

Unbiased Heterogeneous Scene Graph Generation with Relation-Aware Message Passing Neural Network

Kanghoon Yoon^{1*}, Kibum Kim^{1*}, Jinyoung Moon^{3,4}, Chanyoung Park^{1,2†}

¹Dept. of Industrial and Systems Engineering, KAIST, Daejeon, Republic of Korea

²Graduate School of Artificial Intelligence, KAIST, Daejeon, Republic of Korea

³Electronics and Telecommunications Research Institute, 218 Gajeong-ro, Yuseong-gu, Daejeon, Republic of Korea

⁴ ETRI School, University of Science and Technology, 218 Gajeong-ro, Yuseong-gu, Daejeon, Republic of Korea
ykhon08@kaist.ac.kr, kb.kim@kaist.ac.kr, jymoon@etri.re.kr, cy.park@kaist.ac.kr

Abstract

Recent scene graph generation (SGG) frameworks have focused on learning complex relationships among multiple objects in an image. Thanks to the nature of the message passing neural network (MPNN) that models high-order interactions between objects and their neighboring objects, they are dominant representation learning modules for SGG. However, existing MPNN-based frameworks assume the scene graph as a homogeneous graph, which restricts the context-awareness of visual relations between objects. That is, they overlook the fact that the relations tend to be highly dependent on the objects with which the relations are associated. In this paper, we propose an unbiased heterogeneous scene graph generation (HetSGG) framework that captures relation-aware context using message passing neural networks. We devise a novel message passing layer, called relation-aware message passing neural network (RMP), that aggregates the contextual information of an image considering the predicate type between objects. Our extensive evaluations demonstrate that HetSGG outperforms state-of-the-art methods, especially outperforming on tail predicate classes. The source code for HetSGG is available at <https://github.com/KanghoonYoon/hetsgg-torch>.

1 Introduction

Scene graph generation (SGG) is a fundamental visual understanding task, which aims to identify objects from an image and detect their relations (i.e., predicate¹), which can be represented in a triplet format: (subject, predicate, object). A compact structural scene representation is beneficial to various image applications such as visual question answering (Ghosh et al. 2019; Zhang, Chao, and Xuan 2019), image captioning (Yang et al. 2019), and image retrieval (Schroeder and Tripathi 2020; Ramnath et al. 2019). Hence, recent years have seen significant progress in developing methods for SGG.

However, recent SGG methods usually have shown unsatisfactory performance due to the difficulty of learning complex relationships among multiple objects in an image.

*These authors contributed equally.

†Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We use “relation” and “predicate” interchangeably.

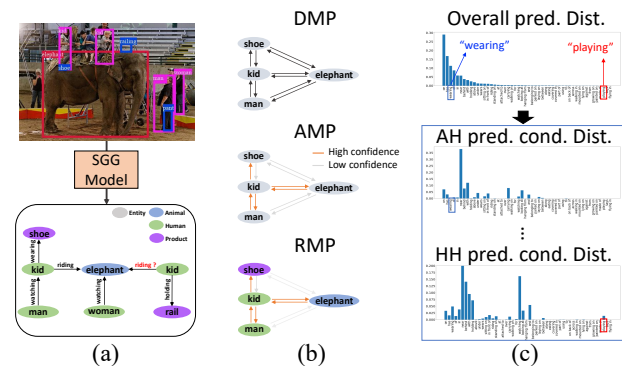


Figure 1: (a): Scene graph represented as a heterogeneous graph. (b) Comparisons of various MPNNs. (c) Overall predicate distribution and predicate distributions conditioned on certain predicate types.

An effective scene graph representation should include the contextual information of an image. For intuitively understanding what it means by contextual information, consider the triplet (kid, riding, elephant) in Figure 1(a). Understanding the context around this triplet, such as “kid holding rail” and “woman watching elephant”, would be helpful for predicting the riding predicate between kid and elephant, compared with the case when kid and elephant are considered independently. Hence, recent studies for SGG mainly focus on capturing the contextual information based on message passing neural networks (MPNNs). Thanks to the nature of the MPNNs that model high-order interactions between objects and their neighboring objects, they can naturally capture the visual context of an object from nearby objects that exist together. Specifically, Graph R-CNN (Yang et al. 2018) designs an MPNN to learn filtered contextual information from neighbors by identifying important objects and relations. Moreover, direction-aware message passing neural network (DMP) (Lin et al. 2020) considers the direction to which the messages are propagated between two objects, and adaptive message passing neural network (AMP) (Li et al. 2021) prevents irrelevant proposal pairs from interacting with each other based on

the confidence of interactions (Figure 1(b)). In summary, recent SGG methods mainly focus on designing new MPNN architectures aiming at capturing the contextual information of an image, thereby increasing the context-awareness of visual relations between objects.

Although the aforementioned existing methods for SGG show effectiveness in understanding the visual context by using advanced MPNNs, they commonly consider the scene graph as a *homogeneous* graph, which in turn restricts the context-awareness of the visual relations between objects. A homogeneous graph considers all its nodes (i.e., objects) and edges (i.e., predicates) to be of a single type. For this reason, existing MPNN-based methods overlook the fact that predicates tend to be highly dependent on the objects with which the predicates are associated. For example, consider the triplet (kid, riding, elephant) in Figure 1(a). Although a kid can ride an elephant, the opposite direction is unlikely to happen, i.e., an elephant usually does not ride a kid because it is usually “Human” that rides “Animal.” However, as existing MPNN-based methods consider the scene graph as a homogeneous graph, “Human”-typed objects and “Animal”-typed objects cannot be distinguished, which eventually fails to explicitly capture such dependencies.

In this paper, we propose an *unbiased* heterogeneous scene graph generation (HetSGG) framework that captures *relation-aware context*. The main idea is to treat each relation differently according to its type. More precisely, we devise a novel message passing layer, called relation-aware message passing neural network (RMP), that aggregates the contextual information of an image considering the predicate type between objects, where the predicate type is determined by the associated object types. For example, given a triplet (subject, predicate, object), if subject and object are assigned “Human (H)” and “Animal (A)” types, respectively, then the type of predicate is “Human-Animal (HA).” More precisely, we first construct a heterogeneous graph based on the objects detected by an off-the-shelf object detector (e.g., Faster R-CNN (Ren et al. 2015)). Then, RMP propagates intra- and inter-relation messages with an attention mechanism to learn the relation-aware context. As described in Figure 1(b), RMP is the only MPNN layer that fully utilizes the heterogeneity of a scene graph thereby capturing the semantics of the relations. It is important to note that RMP is a general framework that subsumes existing mainstream MPNN-based SGG methods, i.e., DMP and AMP. Specifically, RMP is direction-aware by its design, because the predicate types already contain the directional information (e.g., RMP treats kid \rightarrow elephant and elephant \rightarrow kid differently by assigning each relation “HA” and “AH” predicate type, respectively). Besides, RMP is a generalized version of AMP in that the attention function of RMP is predicate type-aware, whereas that of AMP assumes the predicates are of a single type.

