# Fine-Grained Scene Graph Generation via Sample-Level Bias Prediction

Yansheng Li<sup>1</sup>, Tingzhu Wang<sup>1†</sup>, Kang Wu<sup>1</sup>, Linlin Wang<sup>1</sup>, Xin Guo<sup>2</sup>, and Wenbin Wang<sup>3</sup>

<sup>1</sup> School of Remote Sensing and Information Engineering, Wuhan University <sup>2</sup> Ant Group

<sup>3</sup> Hubei Key Laboratory of Intelligent Vision Monitoring for Hydroelectric Engineering, College of Computer and Information Technology, China Three Gorges University

tingzhu.wang@whu.edu.cn

Abstract. Scene Graph Generation (SGG) aims to explore the relationships between objects in images and obtain scene summary graphs, thereby better serving downstream tasks. However, the long-tailed problem has adversely affected the scene graph's quality. The predictions are dominated by coarse-grained relationships, lacking more informative fine-grained ones. The union region of one object pair (i.e., one sample) contains rich and dedicated contextual information, enabling the prediction of the sample-specific bias for refining the original relationship prediction. Therefore, we propose a novel Sample-Level Bias Prediction (SBP) method for fine-grained SGG (SBG). Firstly, we train a classic SGG model and construct a correction bias set by calculating the margin between the ground truth label and the predicted label with one classic SGG model. Then, we devise a Bias-Oriented Generative Adversarial Network (BGAN) that learns to predict the constructed correction biases, which can be utilized to correct the original predictions from coarse-grained relationships to fine-grained ones. The extensive experimental results on VG, GQA, and VG-1800 datasets demonstrate that our SBG outperforms the state-of-the-art methods in terms of Average@K across three mainstream SGG models: Motif, VCtree, and Transformer. Compared to dataset-level correction methods on VG, SBG shows a significant average improvement of 5.6%, 3.9%, and 3.2% on Average@K for tasks PredCls, SGCls, and SGDet, respectively. The code will be available at https://github.com/Zhuzi24/SBG.

Keywords: Scene graph generation  $\cdot$  Long-tailed distribution  $\cdot$  Bias correction  $\cdot$  Fine-grained relationships

# 1 Introduction

In recent years, Scene Graph Generation (SGG) has emerged as a popular area of research. SGG aims to generate a structured summary graph from an im-

<sup>&</sup>lt;sup>†</sup> indicates the corresponding author.



**Fig. 1:** (a) The long-tailed distribution of relationships in the well-known Visual Genome (VG) dataset [17]. (b) Workflow for bias correction: a sample-specific correction bias is predicted by the contextual information from the union region of <man, beach> and the original prediction, then it corrects the prediction from the coarse-grained "on" to the fine-grained "walking on". The y-axis label "classification score" in the figure represents the classification value before the softmax function.

age. This graph consists of nodes and edges, where nodes represent objects in the image, with each node containing category and location information. The edges represent the relationships between the objects. The structured semantic graph clearly describes the objects, their attributes, and the relationships between them in the scene. It offers the abundant semantic information and has the powerful representational capability. Therefore, it can strongly support various downstream tasks, including visual question answering [10, 44], image retrieval [36, 52], and image captioning [12, 48].

However, SGG suffers from the long-tailed problem, which is prevalent across most datasets. Fig. 1(a) depicts the long-tailed distribution of relationships in the well-known Visual Genome (VG) dataset [17]. There is a significant abundance of coarse-grained head categories, such as "on", "has", and "wearing", whereas fine-grained tail categories, like "walking on", "part of", and "flying in", are comparatively scarce. The long-tailed problem significantly impacts the model's predictions for relationships, as depicted in Fig. 1(b). The prediction of the classic SGG model for <man, beach>'s relationship is influenced by this issue. The model predicts a coarse-grained relationship "on", which lacks the sufficient information.

