of both multihop information and a suitable kernel.

Weakly Supervised Scene Graph Generation We study the benefit of proposed weakly-supervised Text-VSPNet generator. Comparing between the 2rd and 3rd row, besides the scene graph generated by a Neural Motif (Zellers et al. 2018) trained on Visual Genome in fully supervised way, we further incorporate another scene graph from weakly-supervised Text-VSPNet trained on VCR dataset. By providing the predicates that are highly relevant to the text domain of VCR, the performance is further improved by 0.4% on Q→A.

4.2 Experiments on Other Dataset

To prove the generalization ability of the proposed framework, we also show some experiments on GQA and SNLI-VE dataset in Tab. 3. It is important to note that we focus on validating the generalized advantage of our method across different datasets, so no pre-training or in-depth parameter tuning is conducted, which may make the accuracy score lower than some SOTA methods (triplet-based masking is not included in this experiment). From Tab. 3, we can demonstrate the benefits of adding predicate features generated and utilizing graphical structure by multihop graph Transformer. It's noted that the domain of GQA is very close to VisualGenome where (Zellers et al. 2018) is trained. Both

	Model	SceneGraph	SceneGraph+	Pretrain-V	HopTrans	Accuracy		
						Q→A	QA→R	Q→AR
1	Baseline					72.9	75.3	54.9
2		✓				73.1	75.6	55.3
3		✓	✓			73.5	76.2	56.0
4	SGEITL	✓	✓	✓		74.4	76.9	57.2
5		✓	✓	✓	✓	74.9	77.2	57.8

Table 1: Ablation results of SGEITL on VCR validation set. SceneGraph represents adding conventional predicates from SGG trained on VG dataset. Pretrain-V means scene-graph-aware pre-training on visual scene graph. HopTrans means applying multihop graph on the input. SceneGraph+ means adding more domain-relevant predicates from SGG weakly supervised on annotated text input.

Model	Number of Hops	Function $F()$	Q → A
w/o HopTrans	-	-	73.80
w/ HopTrans	1	$F(x) = x$	73.48 (-0.32)
	3	$F(x) = x$	73.60 (-0.20)
	3	Gaussian Kernel	73.89 (+0.09)
	3	Rational Quadratic Kernel	74.91 (+1.11)
	6	Rational Quadratic Kernel	74.28 (+0.48)

Table 2: Ablation on number of hops and different kernel functions in proposed multihop graph Transformer.

Dataset	Model	SceneGraph	HopTrans	SceneGraph+	Accuracy
GQA	Baseline	-	-	-	51.63
	SGEITL	✓			52.71 (+1.08)
	SGEITL	✓	✓		53.25 (+1.62)
SNLI-VE	Baseline	-	-	-	74.06
	SGEITL	✓			74.37 (+0.31)
	SGEITL	✓	✓		74.83 (+0.77)
	SGEITL	✓	✓	✓	75.31 (+1.25)

Table 3: Results on GQA (Hudson and Manning 2019) and SNLI-VE (Xie et al. 2018) validation set.

the image and the question all focus on semantically low-level relationships between objects. Thus, SceneGraph+ is not necessary for this task but we do incorporate it in SNLI-VE dataset.

4.3 Comparison with Benchmarks

Tab. 4 shows the comparison between our proposed model and the SOTA methods on the VCR validation set. The baseline here is the original VL-BERT pre-trained on 290k data, referring to the 1st row in Tab. 4. And VL-BERT+SGEITL refers to the last row in Tab. 4. Compared with baseline, the enhancement of visual scene graph can get a boost of 2.1% on average in three tasks. In this work, we focus on proving the advantage of our framework to incorporate visual scene graph on VL models. Our proposed framework is an add-on module that could be applied on top of different Transformer-based VL models. Thus concerning heavy computational cost of generating additional millions of scene graphs, we do not conduct the large-scale pre-training with out-of-domain data as some of the best-performed models on the VCR leaderboard. However, for rigorous fair comparison, we pre-train those top-performed models including UNITER and ERNIE-ViL with the same scale of in-domain data (290k). It turns out that our VL-BERT+SGEITL can still beat powerful SOTA models with

model	# image-caption in pre-training	Q→A	QA→R	Q→AR
VL-BERT	3.3M	75.5	77.9	58.9
UNITER*	9.5M	76.8	80.2	61.6
ERNIE-ViL	3.8M	78.9	83.7	66.4
ERNIE-ViL	290k	74.1	76.9	56.9
UNITER	290k	73.4	76.0	55.8
UNITER+SGEITL	290k	74.8	76.8	57.4
VL-BERT	290k	72.9	75.3	54.9
VL-BERT+SGEITL	290k	74.9	77.2	57.8