Moreover, considering a scene graph as a heterogeneous graph naturally alleviates the biased prediction problem incurred by the long-tail predicate class distribution. Figure 1(c) presents the overall predicate distribution, and the predicate type conditional distributions when considering only the “AH” and “HH” predicate types. We observe that

head predicates in the overall predicate distribution are not anymore considered as head predicates in the “AH” or “HH” predicate type conditional distributions. For example, *wearing* appears at the top-3 in the overall distribution, whereas it rather belongs to a tail predicate class in the “AH” predicate type conditional distribution. On the other hand, *playing* rarely appears in the overall distribution, whereas its proportion increases in the “HH” predicate type conditional distribution. This implies that each predicate type exhibits different head/body/tail predicate class distribution. In this regard, since RMP independently processes all the predicate types, which is followed by an aggregation, HetSGG naturally relieves the bias of the final prediction model. Although existing methods for unbiased prediction (Li et al. 2021; Lin et al. 2020; Tang et al. 2020; Chen et al. 2022; Desai et al. 2021) improve performance on tail predicates, the performance on head predicates is rather sacrificed. On the other hand, HetSGG greatly improves the performance on tail predicates, *while maintaining competitive performance on head predicates*.

Our contributions are summarized as follows:

- To the best of our knowledge, HetSGG is the first work to reformulate the SGG task, in which the scene graph was considered as a homogeneous graph, in the light of a heterogeneous graph.
- HetSGG is model-agnostic in that it can be adopted to any MPNN-based SGG methods. In this work, we adopt HetSGG to two recent SGG methods, i.e., Graph R-CNN (Yang et al. 2018) and BGNN (Li et al. 2021), and demonstrate that HetSGG further improves upon them.
- Through extensive experiments on Visual Genome and Open Images, we demonstrate that HetSGG is superior to state-of-the-art baselines, and that the performance on tail predicate classes in particular improves greatly.

2 Related Work

Scene Graph Generation. Traditional SGG approaches mainly focus on designing advanced neural network architectures to fuse the contextual features in images. They (Chen et al. 2019; Zellers et al. 2018) pass object features of region proposals obtained from an object detector to RNNs aiming at capturing the contextual cues. However, contextualized representations produced by RNNs are dependent on the ordering of the object sequence, which is generated in a heuristic manner without accounting for meaningful contexts. Most recently, MPNN-based methods have shown to be effective for SGG tasks. With an effective architecture that aggregates the neighboring features, MPNNs embed a node-centric representation by combining the contextual elements of neighbors. Several variants of MPNNs have emerged for SGG frameworks as it is essential to design MPNNs with certain characteristics. GPS-Net (Lin et al. 2020) proposes the direction-aware MPNN (DMP) to propagate different messages in two directions, and BGNN (Li et al. 2021) presents confidence-aware adaptive MPNN (AMP) to filter out unnecessary contextual features because the graph constructed from an off-the-shelf

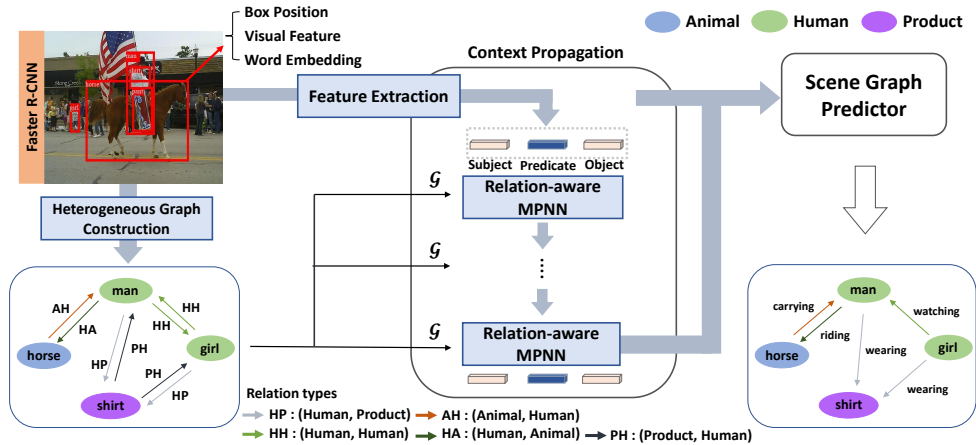


Figure 2: Given an image, a heterogeneous graph is constructed based on the objects detected by an object detector (i.e., Faster R-CNN) from which feature vectors for objects and predicates are extracted. RMP propagates relation-aware messages to the representations of objects and predicates. Finally, the scene graph predictor generates a heterogeneous scene graph.

object detector contains noise. However, the assumption of the existing methods that object and relation types are all equal restricts their performance as the semantics of the relations cannot be captured. In this work, we introduce a new scene graph structure that is based on a heterogeneous graph, and focus on designing a novel MPNN layer, i.e., relation-aware MPNN (RMP), that captures the semantics of predicate types.

Heterogeneous Graph Neural Network. A heterogeneous graph is a powerful tool that can embrace rich semantics and structural information in real-world data. A heterogeneous graph consists of different types of entities (i.e., nodes) and their relations (i.e., edges), which facilitates the modeling of complex semantics of machine learning models. In the graph mining community, a plethora of studies for heterogeneous graph neural networks (HGNN) have been conducted to extract rich information under the heterogeneity by using the *meta-path*, which is a semantically meaningful path defined by relation types (Wang et al. 2019; Sun et al. 2011; Dong, Chawla, and Swami 2017; Hu et al. 2020; Zhang et al. 2019a; Park et al. 2019). As a scene graph represents multiple objects and the relations between them, we argue that it can also be considered as a heterogeneous graph. However, it is non-trivial to apply existing HGNNs developed in the graph mining community for SGG tasks. More precisely, although most existing HGNNs require domain knowledge to generate meta-paths, there are no obvious rules for defining meta-paths in scene graphs. To make the matter worse, since the graph constructed from an off-the-shelf object detector is inherently noisy, it is challenging to apply existing HGNNs that are developed for clean graphs. Hence, in this work, we propose a meta-path-free HGNN that considers the noisy nature of scene graphs. Although there have been comprehensive studies of MPNN layers for SGG, they regard a scene graph as a homogeneous graph, and overlook the node types and predicate types. To the best of our knowledge, this is the first work to consider a scene graph in the SGG task as a

heterogeneous graph.

