

Bipartite Graph Network with Adaptive Message Passing for Unbiased Scene Graph Generation

Rongjie Li^{1,3,4} Songyang Zhang^{1,3,4} Bo Wan^{1,5,*} Xuming He^{1,2}

¹School of Information Science and Technology, ShanghaiTech University

²Shanghai Engineering Research Center of Intelligent Vision and Imaging

³Shanghai Institute of Microsystem and Information Technology, Chinese Academy of Sciences

⁴University of Chinese Academy of Sciences

⁵Department of Electrical Engineering(ESAT), KU Leuven

{lirj2, zhangsy, wanbo, hexm}@shanghaitech.edu.cn

Abstract

Scene graph generation is an important visual understanding task with a broad range of vision applications. Despite recent tremendous progress, it remains challenging due to the intrinsic long-tailed class distribution and large intra-class variation. To address these issues, we introduce a novel confidence-aware bipartite graph neural network with adaptive message propagation mechanism for unbiased scene graph generation. In addition, we propose an efficient bi-level data resampling strategy to alleviate the imbalanced data distribution problem in training our graph network. Our approach achieves superior or competitive performance over previous methods on several challenging datasets, including Visual Genome, Open Images V4/V6, demonstrating its effectiveness and generality.

1. Introduction

Scene graph generation, which aims to detect visual objects and their relationships (or *triplets*: $\langle \text{subject}, \text{predicate}, \text{object} \rangle$) in an image, is a fundamental visual understanding task. Such a compact structural scene representation has potential applications in many vision tasks such as visual question answering [41, 37, 11], image captioning [54] and image retrieval [14]. Tremendous progress has been made recently in scene graph generation [20, 50, 23, 59, 53, 22, 40, 3, 8, 60, 39, 27, 45], thanks to learned visual representations and advances in object detection. However, this task remains particularly challenging due to large variations in visual relationships, extremely

*This work was done when Bo Wan was a master student in ShanghaiTech University. This work was supported by Shanghai NSF Grant (No. 18ZR1425100). Code is available: <https://github.com/Scarecrow0/BGNN-SGG>

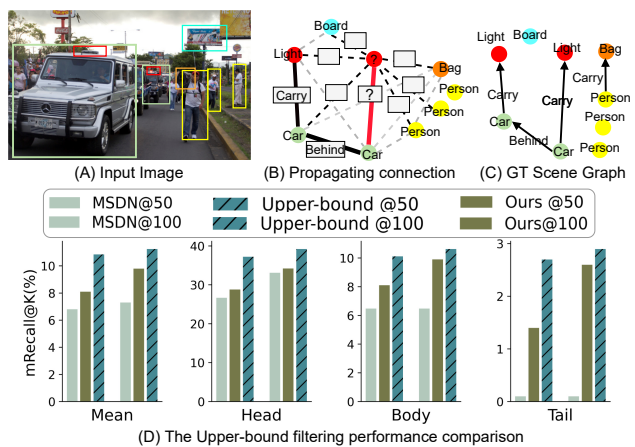


Figure 1: **The illustration of biased scene graph generation and empirical study on Visual Genome.** As shown in (D), the baseline (MSDN [23]) performance is dominated by the head categories due to the imbalanced data. We estimate an upper-bound performance by ignoring negative predicate-entity connections during message propagation, as shown in (B). Its performance (shown in (D)) indicates a large room for improvement in context modeling.

imbalanced object and relation distribution and lack of sufficient annotations for many categories.

One primary challenge, which causes *biased relationship prediction*, is the intrinsic long-tail data distribution. A scene graph model has to simultaneously cope with imbalanced annotations among the head and medium-sized categories, and few-shot learning in the tail categories. A naively learned model will be largely dominated by those few head categories with much degraded performance for many tail categories (as shown in Fig. 1-D). Early work [1, 4] on re-balancing data distribution focus on data

re-sampling or loss re-weighting. However, it is non-trivial to directly apply the image-level re-balancing strategies for such instance-level tasks. Recent efforts try to introduce the re-balancing ideas into object detection [9, 38] and scene graph generation [39], but it remains difficult to achieve a satisfactory trade-off between head and tail categories.

Moreover, those non-head predicate categories typically involve complex semantic meaning and large intra-class variations (*e.g.* play, look) in images, which exacerbates the problems in their representation learning and classification. Many previous works [59, 50, 23, 40, 53, 27] attempt to address this problem by developing context modeling mechanisms, but often suffer from noisy information propagation due to their use of fully connected graphs. More recent efforts [33, 53, 40, 45] aim to improve context modeling by designing a sparse structure, which also limits the model flexibility. To illustrate the impact of noises in graph, we further conduct an empirical analysis, as shown in Fig. 1, which indicates that *a baseline model can achieve notable performance improvement by removing the noisy subject-object associations.*

Based on these findings, we propose a novel confidence-aware graph representation and its learning strategy for unbiased scene graph generation. To this end, we first develop a bipartite graph neural network (BGNN) with the adaptive message propagation for effective context modeling. Specifically, our method takes the hypothesize-and-classify strategy, which first generates a set of visual entity and predicate proposals from a proposal generation network. Then we compute a context-aware representation for those proposals by passing them through a multi-stage BGNN. Our graph network adopts directed edges to model different information flow between entity and relationship proposals as a bipartite graph, and an adaptive message propagation strategy based on relation confidence estimation to reduce the noise in the context modeling. Finally, we use the refined entity and predicate representations to predict their categories with linear classifiers.

To train our multi-stage BGNN for unbiased prediction, we also design a bi-level data resampling strategy to alleviate the imbalanced data distribution problem. Our method combines the image-level over-sampling and instance-level under-sampling ideas [13, 9] for the structured prediction task. Equipped with this strategy, we can achieve a better trade-off between the head and tail categories and learn our bipartite graph neural network more effectively.

We extensively validate our methods on three scene graph generation datasets, including Visual Genome, Open Images V4, and Open Images V6. The empirical results and ablation studies show our method consistently achieves competitive or state-of-the-art performance on all benchmarks. The main contributions of our works are three-folds.