Table 4: Comparison with benchmarks (VLBERT (Su et al. 2020), UNITER (Chen et al. 2020), ERNIE-ViL (Yu et al. 2020)) on VCR validation set. UNITER* denotes that the result is not reported in original paper, but is from official code repository.

the same fair scale of pre-training. These comparisons show that our method is generalizable to other large-scale pre-training scenarios and we leave it for future work.

Submission to VCR Leaderboard: With this submitted version, we further replace our original object detector with a stronger off-the-shell one (Anderson et al. 2018). Then we train our scene graph generator in the same way as before. The test result of VL-BERT+SGEITL with in-domain pre-training and achieve 76.0%/78.0%/59.6% on Q2A/QA2R/Q2AR of VCR leaderboard under the abbreviation, **SGEITL**. The performance coincides with the validation results.

5 Conclusion

We propose a generalized framework to incorporate visual scene graph on top of the current VL models in pre-training, fine-tuning and model structure levels. Our work is the first to effectively prove that visual scene graph could be useful in pre-training and fine-tuning on highly semantic dataset like VCR. Our multihop graph Transformer helps preserve the graphical structure and our scene graph generator is able to use text annotation with weak supervision to generate domain-relevant scene graph. We present extensive experimental results to prove each proposed component’s advantage and compare with the SOTA methods.

Acknowledgements

This work was supported in part by DARPA MCS program under Cooperative Agreement N66001-19-2-4032. The views expressed are those of the authors and do not reflect the official policy of the Department of Defense or the U.S. Government.