Long-tail Visual Recognition. Recent SGG methods mainly focus on relieving the long-tail problem of the predicate class distribution for constructing informative scene graphs. One of the prominent approaches to alleviate the long-tail problem is to employ a cost-sensitive loss for SGG. More precisely, some methods introduce novel reweighted losses that leverage semantic constraints of scene graphs (Knyazev et al. 2020; Lin et al. 2020). Moreover, BGNN (Li et al. 2021) proposes a powerful data re-sampling strategy, called bi-level sampling, which combines both image-level and instance-level re-sampling strategies, aiming to make the entire training predicate class distribution towards a uniform distribution. Recently, TDE (Tang et al. 2020) utilizes causal inference in the prediction stage to alleviate the effect of the long-tail predicate class distribution of the dataset. On the other hand, as illustrated in Figure 1(c), our proposed method naturally alleviates the long-tail problem by considering a scene graph as a heterogeneous graph.

3 Problem Definition

In this section, we describe notations used throughout the paper, and introduce our formulation of the SGG task. Let $\mathcal{G} = \langle \mathcal{V}, \mathcal{E}, \mathcal{T}_{\mathcal{V}}, \mathcal{T}_{\mathcal{E}} \rangle$ be a heterogeneous graph, where \mathcal{V} is the set of objects in an image, \mathcal{E} is the set of relations between objects, $\mathcal{T}_{\mathcal{V}}$ is the set of object types, and $\mathcal{T}_{\mathcal{E}}$ is the set of relation types. Objects $u, v \in \mathcal{V}$ and the relation $e_{u \rightarrow v} \in \mathcal{E}$ between u and v have feature vectors x_u, x_v and $x_{u \rightarrow v}$, respectively. Moreover, we denote $\mathcal{Y} = \langle \mathcal{Y}_o, \mathcal{Y}_r \rangle$ as the set of $|\mathcal{Y}_o|$ object classes and $|\mathcal{Y}_r|$ relation (i.e., predicate) classes.

Our goal is to find a function that maps an image I to a scene graph by maximizing the probability $P(\mathcal{G}, \mathcal{Y} | I)$. Specifically, we aim to estimate the graph structural information \mathcal{G} , the object classes (i.e., \mathcal{Y}_o) and the relation classes (i.e., \mathcal{Y}_r) given an image \mathcal{I} . In this work, we introduce a novel framework for generating scene graphs,

called **Heterogeneous Scene Graph Generation (HetSGG)**, which generates a scene graph with typed objects and relations. $P(\mathcal{G}, \mathcal{Y}|\mathcal{I})$ can be factorized as follows: $P(\mathcal{G}, \mathcal{Y}|\mathcal{I}) = P(\mathcal{G}|\mathcal{I})P(\mathcal{Y}_o|\mathcal{G}, \mathcal{I})P(\mathcal{Y}_r|\mathcal{Y}_o, \mathcal{G}, \mathcal{I})$, where $P(\mathcal{G}|\mathcal{I})$ is a heterogeneous graph construction module, $P(\mathcal{Y}_o|\mathcal{G}, \mathcal{I})$ is an object classifier, and $P(\mathcal{Y}_r|\mathcal{Y}_o, \mathcal{G}, \mathcal{I})$ is a predicate classifier. Note that our formulation is different from that of recent SGG methods (Tang et al. 2019; Li et al. 2021; Tang et al. 2020; Chen et al. 2019; Lin et al. 2020; Tang et al. 2019) in terms of predicting and utilizing the heterogeneous information \mathcal{T}_V and \mathcal{T}_E of \mathcal{G} . Thus, our framework is a generalized version of existing homogeneous SGG methods (i.e., HetSGG degenerates to existing homogeneous SGG methods if $|\mathcal{T}_V| = |\mathcal{T}_E| = 1$). The overall architecture of HetSGG is described in Figure 2.

4 Methodology

4.1 Heterogeneous Graph Construction

Initial Graph Construction. We begin by constructing an initial graph based on the objects detected by an off-the-shelf object detector (e.g., Faster R-CNN (Ren et al. 2015))². The object proposals generated by the object detector are defined as nodes, and node pairs are connected by edges (i.e., the initial graph is a fully-connected graph). Then, the feature vector for object $u \in \mathcal{V}$ (i.e., x_u) is obtained by feed-forwarding the concatenation of the bounding box positions of object u , visual features of object u , and the word embedding (i.e., Glove (Pennington, Socher, and Manning 2014)) of the class name of object u . Moreover, the feature vector of the relation between object u and v (i.e., $x_{u \rightarrow v}$) is extracted from the bounding box positions of $e_{u \rightarrow v}$, and the visual features of the union box of the object pairs (u, v) .

Type Inference. To convert the initial graph into a heterogeneous graph, we assign types to the objects and the relations by utilizing the class logits obtained by the object detector. More precisely, the proposal for object u contains the class logit $p_u \in \mathbb{R}^{|\mathcal{Y}_o|}$ obtained from the object detector, where each element of p_u denotes the logit value for a certain object class. We compute the object type logit vector $q_u \in \mathbb{R}^{|\mathcal{T}_V|}$ with a pre-defined function ϕ that maps an object class to an object type, i.e., $\phi: \mathcal{Y}_o \rightarrow \mathcal{T}_V$. We then infer the type for object u with a simple aggregation function, such as sum, mean, and max as in Figure 3(a). For example, consider two detected objects u and v , which belong to **man** and **dog** object classes, respectively, according to the class logits p_u and p_v . Then, the object types are obtained as: $\phi(u) = \text{“Human” (H)}$ and $\phi(v) = \text{“Animal” (A)}$. Note that we employ **Average**(\cdot) as the aggregation function, but we empirically observed that **Sum**(\cdot) performs similarly. Lastly, the type of the relation (i.e., predicate type) between objects u and v (i.e., $e_{u \rightarrow v}$) is automatically determined by the associated object types with a function $\psi: \mathcal{Y}_o \times \mathcal{Y}_o \rightarrow \mathcal{T}_E$, which is $\psi(u, v) = \text{“HA”}$ in our example. In this work, we mainly consider three object types, i.e., “Human (H)”, “An-

²Although the model performance can be improved by using more advanced object detectors, we rely on Faster R-CNN to clearly validate the benefit of our framework.

imal (A)”, and “Product (P)”, which consequently produce nine predicate types, i.e., “HH”, “HA”, “HP”, “AH”, “AA”, “AP”, “PA”, “PH”, and “PP”.

It is important to note that since the heterogeneous graph on which HetSGG is applied is constructed solely based on the above type inference process, accurately inferring the types is crucial for the performance of HetSGG. In Table 2 of Section 5.1, we empirically show that through the above type assignment process based on Faster R-CNN as the object detector, we achieve around 95% accuracy of the object type inference, and we demonstrate further improvements of HetSGG when ground-truth object types are used, i.e., when the object type inference accuracy is 100%. Furthermore, we show that adding another predicate type, i.e., “Landform (L)”, further improves the model performance provided that the type inference is accurate.