- We introduce a bipartite graph neural network with

adaptive message propagation to alleviate the error propagation and achieve effective context modeling.

- We propose a bi-level data resampling to achieve a better trade-off between head and tail categories for scene graph generation.
- Our method achieves competitive or state-of-the-art performance on various scene graph benchmarks.

2. Related Works

Scene Graph Generation. Traditional methods in scene graph generation typically utilize graph-based context-modeling strategies to learn discriminative representation for node and edge prediction. Most of them either focus on the graph structure design or leveraging scene context via various message propagation mechanisms.

Several types of graph structure have been proposed for context modeling in literature. A popular idea is to model the context based on a sequential model (*e.g.*, LSTM) [59] or a fully-connected graph [50, 5, 23, 55, 48, 44, 27]. In addition, recent works [40, 43, 54, 33] explore sparse graph structures, which are either associated with the downstream tasks (*e.g.* VQA) or built by trimming the relationship proposals according to the category or geometry information of subject-object pairs. However, these works often rely on their specific designs based on the downstream tasks, which limits the flexibility of their representations.

Another direction aims to incorporate context information into existing deep ConvNet models by exploring different message propagation mechanisms. A common strategy is to perform message passing between the entities proposals [59, 40, 45, 48, 33, 44, 27, 3], while the other aggregates the contextual information between the entities and predicates [50, 23, 5, 22, 55, 53, 44, 56], which also produces effective scene graph representations.

Our work considers both message passing and inferring network connectivity in a single framework. In particular, we develop a generic Bipartite Graph Neural Network (BGNN) to effectively model the context of the entity and predicate proposals, and an adaptive message propagation to compute a more flexible representation. The previous SGG models [23, 50, 22, 53] can be considered as special cases of the BGNN.

Long-tail Visual Recognition Previous works in visual recognition typically utilize re-balancing strategies to alleviate biased prediction caused by long-tail distributions. These re-balancing strategies include dataset resampling to achieve balanced class prior [2, 6, 36, 31], and loss re-weighting based on instance frequency or hard-example mining [1, 4, 17, 38, 26, 28, 25, 30]. Recently, [13, 9] propose an instance-level re-sampling strategy for the tasks of object detection and instance segmentation. Other approaches also explore knowledge transfer learning from

head categories for long-tail classification [29, 7] or develop the two-stage learning scheme [63, 16]. However, it is non-trivial to apply a naive re-balancing strategy for scene graph generation. We propose a *bi-level data sampling strategy* by combing the *image-level over-sampling* [9] and *instance-level under-sampling* [13] to achieve a better trade-off between the head and tail categories.

For the task of scene graph generation, several strategies have proposed to tackle the intrinsic long-tail problem. Some researchers propose novel loss designs by leveraging the semantic constraints of scene graph [19, 27, 52, 43]. Others develop new graph structure encoding the context [40, 3, 27] or introduce external commonsense and linguistic knowledge [57, 58, 32] for better representation learning. Recently, Tang *et al.* [39] proposes an unbiased inference method by formulating the recognition process as a causal model. In this work, we aim to improve the context modeling for tail-categories by design an novel graph network and message propagation mechanism.

Graph Neural Network Our work is also related to learning deep networks on graph-structure data. The Graph Neural Network (GNN) is first proposed by [35], which is a powerful method for handling the non-Euclidean data. Previous works have explored different graph structures (e.g. directed graph [15], heterogeneous graph [62, 47]) and aggregation mechanism (e.g. convolutional [18, 10], attention [24, 46, 61, 42]) for various tasks. Several recent efforts attempt to improve the quality of message propagation in GNNs. Hou *et al.* [12] proposes a context-surrounding GNN framework to measure the quality of neighborhood aggregation. Xu *et al.* [51] introduces a dynamical subgraph construction by applying a graphical attention mechanism conditioned on input queries. In this work, we introduce a novel bipartite graph network to learn a robust context-aware feature representation for the predicates.

3. Our Approach

We aim to tackle scene graph generation, which parses an input image into a structural graph representation of object entities and their visual relationship in the scene. In particular, we focus on addressing the challenge of biased scene graph prediction, mainly caused by the intrinsic long-tail distribution of visual relationship in a typical training dataset and the large intra-class variation of predicate categories. To this end, we introduce an adaptive scene graph generation strategy, which simultaneously learns a robust entity/predicate representation and a calibrated classifier for better balanced performance.

In this section, we first present the problem setting of scene graph generation and an overview of our method in Sec. 3.1. We then introduce the details of our predicate representation in Sec. 3.2, followed by our learning strategy on the predicate classifier and representation in Sec. 3.3.

3.1. Problem Setting and Overview

Problem Setting Given an image \mathbf{I} , the task of scene graph generation (SGG) aims to parse the input \mathbf{I} into a scene graph $\mathcal{G}_{scene} = \{\mathcal{V}_o, \mathcal{E}_r\}$, where \mathcal{V}_o is the node set encoding object entities and \mathcal{E}_r is the edge set that represents predicate between an ordered pair of entities. Typically, each node $v_i \in \mathcal{V}_o$ has a category label from a set of entity classes \mathcal{C}_e and a corresponding image location represented by a bounding box, while each edge $e_{i \rightarrow j} \in \mathcal{E}_r$ between a pair of nodes v_i and v_j is associated with a predicate label from a set of predicate classes \mathcal{C}_p in this task.

Method Overview In this work, we adopt a hypothesize-and-classify strategy for the unbiased scene graph generation. Our approach first generates a set of entity and predicate proposals and then computes their context-aware representations, followed by predicting their categories.

Concretely, we introduce a bipartite graph network that explicitly models the interactions between entities and their predicates in order to cope with unreliable contextual information from the noisy proposals. Based on the graph network, we develop an adaptive message passing scheme capable of actively controlling the information flows to reduce the noise impact in the graph and generate a robust context-aware representation for each relationship proposal.