References

- Ahmad, W. U.; Peng, N.; and Chang, K.-W. 2020. GATE: Graph Attention Transformer Encoder for Cross-lingual Relation and Event Extraction. *arXiv preprint arXiv:2010.03009*.
- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6077–6086.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*.
- Chen, T.; Yu, W.; Chen, R.; and Lin, L. 2019. Knowledge-Embedded Routing Network for Scene Graph Generation. In *Conference on Computer Vision and Pattern Recognition*.
- Chen, Y.-C.; Li, L.; Yu, L.; Kholy, A. E.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020. Uniter: Universal image-text representation learning. In *ECCV*.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9268–9277.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Gan, Z.; Chen, Y.-C.; Li, L.; Zhu, C.; Cheng, Y.; and Liu, J. 2020. Large-Scale Adversarial Training for Vision-and-Language Representation Learning. In *NeurIPS*.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017a. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017b. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Huang, Z.; Zeng, Z.; Liu, B.; Fu, D.; and Fu, J. 2020. PixelBERT: Aligning Image Pixels with Text by Deep Multi-Modal Transformers. *arXiv preprint arXiv:2004.00849*.
- Hudson, D. A.; and Manning, C. D. 2019. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jiang, X.; Yu, J.; Qin, Z.; Zhuang, Y.; Zhang, X.; Hu, Y.; and Wu, Q. 2020. DualVD: An Adaptive Dual Encoding Model for Deep Visual Understanding in Visual Dialogue. *AAAI*.
- Jiaxin Shi, J. L., Hanwang Zhang. 2019. Explainable and Explicit Visual Reasoning over Scene Graphs. In *CVPR*.
- Johnson, J.; Gupta, A.; and Fei-Fei, L. 2018. Image Generation from Scene Graphs. In *CVPR*.
- Johnson, J.; Krishna, R.; Stark, M.; Li, L.-J.; Shamma, D.; Bernstein, M.; and Fei-Fei, L. 2015. Image Retrieval Using Scene Graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Krishna, R.; Chami, I.; Bernstein, M.; and Fei-Fei, L. 2018. Referring Relationships. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1): 32–73.
- Lee, K.; He, L.; Lewis, M.; and Zettlemoyer, L. 2017. End-to-end Neural Coreference Resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 188–197. Copenhagen, Denmark: Association for Computational Linguistics.
- Li, L. H.; Yatskar, M.; Yin, D.; Hsieh, C.-J.; and Chang, K.-W. 2019. VisualBERT: A Simple and Performant Baseline for Vision and Language. In *arXiv preprint arXiv:1908.03557*.
- Li, L. H.; You, H.; Wang, Z.; Zareian, A.; Chang, S.-F.; and Chang, K.-W. 2020a. Weakly-supervised VisualBERT: Pre-training without Parallel Images and Captions. *arXiv preprint arXiv:2010.12831*.
- Li, X.; Yin, X.; Li, C.; Hu, X.; Zhang, P.; Zhang, L.; Wang, L.; Hu, H.; Dong, L.; Wei, F.; Choi, Y.; and Gao, J. 2020b. Oscar: Object-Semantics Aligned Pre-training for Vision-Language Tasks. *arXiv preprint arXiv:2004.06165*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. *ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 32 (NIPS 2019)*, 13–23.
- Shi, P.; and Lin, J. 2019. Simple bert models for relation extraction and semantic role labeling. *arXiv preprint arXiv:1904.05255*.
- Su, W.; Zhu, X.; Cao, Y.; Li, B.; Lu, L.; Wei, F.; and Dai, J. 2020. VL-BERT: Pre-training of Generic Visual-Linguistic Representations. In *International Conference on Learning Representations*.
- Suhr, A.; Zhou, S.; Zhang, A.; Zhang, I.; Bai, H.; and Artzi, Y. 2019. A Corpus for Reasoning about Natural Language Grounded in Photographs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6418–6428. Florence, Italy: Association for Computational Linguistics.

Tan, H.; and Bansal, M. 2019. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 5100–5111. Hong Kong, China: Association for Computational Linguistics.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6000–6010.

Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.

Xie, N.; Lai, F.; Doran, D.; and Kadav, A. 2018. Visual entailment task for visually-grounded language learning. *arXiv preprint arXiv:1811.10582*.

Xu, D.; Zhu, Y.; Choy, C.; and Fei-Fei, L. 2017. Scene Graph Generation by Iterative Message Passing. In *Computer Vision and Pattern Recognition (CVPR)*.

Yang, J.; Lu, J.; Lee, S.; Batra, D.; and Parikh, D. 2018. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 670–685.

Yang, X.; Tang, K.; Zhang, H.; and Cai, J. 2019. Auto-Encoding Scene Graphs for Image Captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yu, F.; Tang, J.; Yin, W.; Sun, Y.; Tian, H.; Wu, H.; and Wang, H. 2020. ERNIE-ViL: Knowledge Enhanced Vision-Language Representations Through Scene Graph. *arXiv preprint arXiv:2006.16934*.

Zareian, A.; Karaman, S.; and Chang, S.-F. 2020. Weakly Supervised Visual Semantic Parsing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Zareian, A.; You, H.; Wang, Z.; and Chang, S.-F. 2020. Learning Visual Commonsense for Robust Scene Graph Generation. In *ECCV*.

Zellers, R.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. From Recognition to Cognition: Visual Commonsense Reasoning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Zellers, R.; Yatskar, M.; Thomson, S.; and Choi, Y. 2018. Neural Motifs: Scene Graph Parsing with Global Context. In *Conference on Computer Vision and Pattern Recognition*.

Zhang, P.; Goyal, Y.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2016. Yin and Yang: Balancing and Answering Binary Visual Questions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Zhu, Z.; Yu, J.; Sun, Y.; Hu, Y.; Wang, Y.; and Wu, Q. 2020. Mucko: Multi-Layer Cross-Modal Knowledge Reasoning for Fact-based Visual Question Answering. In *International Joint Conference on Artificial Intelligence (IJCAI)*.