4.2 Relation-aware Contextual Representation Learning

Now that we have constructed a heterogeneous graph, we need to learn the representations of objects and their relations. To this end, we propose a novel message passing layer, called **Relation-aware Message Passing** neural network (RMP), to capture relation-aware context. The main idea is to treat each relation differently according to its type. A naïve approach is to assign a type-specific projection matrix W_t to each relation type $t \in \mathcal{T}_E$. However, using a distinct parameter for each type is not only computationally inefficient as the model complexity grows proportional to the number of relation types \mathcal{T}_E , but also is prone to overfitting. Hence, we compose an efficient relation type-specific projection matrix $W_t \in \mathbb{R}^{d \times d}$ as a linear combination of shared basis matrices (Schlichtkrull et al. 2018): $W_t = \sum_{i=1}^b a_{ti} B_i$, where $B_i \in \mathbb{R}^{d \times d}$ is a trainable matrix for basis i , b is the number of bases, and $a_{ti} \in \mathbb{R}$ is a trainable coefficient for relation type t and basis i , which captures the relation type information. A large a_{ti} implies that B_i makes a large contribution when we compose the relation type-specific projection matrix W_t . Based on the reformulation of relation type-specific projection matrix, we can greatly reduce the number of parameters (i.e., from $\mathcal{O}(d^2|\mathcal{T}_E|)$ to $\mathcal{O}(d^2b + |\mathcal{T}_E|b)$, where $b \ll |\mathcal{T}_E|$), which not only facilitates efficient training, but also alleviates the overfitting issue.

The goal of RMP is to capture the relation-aware context and update the representation of objects and relations using the context. In this regard, RMP consists of the following two steps: **1) Edge-wise update for relations**, and **2) Node-wise update for objects**. In a nutshell, in the edge-wise update step, RMP generates relation-specific messages between objects to refine the relation representations. In the node-wise update step, RMP aggregates messages from neighboring relations according to the relation types (i.e., intra-relation aggregation), and then aggregates all the relation type-specific object representations to obtain the final object representations (i.e., inter-relation aggregation). Note that *RMP is model-agnostic in that it can be adopted to any MPNN-based SGG methods*, such as Graph R-CNN (Yang

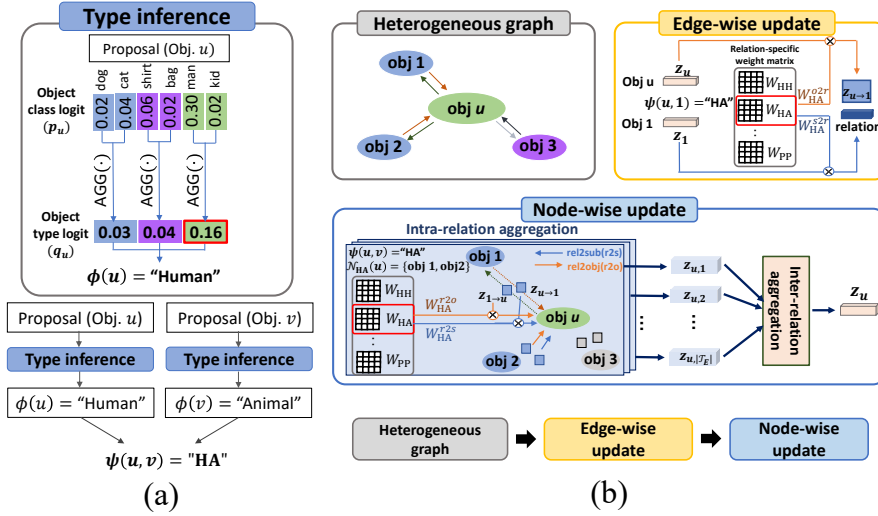


Figure 3: (a) From the class logits of an object proposal, Type inference module computes the probability for each object type with pre-defined mapping. The relation type is defined as Cartesian product of object types. (b) For the edge-wise update, relation type-specific weight matrices are used to generate two-way messages, which are then used to obtain the relation representation. For the node-wise update, the intra-relation aggregation produces the relation-specific messages using the predicate features. Then, the inter-relation aggregation gathers all the relation to make the relation-aware object representation.

et al. 2018) and BGNN (Li et al. 2021). Due to the space limitation, we only explain RMP adopted to Graph R-CNN (i.e., HetSGG) in this paper.

Step 1) Edge-wise update for relations. In this step, given two objects u and v , and the relation between the objects, i.e., $e_{u \rightarrow v}$, the objects propagate contextual information to the relation. More precisely, two directional messages are generated: one is from object u to relation $e_{u \rightarrow v}$, and the other is from object v to relation $e_{u \rightarrow v}$. For an edge $e_{u \rightarrow v}$ whose direction is from object u to object v , the relation representation is computed as follows:

$$z_{u \rightarrow v}^{(l+1)} = z_{u \rightarrow v}^{(l)} + \sigma(\alpha(u, v) W_{\psi(u, v)}^{s2r} z_u^{(l)} + (1 - \alpha(u, v)) W_{\psi(u, v)}^{o2r} z_v^{(l)}) \quad (1)$$

where $z_{u \rightarrow v}^{(l+1)} \in \mathbb{R}^d$ is the relation representation of $e_{u \rightarrow v}$ at the $(l+1)$ -th layer, $W_{\psi(u, v)}^{s2r} \in \mathbb{R}^{d \times d}$ and $W_{\psi(u, v)}^{o2r} \in \mathbb{R}^{d \times d}$ are weight matrices of the relation type $\psi(u, v)$ for the two-way messages (i.e., subject-to-relation and object-to-relation messages given a triplet (subject, predicate, object)), and σ is a non-linear activation function. The initial representations for objects and relations are feature vectors of objects and relations (i.e., $z_{u \rightarrow v}^{(0)} = x_{u \rightarrow v}$ and $z_u^{(0)} = x_u$). Moreover, $\alpha(u, v)$ determines the importance of the messages propagated from subject (i.e., u) and object (i.e., v), and it is formulated as follows: $\alpha(u, v) = \frac{\exp(w^T z_u^{(l)})}{\exp(w^T z_u^{(l)}) + \exp(w^T z_v^{(l)})}$, where $w \in \mathbb{R}^d$ is an attention vector. A large $\alpha(u, v)$ implies that object u is more important than object v for generating the representation of the relation $e_{u \rightarrow v}$.