Taking this representation, we then learn predicate and entity classifiers to predict the categories of predicate and entity within relationship proposals. To alleviate the bias effect of the imbalanced data, we also design an efficient bi-level data resampling strategy for the model training, which enables us to achieve a better trade-off between the head and tail categories. An overview of our method is illustrated in Fig. 2, and we will start from a detailed description of our model architecture below.

3.2. Model Architecture

Our scene graph generation model is a modular deep network consisting of three main submodules: 1) a *proposal generation network* to generate entity and relationship proposals and compute their initial representation (Sec. 3.2.1); 2) a *bipartite graph neural network* to encode the scene context with adaptive message propagation and multi-stage iterative refinement (Sec. 3.2.2); and 3) a *scene graph predictor* to decode the scene graph from the context-aware representations of relationship proposals (Sec. 3.2.3).

3.2.1 Proposal Generation Network

Following [59, 60], we utilize an object detector network (e.g., Faster R-CNN [34]) to generate a set of entity and relationship proposals. The entity proposals are taken directly from the detection output with their categories and classification scores, while the relationship proposals are generated by forming ordered pairs of all the entity proposals.

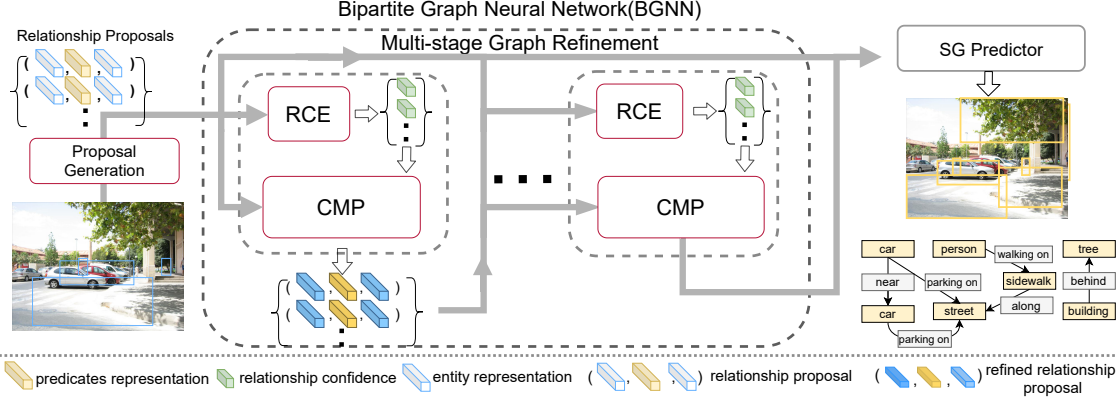


Figure 2: **Illustration of overall pipeline of our BGNN model.** RCE denotes the relationship confidence estimation module. CMP denotes the confidence-aware message propagation model. SG Predictor is the scene graph predictor for the final prediction.

Given the relationship proposals, we then compute an initial representation for both entities and predicates. Specifically, for the i -th entity proposal, we denote its convolution feature as \mathbf{v}_i , its bounding box as b_i and its detected class as c_i . The entity representation \mathbf{e}_i uses a fully-connected network f_e to integrate its visual, geometric and semantic features as,

$$\mathbf{e}_i = f_e(\mathbf{v}_i \oplus \mathbf{g}_i \oplus \mathbf{w}_i) \quad (1)$$

where \mathbf{g}_i is a geometric feature based on its bounding box b_i , \mathbf{w}_i is a semantic feature based on a word embedding of its class c_i , and \oplus is the concatenation operation.

For the relationship proposal from entity i to j , we combines the entity representations $\mathbf{e}_i \oplus \mathbf{e}_j$ with the convolutional feature of their union region (denoted as $\mathbf{u}_{i,j}^p$). Formally, we compute the predicate representation $\mathbf{r}_{i \rightarrow j}$ as

$$\mathbf{r}_{i \rightarrow j} = f_u(\mathbf{u}_{i,j}^p) + f_p(\mathbf{e}_i \oplus \mathbf{e}_j) \quad (2)$$

where f_u and f_p are two fully-connected networks.

3.2.2 Bipartite Graph Neural Network

Given the relationship proposals, we build a graph structure to capture the dependency between entities and predicates. To this end, we introduce a bipartite graph \mathcal{G}_b with directed edges, which enables us to model the different information flow directions between entity and predicate representations. Specifically, the graph consists of two groups of nodes $\mathcal{V}_e, \mathcal{V}_p$, which correspond to entity representations and predicate representations respectively. Those two groups of nodes are connected by two sets of directed edges $\mathcal{E}_{e \rightarrow p}$ and $\mathcal{E}_{p \rightarrow e}$ representing information flows from the entities to predicates and vice versa. Hence the bipartite graph has a form as $\mathcal{G}_b = \{\mathcal{V}_e, \mathcal{V}_p, \mathcal{E}_{e \rightarrow p}, \mathcal{E}_{p \rightarrow e}\}$.

To effectively model the context of the entity and predicate proposals, we develop a Bipartite Graph Neural Network (BGNN) on the graph \mathcal{G}_b . Our BGNN conducts a multi-stage message propagation and each stage consists of 1) a *relationship confidence estimation* module to provide a confidence estimate on relationship; 2) a *confidence-aware message propagation* to incorporate scene context and semantic cues into the entity/predicate proposals. The overview of our BGNN is illustrated in Fig. 2. We will focus on a single stage of our network in the rest of this section.

Relationship Confidence Estimation (RCE) Module In order to reduce the noise in context modeling, we introduce a relationship confidence estimation (RCE) module. It predicts a confidence score for each relationship proposal to control the information flow in the message propagation.