Step 2) Node-wise update for objects. The node update is

performed based on the relation representation $z_{u \rightarrow v}^{(l+1)}$ obtained in the previous step. The main idea is to aggregate messages from neighboring objects that share the same relation type. More precisely, given an object u and its neighboring objects with the relation type t (i.e., $\mathcal{N}_t(u)$), the representation of the object u regarding the relation type t at $(l+1)$ -th layer (i.e., $z_{u, t}^{(l+1)} \in \mathbb{R}^d$) is computed as follows:

$$z_{u, t}^{(l+1)} = \sum_{v \in \mathcal{N}_t(u)} \alpha_{r2s}(v, t) W_{\psi(u, v)}^{r2s} z_{u \rightarrow v}^{(l+1)} + \alpha_{r2o}(v, t) W_{\psi(u, v)}^{r2o} z_{v \rightarrow u}^{(l+1)} \quad (2)$$

where $W_{\psi(u, v)}^{r2s} \in \mathbb{R}^{d \times d}$ and $W_{\psi(u, v)}^{r2o} \in \mathbb{R}^{d \times d}$ are weight matrices for the relation-to-subject and relation-to-object messages, respectively. The first term produces a message in the perspective of u being the subject in a triplet, and the second term produces a message in the perspective of u being the object in a triplet. We denote the above process as *intra-relation aggregation*. Moreover, $\alpha_{r2s}(v, t)$ and $\alpha_{r2o}(v, t)$ denote the importance of object v among the set of neighbors of object u with relation type t (i.e., $\mathcal{N}_t(u)$). They are defined as follows: $\alpha_{r2s}(v, t) = \frac{\exp(w_{r2s, t}^T z_{u \rightarrow v}^{(l+1)})}{\sum_{q \in \mathcal{N}_t(u)} \exp(w_{r2s, t}^T z_{u \rightarrow q}^{(l+1)})}$,

$\alpha_{r2o}(v, t) = \frac{\exp(w_{r2o, t}^T z_{v \rightarrow u}^{(l+1)})}{\sum_{q \in \mathcal{N}_t(u)} \exp(w_{r2o, t}^T z_{q \rightarrow u}^{(l+1)})}$, where $w_{r2s, t}$ and $w_{r2o, t} \in \mathbb{R}^d$ are attention vectors for relation-to-subject updates and relation-to-object updates considering the relation type t , respectively. A large $\alpha_{r2s}(v, t)$ and $\alpha_{r2o}(v, t)$ imply that under the relation type t , the relation from u to v (i.e., $e_{u \rightarrow v}$) and the relation from v to u (i.e., $e_{v \rightarrow u}$) are crucial for object u , respectively. Finally, we aggregate all the relation type-specific object representations to obtain the final representation of object u (i.e., $z_u^{(l+1)} \in \mathbb{R}^d$): $z_u^{(l+1)} =$

Models	PredCls		SGCls		SGGen	
	mR@50/100	R@50/100	mR@50/100	R@50/100	mR@50/100	R@50/100
RelDN (Zhang et al. 2019b)	15.8/17.2	64.8/66.7	9.3/9.6	38.1/39.3	6.0/7.3	31.4/35.9
Motifs (Zellers et al. 2018)	14.6/15.8	66.0/67.9	8.0/8.5	39.1/39.9	5.5/6.8	32.1/36.9
VCtree (Tang et al. 2019)	15.4/16.6	65.5/67.4	7.4/7.9	38.9/39.8	6.6/7.7	31.8/36.1
G-RCNN (Yang et al. 2018)	16.4/17.2	65.4/67.2	9.0/9.5	37.0/38.5	5.8/6.6	29.7/32.8
MSDN (Li et al. 2017)	15.9/17.5	64.6/66.6	9.3/9.7	38.4/39.8	6.1/7.2	31.9/36.6
Unbiased (Tang et al. 2020)	25.4/28.7	47.2/51.6	12.2/14.0	25.4/27.9	9.3/11.1	19.4/23.2
GPS-Net (Lin et al. 2020)	15.2/16.6	65.2/67.1	8.5/9.1	37.8/39.2	6.7/8.6	31.1/35.9
GPS-Net [‡] (Lin et al. 2020)	29.2/31.4	55.2/57.6	15.9/16.9	36.4/37.5	8.1/9.6	28.4/33.4
NICE-Motif(Li et al. 2022a)	29.9/32.3	55.1/57.2	16.6/17.9	33.1/34.0	12.2/14.4	27.8/31.8
PPDL(Li et al. 2022b)	32.2/33.3	47.2/47.6	17.5/18.2	28.4/29.3	11.4/13.5	21.2/23.9
BGNN [‡] (Li et al. 2021)	30.4/32.9	59.2/61.3	14.3/16.5	37.4/38.5	10.7/12.6	31.0/35.8
BGNN ^{*‡} (Li et al. 2021)	29.2/31.7	57.8/60.0	14.6/16.0	36.9/38.1	10.9/13.1	30.2/34.9
HetSGG [‡]	31.6/33.5	57.8/59.1	17.2/18.7	37.6/38.7	12.2/14.4	30.0/34.6
HetSGG ^{‡+}	32.3/34.5	57.1/59.4	15.8/17.7	37.6/38.5	11.5/13.5	30.2/34.5
Improv.(%)	10.6/8.8	0.0/-1.0	17.8/16.9	1.9/1.6	11.9/9.9	0.0/-0.8

Table 1: Results on Visual Genome (Krishna et al. 2017). Improv. denotes improvements of HetSGG compared with BGNN^{*‡}. ‡ denotes bi-level sampling (Li et al. 2021) is applied, and * denotes results reproduced with authors’ code.




Figure 4: R@100 improvement per class of HetSGG[‡] over BGNN^{*‡} in SGCls task.

$z_u^{(l)} + \frac{1}{|\mathcal{T}_\varepsilon|} \sum_{t=1}^{|\mathcal{T}_\varepsilon|} \sigma(z_{u,t}^{(l+1)})$. We denote the above process as *inter-relation aggregation*. In summary, HetSGG generates relation-specific context through intra- and inter-relation aggregations. Moreover, high-order interactions between objects can be captured by stacking multiple RMP layers.

4.3 Scene Graph Predictor and Model Training

Scene Graph Predictor. After obtaining the representations for objects and relations of an image, we predict their classes to generate the scene graph of the image. First, given the representation of object u (i.e., z_u), we employ a simple linear classifier to obtain the object class probability as follows: $p_u = \text{softmax}(W^{\text{obj}} z_u) \in \mathbb{R}^{|\mathcal{Y}_o|}$, where $W^{\text{obj}} \in \mathbb{R}^{|\mathcal{Y}_o| \times d}$ is the weight matrix for the linear classifier. Next, given the representation of the relation between objects u and v (i.e., $z_{u \rightarrow v}$), we employ a simple linear classifier along with an added bias term regarding the class frequency prior (i.e., $\hat{p}_{u \rightarrow v}$) (Lin et al. 2020; Li et al. 2021; Zellers et al. 2018). The relation class probability is computed as follows: $p_{u \rightarrow v} = \text{softmax}(W^{\text{rel}} z_{u \rightarrow v} + \log \hat{p}_{u \rightarrow v}) \in \mathbb{R}^{|\mathcal{Y}_r|}$, where $W^{\text{rel}} \in \mathbb{R}^{|\mathcal{Y}_r| \times d}$ is the weight matrix for the linear classifier, and $\hat{p}_{u \rightarrow v} \in \mathbb{R}^{|\mathcal{Y}_r|}$ denotes the frequency distribution of predicates given two objects u and v , which is pre-computed from the training data.




Figure 5: Results on the overall, head, body, and tail predicate classes in SGGen task.

Model Training. Using the object class probability (i.e., p_u) and relation class probability (i.e., $p_{u \rightarrow v}$), HetSGG is trained by minimizing the conventional cross-entropy losses for objects (i.e., \mathcal{L}_{obj}) and relations (i.e., \mathcal{L}_{rel}). The final objective function is defined as follows: $\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{obj}} + \mathcal{L}_{\text{rel}}$.

5 Experiment

We evaluate HetSGG compared with state-of-the-arts methods on commonly used benchmark datasets, i.e., Visual Genome (VG) (Krishna et al. 2017) (Section 5.1), and Open Images (OI) V6 (Kuznetsova et al. 2020b) (Section 5.2).

Evaluation Metric. Due to the long-tail problem in SGG tasks, existing methods perform poorly on less frequently appearing predicates. Therefore, following the evaluation protocol of recent SGG methods (Tang et al. 2020), we evaluate SGG models on **mean Recall@K (mR@K)** in addition to the conventional measure **Recall@K (R@K)** (Li et al. 2017). For Open Images dataset, we follow the evaluation protocols of previous works (Kuznetsova et al. 2020b; Lin et al. 2020), and we additionally report weighted mean AP of relationships (wmAP_{rel}), weighted mean AP of phrase (wmAP_{phr}), and the weighted metric score ($\text{score}_{\text{wtd}}$), which is calculated as: $0.2 \times \text{R@50} + 0.4 \times \text{wmAP}_{\text{rel}} + 0.4 \times \text{wmAP}_{\text{phr}}$.

Evaluation Protocol. We evaluate on three conventional SGG tasks (Xu et al. 2017): (1) **Predicate Classifica-**

tion (PredCls), (2) Scene Graph Classification (SGCls), and (3) Scene Graph Generation (SGGen). Note that for SGGen, an object is considered to be correctly detected when its IoU (Intersection over Union) with the ground truth bounding box is greater than 0.5.

Implementation Details. For fair comparisons, we adopt ResNeXt-101-FPN (Xie et al. 2017) and Faster R-CNN (Ren et al. 2015) as the object detector, whose parameters are frozen while training the SGG model. For SGGen task, we select the top 80 object proposals sorted by object scores, and use per-class non-maximal suppression (NMS) (Zellers et al. 2018) at IoU 0.5. To obtain the relation proposals, we localize the union box of the bounding boxes of two objects, and obtain the ROI features using pre-trained Faster R-CNN. For RMP of HetSGG, we set the number of bases (i.e., b) to 8 in VG and 4 in OI, and use four MPNN layers (i.e., $l = 4$). Since HetSGG framework is model-agnostic, we evaluate two versions of HetSGG: i) RMP adopted to Graph R-CNN (Yang et al. 2018) (i.e., HetSGG), and ii) RMP adopted to BGNN (Li et al. 2021) (i.e., HetSGG₊₊).

5.1 Visual Genome

Dataset Details. We follow the same pre-processing strategy that has been widely used for evaluations of SGG (Xu et al. 2017). Specifically, the most frequently appearing 150 object classes and 50 predicate classes are used for evaluation. After preprocessing, each image contains 11.6 objects and 6.2 predicates on average. A total of 108k images are split into training set (70%) and test set (30%).

Comparisons with State-of-the-Art Methods. Table 1 shows the results on various SGG tasks in terms of mR@50/100 and R@50/100. We have the following observations: **1)** HetSGG and HetSGG₊₊ generally outperform all baseline models in various tasks on both metrics. More precisely, HetSGG and HetSGG₊₊ greatly improve mR@50/100, while performing competitively on R@50/100. This verifies that HetSGG effectively relieves the long-tail problem of the predicate class distribution by considering scene graphs as heterogeneous graphs. **2)** It is important to note that for fair comparisons among various MPNN architectures (i.e., DMP of GPS-Net (Lin et al. 2020), AMP of BGNN (Li et al. 2021), and RMP of HetSGG), we also compare with the version of GPS-Net to which bi-level sampling (Li et al. 2021) is applied (i.e., GPS-Net[‡]). We observe that both HetSGG and HetSGG₊₊ outperform GPS-Net[‡] and BGNN^{*‡}, which demonstrates the superiority of RMP of HetSGG. In summary, these results imply that using the heterogeneous information inherent in scene graphs has more powerful predictive performance compared with the methods that utilize MPNNs based on homogeneous scene graphs.

Moreover, Figure 4 shows improvements of HetSGG over BGNN per predicate class. The order of the predicates (i.e., x -axis) is sorted by the frequency of predicates in the training data. We observe that HetSGG generally achieves improvements on all head, body and tail predicate classes, while particularly showing great improvements on tail predicate classes. Similar results are shown in Figure 5. It is im-

portant to note that there is a clear trade-off between the head and tail performance as shown in existing studies (Chen et al. 2022; Desai et al. 2021; Li et al. 2021). However, since the predicates in head classes (e.g., on, has, of, in, etc) are less informative for generating meaningful scene graphs, it is important to achieve a high performance on tail classes. In this regard, HetSGG is superior to other baselines in that *it outperforms baselines in terms of tail performance, while also maintaining a competitive head performance.*

Object Types	Model	SGCls		Type Inf. Acc.(%)
		mR@50/100	R@50/100	
P,H,A	HetSGG [‡]	17.2 / 18.7	37.6 / 38.7	95.3
	HetSGG _{GT} [‡]	17.4 / 19.1	38.0 / 39.0	100
P,H,A,L	HetSGG [‡]	15.9 / 18.2	37.5 / 38.4	90.9
	HetSGG _{GT} [‡]	18.2 / 19.4	39.4 / 40.5	100

Table 2: Analysis on the object types and the accuracy of the type inference module when different object types are defined: P-Product, H-Human, A-Animal, L-Landform.

Analysis on Object Types. In this section, we conduct the following two experiments, and show results in Table 2: i) To investigate the importance of accurate object type inference on the performance of HetSGG, we train HetSGG using the ground-truth object types (i.e., HetSGG_{GT}), and ii) To investigate the benefit of considering the object types in a more fine-grained manner, we add another object type, i.e., “Landform³.” We have the following observations: **1)** HetSGG_{GT} consistently outperforms HetSGG. This implies that accurately inferring the object types is crucial, and that an advanced object detector can further improve HetSGG. **2)** Comparing the performance of HetSGG_{GT}, we observe that considering four object types is superior to considering three object types. This implies that it is beneficial to consider the object types in a more fine-grained manner. **3)** However, comparing the performance of HetSGGs, we observe that when the type inference accuracy is not high enough, the performance degrades when considering more types. This again demonstrates the importance of accurately inferring the object types.