Concretely, for a predicate node from entity i to j , the RCE module takes as input the predicate proposal features $\mathbf{r}_{i \rightarrow j}$ and its associated entities' class scores, and predicts a confidence score for each predicate class as below,

$$\mathbf{s}_{i \rightarrow j}^m = g_x(\mathbf{r}_{i \rightarrow j} \oplus \mathbf{p}_i \oplus \mathbf{p}_j) \in \mathbb{R}^{|\mathcal{C}_p|} \quad (3)$$

where $\mathbf{p}_i, \mathbf{p}_j \in \mathbb{R}^{|\mathcal{C}_e|}$ are the class probabilities for entity \mathbf{e}_i and \mathbf{e}_j from the detection, and g_x is a multilayer fully-connected network. We then fuse those confidence scores into a global confidence score for the predicate node as

$$\mathbf{s}_{i \rightarrow j}^b = \sigma(\mathbf{w}_b^T \mathbf{s}_{i \rightarrow j}^m), \quad \mathbf{w}_b \in \mathbb{R}^{|\mathcal{C}_p|} \quad (4)$$

where σ is the sigmoid activation function and \mathbf{w}_b are the parameters for the fusion.

Confidence-aware Message Propagation We now introduce our adaptive message propagation for capturing the scene context. Specifically, we design two types of message passing update, including an *entity-to-predicate* message and a *predict-to-entity* message according to the edge directions, as illustrated in Fig. 3. Below we consider an iteration from l to $l + 1$ in the message passing.




Figure 3: **Two kinds of message propagation with confidence gating of bipartite graph neural network.** The dashed arrows mean information flow is blocked when the source is uncertain. The different alpha value stands for different aggregation weights for rest connections.

1) **Entity-to-predicate Message Propagation:** We update the representation of a predicate node $\mathbf{r}_{i \rightarrow j}$ by fusing its neighboring entity nodes:

$$\mathbf{r}_{i \rightarrow j}^{(l+1)} = \mathbf{r}_{i \rightarrow j}^{(l)} + \phi(d_s \mathbf{W}_r^\top \mathbf{e}_i^{(l)} + d_o \mathbf{W}_r^\top \mathbf{e}_j^{(l)}) \quad (5)$$

$$d_s = \sigma(\mathbf{w}_s^\top [\mathbf{r}_{i \rightarrow j}^{(l)} \oplus \mathbf{e}_i^{(l)}]), \quad d_o = \sigma(\mathbf{w}_o^\top [\mathbf{r}_{i \rightarrow j}^{(l)} \oplus \mathbf{e}_j^{(l)}]) \quad (6)$$

where \mathbf{W}_r is a linear transformation, ϕ is an activation function (e.g. ReLU). and d_s, d_o are learnable affinity functions of entity and predicate, where \mathbf{w}_s and \mathbf{w}_o are their parameters.

2) **Predicate-to-entity Message Propagation:** As the predicate node set \mathcal{V}_p typically includes a considerable amount of false positive predicate proposals, we develop a *confidence-aware* adaptive message propagation for entity nodes update to alleviate such noise effect. Specifically, we first introduce a confidence gating function to control the information flow from an entity's neighbor $\mathbf{r}_{i \rightarrow j}$ as:

$$\gamma_{i \rightarrow j} = \mathcal{T}(s_{i \rightarrow j}^b), \quad \mathcal{T}(x) = \begin{cases} 0 & x \leq \beta \\ \alpha x - \alpha\beta & \beta < x < 1/\alpha + \beta \\ 1 & x \geq 1/\alpha + \beta \end{cases} \quad (7)$$

where α and β is learnable threshold parameters. The gating function $\mathcal{T}(x)$ is designed for achieving a hard control for the predicate proposals with high or low scores (confidently positive or negative), and a soft control for the predicates with intermediate scores.

For each entity node \mathbf{e}_i , we divide its neighboring predicates into two sets: $\mathcal{B}_s(i)$ for \mathbf{e}_i as the *subject* and $\mathcal{B}_o(i)$ for \mathbf{e}_i as the *object*. We update the entity representation \mathbf{e}_i by

aggregating its neighbors' messages:

$$\mathbf{e}_i^{(l+1)} = \mathbf{e}_i^{(l)} + \phi \left(\frac{1}{|\mathcal{B}_s(i)|} \sum_{k \in \mathcal{B}_s(i)} \gamma_k d_s \mathbf{W}_e^\top \mathbf{r}_k^{(l)} \right) \quad (8)$$

$$+ \frac{1}{|\mathcal{B}_o(i)|} \sum_{k \in \mathcal{B}_o(i)} \gamma_k d_o \mathbf{W}_e^\top \mathbf{r}_k^{(l)} \right) \quad (9)$$

where \mathbf{W}_e is the parameter of a linear transformation.

In each stage, we typically perform N_i iterations of the above two message propagations to capture context in a sufficiently large scope.

3.2.3 Scene Graph Prediction

To generate the scene graph of the given image, we introduce two linear classifiers to predict the class of the entities and predicates based on their refined representations. Concretely, for each relationship proposal, our classifier integrates the final representation of predicates proposal from our BGNN, denoted as $\hat{\mathbf{r}}_{i \rightarrow j}$, and a class frequency prior [59], $\hat{\mathbf{p}}_{\mathbf{r}_{i \rightarrow j}}$, for classification:

$$\mathbf{p}_{\mathbf{r}_{i \rightarrow j}} = \text{softmax}(\mathbf{W}_{rel}^\top \hat{\mathbf{r}}_{i \rightarrow j} + \log(\hat{\mathbf{p}}_{\mathbf{r}_{i \rightarrow j}})) \in \mathbb{R}^{C_p} \quad (10)$$

For each entity, we introduce a learnable weight to fuse the initial visual features \mathbf{v}_i and enhanced features $\hat{\mathbf{e}}_i$ output by our BGNN. The final entity classification is computed as:

$$\mathbf{p}_{\mathbf{e}_i} = \text{softmax}(\mathbf{W}_{ent}^\top (\rho \hat{\mathbf{e}}_i + (1 - \rho) \mathbf{v}_i)) \in \mathbb{R}^{C_e} \quad (11)$$

where ρ is a weight in $[0, 1]$, and \mathbf{W}_{rel} and \mathbf{W}_{ent} are the parameters of two classifiers.

3.3. Learning with Bi-level Data Sampling

We now present our learning strategy for unbiased scene graph generation. We will first develop a bi-level data sampling strategy to balance the data distribution of entities and predicates, and then describe a multitask loss for learning the adaptive BGNN.

Bi-level Data Resampling Unlike in other vision tasks, the scene graph annotations have varying structures, which makes it non-trivial to adopt either the instance-level replay strategy [13] or images-level resampling method LVIS[9]. To tackle the intrinsic long-tail data distribution of entity and relation, we design a two-level data sampling strategy that integrates the above two ideas on rebalancing. Specifically, our data sampling strategy consists of two steps:

1) **Image-level over-sampling:** We adopt the repeat factor sampling in [9] to sample images first. We start from a class-specific repeat number, $r^c = \max(1, \sqrt{t/f^c})$, where c is the category, f^c is its frequency on the entire dataset

B	Models	PredCls		SGCls		SGGen	
		mR@50 / 100	R@50 / 100	mR@50 / 100	R@50 / 100	mR@50 / 100	R@50 / 100
VGG16	Motifs[59, 40]	14.0 / 15.3	65.2 / 67.1	7.7 / 8.2	35.8 / 36.5	5.7 / 6.6	27.2 / 30.3
	FREQ [59, 40]	13.0 / 16.0	60.6 / 62.2	7.2 / 8.5	32.3 / 32.9	6.1 / 7.1	26.2 / 30.1
	G-RCNN [53]	- / -	54.2 / 59.1	- / -	31.6 / 29.6	- / -	11.4 / 13.7
	VCTree[40]	17.9 / 19.4	66.4 / 68.1	10.1 / 10.8	38.1 / 38.8	6.9 / 8.0	27.9 / 31.3
	RelDN[60]	- / -	68.4 / 68.4	- / -	36.8 / 36.8	- / -	28.3 / 32.7
	KERN [3]	17.7 / 19.2	67.6 / 65.8	9.4 / 10.0	36.7 / 37.4	6.4 / 7.3	29.8 / 27.1
	GPS-Net[27]	- / 22.8	66.9 / 68.8	- / 12.6	39.2 / 40.1	- / 9.8	28.4 / 31.7
	PCPL[52]	35.2 / 37.8	50.8 / 52.6	18.6 / 19.6	27.6 / 28.4	9.5 / 11.7	14.6 / 18.6
X-101-FPN	RelDN [†]	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[39]	14.6 / 15.8	66.0 / 67.9	8.0 / 8.5	39.1 / 39.9	5.5 / 6.8	32.1 / 36.9
	Motifs*[39]	18.5 / 20.0	64.6 / 66.7	11.1 / 11.8	37.9 / 38.8	8.2 / 9.7	30.5 / 35.4
	VCTree[39]	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 [†]	16.4 / 17.2	65.4 / 67.2	9.0 / 9.5	38.5 / 37.0	5.8 / 6.6	29.7 / 32.8
	MSDN [†] [23]	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[39]	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 [†]	15.2 / 16.6	65.2 / 67.1	8.5 / 9.1	39.2 / 37.8	6.7 / 8.6	31.1 / 35.9
	GPS-Net [†] *	19.2 / 21.4	64.4 / 66.7	11.7 / 12.5	37.5 / 38.6	7.4 / 9.5	27.8 / 32.1
BGNN	30.4 / 32.9	59.2 / 61.3	14.3 / 16.5	37.4 / 38.5	10.7 / 12.6	31.0 / 35.8	

Table 1: **The SGG performance of three tasks with graph constraints setting.** [†] denote results reproduced with the authors’ code. * denotes the resampling [9] is applied for this model.

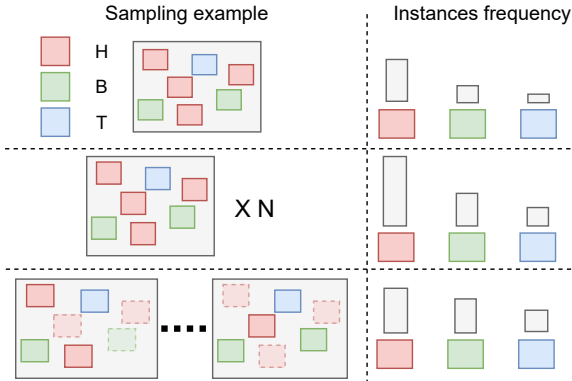


Figure 4: **Illustration of bi-level data sampling for one image.** The top row is the instance frequency of head(H), body(B), and tail(T) categories in the image. The middle row shows image-level oversampling with repeat factor N . The bottom row shows the instance-level under-sampling for instances of different categories.

and t is a hyper-parameter that controls when oversampling starts. For i -th image, we set $r_i = \max_{c \in i} r^c$, where $\{c \in i\}$ are the categories labeled in i .

2) *Instance-level under-sampling*: Given the sampled images, we further design an instance-level sampling strategy for predicates. Concretely, we compute a drop-out probability for instances of different predicate classes in each image. The drop-out rate d_i^c for instances in i -th image, with category label c is calculated by $d_i^c = \max((r_i -$

$r^c)/r_i * \gamma_d, 1.0)$, and γ_d is the hyper-parameter for adjusting the drop-out rate. With this strategy, our two-level data resampling can achieve an effective trade-off between the head and tail categories.

Training Losses To train our BGNN model, we design a multitask loss that consists of three components, including \mathcal{L}_{rce} for relation confidence estimation module (RCE), \mathcal{L}_p for predicate proposal classification and \mathcal{L}_e for entity proposal classification. Formally,

$$\mathcal{L}_{total} = \mathcal{L}_p + \lambda_{rce}\mathcal{L}_{rce} + \lambda_e\mathcal{L}_e \quad (12)$$

where λ_{rce}, λ_e are weight parameters for calibrating the supervision from each sub-task.