Ablation Studies. **i) On the relation type-specific weight matrices:** To verify the importance of capturing the semantics of relations in the edge-wise and update node-wise update, we remove the relation type-specific weight matrices in each update. More specifically, in the edge-wise step, we replace the $W_{\phi(u,v)}^{s2r}$ and $W_{\phi(u,v)}^{o2r}$ with W^{s2r} and W^{o2r} , respectively, which implies that we treat all the relation types equivalently (i.e., Edge \mathcal{X}). Likewise, in the node-wise update step, we replace $W_{\phi(u,v)}^{r2s}$ and $W_{\phi(u,v)}^{r2o}$ with W^{r2s} and W^{r2o} , respectively (i.e., Node \mathcal{X}). We have the following observations in Table 3: **1)** Adding the relation-specific weight matrix to either edge- and node-wise updates improves the

³We sample “Landform” predicates from “Product” predicates, while “Human” and “Animal” predicates remain unchanged.

Component		Metric	
Edge	Node	mR@100	R@100
✗	✗	15.9	38.6
✓	✗	16.2	38.7
✗	✓	17.7	38.7
✓	✓	18.7	38.7

Table 3: Ablation study on W_t .

# basis (b)	$ \mathcal{T}_{\mathcal{E}} = 9 (= 3^2)$		$ \mathcal{T}_{\mathcal{E}} = 16 (= 4^2)$	
	mR@100	R@100	mR@100	R@100
4	17.2	38.5	17.2	38.5
8	18.7	38.7	18.2	38.4
12	17.6	38.3	17.6	38.5
16	18.2	38.9	17.6	38.3

Table 4: Analysis on the number of basis matrices b .

overall performance of HetSGG. **2)** Considering the relations in both edge- and node-wise updates performs the best, which verifies the benefit of capturing relation-aware context for the SGG task. **ii) On the efficiency of shared basis matrices:** To verify the efficiency of composing a relation type-specific projection matrix as a linear combination of shared basis matrices, we evaluate HetSGG over various number of basis (i.e., b) and relation types (i.e., $|\mathcal{T}_{\mathcal{E}}|$). Table 4 shows that even though $|\mathcal{T}_{\mathcal{E}}|$ increases from 9 to 16, the best performing number of basis matrices does not change, i.e., $b = 8$ performs the best. This implies that RMP efficiently captures the semantics of relations even with a small number of parameters. Hence, we argue that a further benefit of adopting the basis matrices is that the complexity is expected to remain practical even if the number of relation types increases.

Qualitative Results. To verify that HetSGG indeed captures the relation-aware context, we qualitatively compare scene graphs generated by BGNN (Li et al. 2021) and HetSGG for test images. We have the following observations: **1)** Consider the triplet $\langle \text{hand}, \text{of}, \text{boy} \rangle$ in Figure 6(a). We observe that BGNN generates an (incorrect) predicate, i.e., *hold*, even though it does not make sense for “Product” (hand) to hold “Human” (boy). On the other hand, since HetSGG considers the object types, it can avoid such irrational cases, and generate the correct predicate, i.e., *of*. **2)** Consider the triplet $\langle \text{tree}, \text{on}, \text{hill} \rangle$ in Figure 6(b). Although BGNN generates a correct prediction, i.e., *on*, it simply predicts all the predicates related to *tree* as *on*, which is the most frequently appearing predicate in the dataset. It is interesting to see that the prediction generated by HetSGG, i.e., *growing on*, though incorrect, is in fact more realistic. This again verifies that HetSGG generates predicates that appear less frequently in the dataset.

5.2 Open Images

Dataset Details. We closely follow the data processing and evaluation protocols of previous works (Kuznetsova et al. 2020a; Lin et al. 2020). After preprocessing, OI V6 has 301 object classes, and 31 predicate classes, and is split into




Figure 6: Generated scene graphs of HetSGG[†] and BGNN^{*†} in SGCs. (Red: incorrect predictions by BGNN^{*†}, Blue: correct predictions by HetSGG[†], but incorrect by BGNN.)

126,368 train images, 1,813 validation images, and 6,322 test images.

Comparisons with State-of-the-Art Methods. Table 5 demonstrates the experimental results in SGGGen task on Open Images V6. We have the following observations: **1)** HetSGG and HetSGG₊₊ show large improvements on mR@50 implying that the bias prediction problem is greatly alleviated by our proposed framework. **2)** HetSGG and HetSGG₊₊ show competitive performance on R@50 and weighted mAP (wmAP). Considering that R@50 and wmAP are metrics that contradict with the goal of addressing the long-tail problem, this implies that HetSGG alleviates the biased prediction problem, while also maintaining the performance on head classes.

Model	mR@50	R@50	wmAP _{rel}	wmAP _{phr}	score _{wtd}
RelDN	37.2	75.3	32.2	33.4	42.0
VCTree	33.9	74.1	34.2	33.1	40.2
G-RCNN	34.0	74.5	33.2	34.2	41.8
Motifs	32.7	71.6	29.9	31.6	38.9
Unbiased	35.5	69.3	30.7	32.8	39.3
GPS-Net	38.9	74.7	32.8	33.9	41.6
BGNN [†]	40.5	75.0	33.5	34.1	42.1
HetSGG [†]	42.7	76.8	34.6	35.5	43.3
HetSGG [†] ₊₊	43.2	74.8	33.5	34.5	42.2

Table 5: Results on Open Images V6 in SGGGen task.

6 Conclusion

In this work, we proposed an *unbiased heterogeneous* scene graph generation framework, called HetSGG. We devised a novel MPNN architecture, called relation-aware message passing network (RMP), that captures the relation-aware context given the types of objects and associated predicates. By considering a scene graph as a heterogeneous graph, HetSGG alleviated the biased prediction problem incurred by the long-tail predicate class distribution. HetSGG is model-agnostic in that it can be adopted to any MPNN-based SGG methods. Through extensive experiments, we verified that HetSGG outperforms existing SGG methods and generates more realistic scene graphs.

Acknowledgements

This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. 2020-0-00004, Development of Previsional Intelligence based on Long-term Visual Memory Network, and No.2022-0-00077).