Here $\mathcal{L}_p, \mathcal{L}_e$ are the standard cross entropy loss for multi-class classification (foreground categories plus background). The loss of RCE \mathcal{L}_{rce} is composed by two terms: $\mathcal{L}_{rce} = \mathcal{L}_m + \lambda \cdot \mathcal{L}_b$, where λ is a weight parameter, and \mathcal{L}_m and \mathcal{L}_b are losses for the class-specific and overall relation confidence estimation s^m, s^b respectively. Both predictions have explicit supervision as in the training of relationship predictor in the graph refinement stage. We adopt the focal loss [26] to alleviate positive-negative imbalance in the relationship confidence estimation.

4. Experiments

In this section, we conduct a series of comprehensive experiments to validate the effectiveness of our method. Below we first present our experimental analysis and ablative




Figure 5: The per-group results on VG dataset (SGGen task).

study on Visual Genome [20] dataset (in Sec. 4.1), then our results on the Open Images V4/V6 [21] (In Sec. 4.2). In each dataset, we first introduce the implementation details of our method and then report comparisons of quantitative results in detail.

4.1. Visual Genome

4.1.1 Experiments Configurations

Dataset Details For Visual Genome [20] dataset, we take the same split protocol as [50, 59]. The most frequent 150 object categories and 50 predicates are adopted for evaluation. Following the similar protocol in [29], we divided the categories into three disjoint groups according to the instance number in training split: *head* (more than 10k), *body* (0.5k ~ 10k), *tail* (less than 0.5k), more details are reported in Suppl.

Evaluation Protocol We evaluate the model on three sub-tasks as [50, 59]: 1) predicate classification (PredCls), 2) scene graph classification (SGCls), 3) scene graph generation (SGGen, also denote as SGDNet). Following the previous works [3, 40, 39, 27] that concentrate on the long-tail distribution of entity and predicate categories on Visual Genome, we takes **recall@K(R@K)** and **mean recall@K (mR@K)** as evaluation metrics, and also report the mR@K on each long-tail category groups (*head*, *body* and *tail*).

Implementation Details Similar to Tang [39], we adopt ResNeXt-101-FPN [49] as backbone network and FasterRCNN [34] as object detector, whose model parameters are kept frozen in model training. In our resampling strategy, we set the repeat factor $t=0.07$, instances drop rate $\gamma_d=0.7$, and weight of fusion the entities features $\rho=-5$. The α, β are initialized as 2.2 and 0.025 respectively.

4.1.2 Comparisons with State-of-the-Art Methods

As shown in Tab. 1, BGNN achieves the state of the art in all three sub-tasks (PredCls, SGCls, SGGen) evaluated by mR@50/100 on X-101-FPN backbone, and outperforms Unbiased [39] with a significant margin of **5.0** and **4.2** on PredCls. Besides, BGNN shows a comparable performance with previous SOTA in SGCls and SGGen sub-tasks

Module			SGGen				
B	C	M	mR@100	R@100	Head	Body	Tail
\times	\times	\times	9.7	34.1	32.1	9.3	3.0
\checkmark	\times	\times	10.5	34.8	32.4	10.7	4.5
\checkmark	\checkmark	\times	11.7	35.3	33.6	11.4	5.2
\checkmark	\checkmark	\checkmark	12.6	35.8	34.0	12.9	6.0

Table 2: Ablation study of model structure. **B**: bipartite graph neural network with plain message passing; **C**: confidence-aware message propagation mechanism; **M**: multi-stage refinement

Resample	SGGen				
	mR@100	R@100	Head	Body	Tail
None	9.7	36.1	34.2	9.9	2.6
RFS[9]	10.7	34.6	33.2	9.7	3.5
BLS	12.6	35.8	34.0	12.9	6.0

Table 3: The ablation for the resampling strategy. RFS: repeat factor sampling; BLS: bi-level data sampling.

on R@50/100, which demonstrate the effectiveness of our methods.

On the VGG16 backbone, we achieves comparable result with the state-of-the-art, PCPL [52] on the SGCls and PredCls sub-tasks by mR@50/100. However, BGNN outperforms the PCPL on SGGen sub-tasks by a large margin, which shows that our proposed BGNN is capable of handling challenging SGGen task with more noisy proposals.

Moreover, as shown in Fig. 5, we compute the mean recall on each long-tail category groups in SGGen sub-task and find BGNN significantly outperforms the prior works [27, 39] on *tail* group, as a result it achieves the highest mean recall over all categories. We provide more visualization of our results and comparisons in the Suppl.

4.1.3 Ablation Study

Model Components We first evaluate the importance of each component in our BGNN. As shown in Tab. 2, we incrementally add one component to the plain baseline to validate their effectiveness. We observe the plain message propagation (without adaptive confidence-aware mechanism) on bipartite graph neural network improves the baseline with **8%** and achieves **10.5** on mR@100. Besides, the confidence-aware message propagation mechanism promotes the performance to **11.7**, and multiple stages refinement further improves the final results to **12.6**.

Re-sampling Strategies We compare the widely-used repeat factor sampling [9] with BLS for validating the proposed bi-level data sampling strategy (BLS). As shown in

		SGGen				
N_t	N_i	mR@100	R@100	Head	Body	Tail
1	1	10.0	35.3	33.5	9.3	4.0
1	2	10.5	35.5	34.0	9.4	4.2
1	3	10.8	35.3	33.8	10.6	4.6
1	3	10.8	35.3	33.8	10.6	4.6
2	3	11.1	35.6	34.0	11.3	5.3
3	3	12.6	35.8	34.0	12.9	6.0
4	3	12.5	35.2	34.4	12.2	5.7

Table 4: **The ablation for the different graph iteration parameters.** The ablation of different iteration number for iterative refinement models in our methods.

D	Models	mR@50	R@50	wmAP		score _{wtd}
				rel	phr	
V4	RelDN[27]	-	74.94	35.54	38.52	44.61
	GPS-Net[27]	-	77.27	38.78	40.15	47.03
	RelDN [†]	70.40	75.66	36.13	39.91	45.21
	GPS-Net [†]	69.50	74.65	35.02	39.40	44.70
	BGNN	72.11	75.46	37.76	41.70	46.87
V6	RelDN [†]	33.98	73.08	32.16	33.39	40.84
	RelDN ^{†*}	37.20	75.34	33.21	34.31	41.97
	VCTree [†]	33.91	74.08	34.16	33.11	40.21
	G-RCNN [†]	34.04	74.51	33.15	34.21	41.84
	Motifs [†]	32.68	71.63	29.91	31.59	38.93
	Unbiased [†]	35.47	69.30	30.74	32.80	39.27
	GPS-Net [†]	35.26	74.81	32.85	33.98	41.69
	GPS-Net ^{†*}	38.93	74.74	32.77	33.87	41.60
	BGNN	40.45	74.98	33.51	34.15	42.06

Table 5: **The Performance table of Open Images Dataset.** * denotes the resampling [9] is applied for this model. † denote results reproduced with the authors’ code.

in Tab. 3, we find our proposed BLS outperforms the baseline and RFS with a large margin, especially on *body* and *tail* categories. Besides, unlike other methods, BLS maintains the performance on head categories. This indicates our sampling can balance the prediction on all category groups.

Stages and Iterations of BGNN To investigate the multiple stage refinement and iterative context encoding in BGNN, we incrementally apply several sets of hyperparameters on number of stages N_t and iterations N_i in model design. The quantitative results in Tab. 4 indicate the message propagation with 3 iterations achieves the best performance in one-stage BGNN. Furthermore, by freezing the iteration number N_i as 3, we find the performance increases with more stages, and will saturate when $N_t=3$.

4.2. Open Images

4.2.1 Experiments Setting

Dataset Details. The **Open Images** dataset [21] is a large-scale dataset proposed by Google recently. Com-

pared with Visual Genome dataset, it has a superior annotation quality for the scene graph generation. In this work, we conduct experiments on Open Images V4&V6, we follow the similar data processing and evaluation protocols in [60, 21, 27].

The Open Images V4 is introduced as a benchmark for scene graph generation by [60] and [27], which has 53,953 and 3,234 images for the train and val sets, 57 objects categories, and 9 predicate categories in total. The Open Images V6 has 126,368 images used for training, 1813 and 5322 images for validation and test, respectively, with 301 object categories and 31 predicate categories. This dataset has a comparable amount of semantics categories with the VG.

Evaluation Protocol For Open Images V4&V6, we follow the same data processing and evaluation protocols in [60, 21, 27]. The mean Recall@50 (mR@50), Recall@50 (R@50), weighted mean AP of relationships (wmAP_{rel}), and weighted mean AP of phrase (wmAP_{phr}) are used as evaluation metrics. Following standard evaluation metrics of Open Images refer to [21, 60, 27], the weight metric score_{wtd} is computed as: score_{wtd} = 0.2 × R@50 + 0.4 × wmAP_{rel} + 0.4 × wmAP_{phr}. Besides, we also report mRecall@K like Visual Genome as a balanced metric for comprehensive comparison.

4.2.2 Quantitative Results

Performance on these two datasets are reported in Tab. 5. For Open Images V4, which includes only 9 predicate categories in total, make it constrained to explore the biased scene graph generation tasks. Our method can still achieve competitive results on weighted metric score, and outperform the previous work on mean recall with a significant margin. For Open Images V6 with 31 predicate categories, we reimplement several recent works [27, 39, 40, 59, 53] for fair comparison. As shown in Tab. 5, our method achieves the state-of-the-art performance on mean recall and competitive results on weighted metric score. Results on both datasets demonstrate the efficacy of our approach.

5. Conclusion

In this work, we have proposed a novel bipartite graph neural network (BGNN) for unbiased scene graph generation. Compared to previous methods, our main contribution consists of two key components as follows. We first develop a confidence-aware message passing mechanism for our BGNN to encode scene context in an effective manner; Moreover, we design a bi-level resampling strategy to mitigate the imbalanced data distribution during the training. The results evidently show that our BGNN achieves the superior or comparable performances over the prior state-of-the-art approaches on all three scene graph datasets.

References

- [1] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems*, pages 1567–1578, 2019. 1, 2
- [2] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002. 2
- [3] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6163–6171, 2019. 1, 2, 3, 6, 7
- [4] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019. 1, 2
- [5] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *Proceedings of the IEEE conference on computer vision and Pattern recognition*, pages 3076–3086, 2017. 2
- [6] Chris Drummond, Robert C Holte, et al. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on learning from imbalanced datasets II*, volume 11, pages 1–8. Citeseer, 2003. 2
- [7] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018. 3
- [8] Jiuxiang Gu, Handong Zhao, Zhe Lin, Sheng Li, Jianfei Cai, and Mingyang Ling. Scene graph generation with external knowledge and image reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2019. 1
- [9] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5356–5364, 2019. 2, 3, 5, 6, 7, 8
- [10] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pages 1024–1034, 2017. 3
- [11] Marcel Hildebrandt, Hang Li, Rajat Koner, Volker Tresp, and Stephan Günnemann. Scene graph reasoning for visual question answering. *arXiv preprint arXiv:2007.01072*, 2020. 1
- [12] Yifan Hou, Jian Zhang, James Cheng, Kaili Ma, Richard TB Ma, Hongzhi Chen, and Ming-Chang Yang. Measuring and improving the use of graph information in graph neural networks. In *International Conference on Learning Representations*, 2019. 3
- [13] Xinting Hu, Yi Jiang, Kaihua Tang, Jingyuan Chen, Chunyan Miao, and Hanwang Zhang. Learning to segment the tail. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14045–14054, 2020. 2, 3, 5
- [14] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015. 1
- [15] Michael Kampffmeyer, Yinbo Chen, Xiaodan Liang, Hao Wang, Yujia Zhang, and Eric P Xing. Rethinking knowledge graph propagation for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11487–11496, 2019. 3
- [16] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *International Conference on Learning Representations (ICLR)*, 2019. 3
- [17] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems*, 29(8):3573–3587, 2017. 2
- [18] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. 2016. 3
- [19] Boris Knyazev, Harm de Vries, Cătălina Cangea, Graham W. Taylor, Aaron Courville, and Eugene Belilovsky. Graph Density-Aware Losses for Novel Compositions in Scene Graph Generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2017. 3
- [20] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 1, 7
- [21] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision (IJCV)*, 2020. 7, 8
- [22] Yikang Li, Wanli Ouyang, Bolei Zhou, Jianping Shi, Chao Zhang, and Xiaogang Wang. Factorizable net: an efficient subgraph-based framework for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 335–351, 2018. 1, 2
- [23] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1261–1270, 2017. 1, 2, 6
- [24] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. In *International Conference on Learning Representations (ICLR)*, 2015. 3
- [25] Yu Li, Tao Wang, Bingyi Kang, Sheng Tang, Chunfeng Wang, Jintao Li, and Jiashi Feng. Overcoming classifier im-