References

- Chen, C.; Zhan, Y.; Yu, B.; Liu, L.; Luo, Y.; and Du, B. 2022. Resistance Training using Prior Bias: toward Unbiased Scene Graph Generation. *AAAI*.
- Chen, T.; Yu, W.; Chen, R.; and Lin, L. 2019. Knowledge-embedded routing network for scene graph generation. *CVPR*.
- Chiou, M.-J.; Ding, H.; Yan, H.; Wang, C.; Zimmermann, R.; and Feng, J. 2021. Recovering the unbiased scene graphs from the biased ones. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1581–1590.
- Desai, A.; Wu, T.-Y.; Tripathi, S.; and Vasconcelos, N. 2021. Learning of Visual Relations: The Devil Is in the Tails. In *ICCV*, 15404–15413.
- Dong, X.; Gan, T.; Song, X.; Wu, J.; Cheng, Y.; and Nie, L. 2022. Stacked Hybrid-Attention and Group Collaborative Learning for Unbiased Scene Graph Generation. In *CVPR*.
- Dong, Y.; Chawla, N. V.; and Swami, A. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining(KDD)*, 135–144.
- Gao, L.; Wang, B.; and Wang, W. 2018. Image captioning with scene-graph based semantic concepts. *ICMLC*.
- Ghosh, S.; Burachas, G.; Ray, A.; and Ziskind, A. 2019. Generating natural language explanations for visual question answering using scene graphs and visual attention. *arXiv preprint arXiv:1902.05715*.
- Hu, Z.; Dong, Y.; Wang, K.; and Sun, Y. 2020. Heterogeneous graph transformer. In *WWW*.
- Jing, B.; Park, C.; and Tong, H. 2021. HDMI: High-Order Deep Multiplex Infomax. In *Proceedings of the Web Conference 2021*, WWW '21, 2414–2424. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383127.
- Knyazev, B.; de Vries, H.; Cangea, C.; Taylor, G. W.; Courville, A.; and Belilovsky, E. 2020. Graph density-aware losses for novel compositions in scene graph generation. *BMVC*.
- 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. *ICCV*.
- Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Kolesnikov, A.; et al. 2020a. The open images dataset v4. *International Journal of Computer Vision*, 128(7): 1956–1981.
- Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Kolesnikov, A.; et al. 2020b. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *ICCV*.
- Li, L.; Chen, L.; Huang, Y.; Zhang, Z.; Zhang, S.; and Xiao, J. 2022a. The Devil is in the Labels: Noisy Label Correction for Robust Scene Graph Generation. In *CVPR*, 18869–18878.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Thirty-Second AAAI conference on artificial intelligence*.
- Li, R.; Zhang, S.; Wan, B.; and He, X. 2021. Bipartite graph network with adaptive message passing for unbiased scene graph generation. *CVPR*.
- Li, W.; Zhang, H.; Bai, Q.; Zhao, G.; Jiang, N.; and Yuan, X. 2022b. PPDL: Predicate Probability Distribution Based Loss for Unbiased Scene Graph Generation. In *CVPR*.
- Li, Y.; Ouyang, W.; Zhou, B.; Wang, K.; and Wang, X. 2017. Scene graph generation from objects, phrases and region captions. *ICCV*, 1261–1270.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. *ICCV*.
- Lin, X.; Ding, C.; Zeng, J.; and Tao, D. 2020. Gps-net: Graph property sensing network for scene graph generation. *CVPR*, 3746–3753.
- Park, C.; Kim, D.; Han, J.; and Yu, H. 2020. Unsupervised Attributed Multiplex Network Embedding. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 5371–5378. AAAI Press.
- Park, C.; Kim, D.; Zhu, Q.; Han, J.; and Yu, H. 2019. Task-guided pair embedding in heterogeneous network. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 489–498.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*.
- Qi, M.; Li, W.; Yang, Z.; Wang, Y.; and Luo, J. 2019. Attentive relational networks for mapping images to scene graphs. *CVPR*, 3957–3966.
- Ramnath, S.; Saha, A.; Chakrabarti, S.; and Khapra, M. M. 2019. Scene Graph based Image Retrieval—A case study on the CLEVR Dataset. *ICCV Workshops*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Berg, R. v. d.; Titov, I.; and Welling, M. 2018. Modeling relational data with graph convolutional networks. *European semantic web conference(ESWC)*, 593–607.
- Schroeder, B.; and Tripathi, S. 2020. Structured query-based image retrieval using scene graphs. *CVPR Workshops*, 178–179.
- Sun, Y.; Han, J.; Yan, X.; Yu, P. S.; and Wu, T. 2011. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. *Proceedings of the VLDB Endowment*, 4(11): 992–1003.
- Sur, C. 2019. Tpsgr: Neural-symbolic tensor product scene-graph-triplet representation for image captioning. *arXiv preprint arXiv:1911.10115*.
- Tang, K.; Niu, Y.; Huang, J.; Shi, J.; and Zhang, H. 2020. Unbiased scene graph generation from biased training. *CVPR*, 3716–3725.
- Tang, K.; Zhang, H.; Wu, B.; Luo, W.; and Liu, W. 2019. Learning to compose dynamic tree structures for visual contexts. *CVPR*.
- Wang, X.; Ji, H.; Shi, C.; Wang, B.; Ye, Y.; Cui, P.; and Yu, P. S. 2019. Heterogeneous graph attention network. *The world wide web conference(WWW)*, 2022–2032.

- Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. *Proceedings of the IEEE conference on computer vision and pattern recognition(CVPR)*, 1492–1500.
- Xu, D.; Zhu, Y.; Choy, C. B.; and Fei-Fei, L. 2017. Scene graph generation by iterative message passing. *CVPR*.
- Yan, S.; Shen, C.; Jin, Z.; Huang, J.; Jiang, R.; Chen, Y.; and Hua, X.-S. 2020. PcpI: Predicate-correlation perception learning for unbiased scene graph generation. In *Proceedings of the 28th ACM International Conference on Multimedia*, 265–273.
- Yang, J.; Lu, J.; Lee, S.; Batra, D.; and Parikh, D. 2018. Graph r-cnn for scene graph generation. *ECCV*, 128(7): 670–685.
- Yang, X.; Tang, K.; Zhang, H.; and Cai, J. 2019. Auto-encoding scene graphs for image captioning. *CVPR*, 10685–10694.
- Yun, S.; Kim, K.; Yoon, K.; and Park, C. 2022. LTE4G: Long-Tail Experts for Graph Neural Networks. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2434–2443.
- Zellers, R.; Yatskar, M.; Thomson, S.; and Choi, Y. 2018. Neural motifs: Scene graph parsing with global context. *CVPR*, 5831–5840.
- Zhang, C.; Chao, W.-L.; and Xuan, D. 2019. An empirical study on leveraging scene graphs for visual question answering. *BMVC*.
- Zhang, C.; Song, D.; Huang, C.; Swami, A.; and Chawla, N. V. 2019a. Heterogeneous graph neural network. In *KDD*.
- Zhang, J.; Shih, K. J.; Elgammal, A.; Tao, A.; and Catanzaro, B. 2019b. Graphical contrastive losses for scene graph parsing. *CVPR*, 11535–11543.