- balance for long-tail object detection with balanced group softmax. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10991–11000, 2020. [2](#)
- [26] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. [2](#), [6](#)
- [27] Xin Lin, Changxing Ding, Jinquan Zeng, and Dacheng Tao. Gps-net: Graph property sensing network for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3746–3753, 2020. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [28] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision (ECCV)*, pages 21–37. Springer, 2016. [2](#)
- [29] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2537–2546, 2019. [3](#), [7](#)
- [30] Xin Lu, Buyu Li, Yuxin Yue, Quanquan Li, and Junjie Yan. Grid r-cnn. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7363–7372, 2019. [2](#)
- [31] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 181–196, 2018. [2](#)
- [32] Julia Peyre, Ivan Laptev, Cordelia Schmid, and Josef Sivic. Detecting unseen visual relations using analogies. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1981–1990, 2019. [3](#)
- [33] Mengshi Qi, Weijian Li, Zhengyuan Yang, Yunhong Wang, and Jiebo Luo. Attentive relational networks for mapping images to scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3957–3966, 2019. [2](#)
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. [3](#), [7](#)
- [35] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008. [3](#)
- [36] Li Shen, Zhouchen Lin, and Qingming Huang. Relay back-propagation for effective learning of deep convolutional neural networks. In *European conference on computer vision*, pages 467–482. Springer, 2016. [2](#)
- [37] Jiaxin Shi, Hanwang Zhang, and Juanzi Li. Explainable and explicit visual reasoning over scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8376–8384, 2019. [1](#)
- [38] Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, and Junjie Yan. Equalization loss for long-tailed object recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11662–11671, 2020. [2](#)
- [39] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3716–3725, 2020. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [40] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6619–6628, 2019. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [41] Damien Teney, Lingqiao Liu, and Anton van Den Hengel. Graph-structured representations for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2017. [1](#)
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [3](#)
- [43] Tzu-Jui Julius Wang, Selen Pehlivan, and Jorma Laaksonen. Tackling the unannotated: Scene graph generation with bias-reduced models. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2020. [2](#), [3](#)
- [44] Wenbin Wang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Exploring context and visual pattern of relationship for scene graph generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2019. [2](#)
- [45] Wenbin Wang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Sketching image gist: Human-mimetic hierarchical scene graph generation. *arXiv preprint arXiv:2007.08760*, 2020. [1](#), [2](#)
- [46] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. [3](#)
- [47] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The World Wide Web Conference*, pages 2022–2032, 2019. [3](#)
- [48] Sanghyun Woo, Dahun Kim, Donghyeon Cho, and In So Kweon. Linknet: Relational embedding for scene graph. In *Advances in Neural Information Processing Systems*, pages 560–570, 2018. [2](#)
- [49] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. [7](#)
- [50] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5410–5419, 2017. [1](#), [2](#), [7](#)

- [51] Xiaoran Xu, Wei Feng, Yunsheng Jiang, Xiaohui Xie, Zhiqing Sun, and Zhi-Hong Deng. Dynamically pruned message passing networks for large-scale knowledge graph reasoning. *arXiv preprint arXiv:1909.11334*, 2019. [3](#)
- [52] Shaotian Yan, Chen Shen, Zhongming Jin, Jianqiang Huang, Rongxin Jiang, Yaowu Chen, and Xian-Sheng Hua. Pcp1: Predicate-correlation perception learning for unbiased scene graph generation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 265–273, 2020. [3](#), [6](#), [7](#)
- [53] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 670–685, 2018. [1](#), [2](#), [6](#), [8](#)
- [54] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10685–10694, 2019. [1](#), [2](#)
- [55] Guojun Yin, Lu Sheng, Bin Liu, Nenghai Yu, Xiaogang Wang, Jing Shao, and Chen Change Loy. Zoom-net: Mining deep feature interactions for visual relationship recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 322–338, 2018. [2](#)
- [56] Cong Yuren, Hanno Ackermann, Wentong Liao, Michael Ying Yang, and Bodo Rosenhahn. Nodis: Neural ordinary differential scene understanding. *arXiv preprint arXiv:2001.04735*, 2020. [2](#)
- [57] Alireza Zareian, Svebor Karaman, and Shih-Fu Chang. Bridging knowledge graphs to generate scene graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [3](#)
- [58] Alireza Zareian, Haoxuan You, Zhecan Wang, and Shih-Fu Chang. Learning visual commonsense for robust scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [3](#)
- [59] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [60] Ji Zhang, Kevin J. Shih, Ahmed Elgammal, Andrew Tao, and Bryan Catanzaro. Graphical Contrastive Losses for Scene Graph Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#), [3](#), [6](#), [8](#)
- [61] Songyang Zhang, Xuming He, and Shipeng Yan. Latentgnn: Learning efficient non-local relations for visual recognition. In *International Conference on Machine Learning*, pages 7374–7383, 2019. [3](#)
- [62] Yizhou Zhang, Yun Xiong, Xiangnan Kong, Shanshan Li, Jinhong Mi, and Yangyong Zhu. Deep collective classification in heterogeneous information networks. In *Proceedings of the 2018 World Wide Web Conference*, pages 399–408, 2018. [3](#)
- [63] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *Proceedings of*

the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9719–9728, 2020. [3](#)