Several methods [1, 2, 4, 7, 40] have been proposed to mitigate the impact posed by the long-tailed problem. DLFE [4] and RTPB [2] address this problem by correcting relationship predictions. DLEF employs the estimated per-class label frequency to recover the unbiased relationship probabilities from the biased ones. RTPB leverages the resistance bias to improve the model's detecting ability for tail categories. Both of them share a common limitation in that they apply the estimated per-class label frequency and resistance bias for all object pairs' relationship predictions, neglecting the specificity of each object pair (i.e., each sample). They can be categorized as the dataset-level correction, as the estimated per-class label frequency and resistance bias are the overall information of the dataset. In contrast to them, we are dedicated to focusing on capturing the characteristics of each object pair. The contextual information differs among different object pairs' union regions, enabling us to predict sample-specific biases for correcting the corresponding relationship predictions. This is also denoted as the sample-level bias correction. The prior information of relationships typically facilitates the generation of the scene graph, so we incorporate relationships' global bias into the bias prediction process. The global bias is the relationship's prior bias obtained through the statistics of the dataset. The process of bias correction is depicted in Fig. 1(b), a sample-specific bias is predicted to correct the coarse-grained "on" to the fine-grained "walking on".

Inspired by the above idea, we propose a novel Sample-Level Bias Prediction (SBP) method for fine-grained SGG (SBG). We first train a classic SGG model and construct a correction bias set by calculating the margin between the ground truth label and the predicted label with one classic SGG model. Each correction bias in the set represents a vector that contains bias values for each relationship category and is designed to accurately correct the corresponding original prediction. Next, we aim to predict a vector that closely approximates the constructed correction bias. The conventional methods, such as fully connected networks [37] and one-dimensional convolutional networks [34], have relatively simple loss constraints and weak generative capabilities. Therefore, we devise a Bias-Oriented Generative Adversarial Network (BGAN). BGAN employs an adversarial training approach where the generator and discriminator compete with each other. This adversarial learning process empowers the BGAN with robust generative capabilities to generate the desired vector. The constructed correction bias set is utilized to train the generator and discriminator in BGAN. Finally, the generator in BGAN predicts sample-specific bias for refining the original predictions from the classic SGG model, further obtaining fine-grained relationships.

To assess the effectiveness of our SBG, we introduce the Average@K [7,18], which can better reflect the overall performance of all relationships. In sum, this paper provides the following valuable contributions to the SGG community in the hope of advancing the fine-grained SGG research:

- This paper, for the first time, explores sample-level bias correction to tackle the long-tailed problem, promoting refining the coarse-grained relationships into the fine-grained ones.
- We devise a novel Bias-Oriented Generative Adversarial Network (BGAN), which leverages the contextual information to predict the sample-specific correction bias.

- 4 Y. Li et al.
- The experiments on VG, GQA, and VG-1800 datasets demonstrate that our SBG outperforms the state-of-the-art methods and shows a significant improvement compared to dataset-level correction methods on Average@K.

# 2 Related Work

## 2.1 Scene Graph Generation

SGG [24, 28, 45, 47] decodes images into structured semantic graphs, supporting various downstream tasks [10, 12, 31, 36, 52]. Early works were primarily dedicated to designing models to achieve the relationship predictions between the objects. [29] leveraged language priors from semantic word embeddings to finetune the predicted relationships, [5] exploited the statistical dependencies between objects and their relationships, and [25] designed a ranking objective function by enforcing the annotated relationships to have higher relevance scores. The method described above neglected the rich visual context information from images. To address this problem, some methods designed elaborate feature refinement modules to encode the visual context information, such as message passing strategies [21, 46], sequential LSTMs [41, 51], graph neural networks [3, 27, 50], and self-attention networks [11, 30, 38]. In addition, there are now some datasets as well as methods [19, 23] emerging in the remote sensing field.

Later works focused on designing optimization frameworks to further improve the performance of SGG models. Many optimization frameworks were proposed to tackle the long-tailed problem of relationships. [40] used causal inference for unbiased SGG, [39] proposed an energy-based constraint loss to learn relationships in small numbers. [7,9] employed hierarchical or grouped strategies to gradually learn, thereby mitigating the impact of the long-tailed effect. [2,4] corrected the relationship predictions to obtain the unbiased scene graph. [1] leveraged a within-triplet Bayesian network to eradicate the long-tailed effect. [18] proposed a compositional feature augmentation strategy to mitigate the bias from the perspective of increasing the diversity of triplet features. There were also some methods using data augmentation strategies to improve the performance of SGG models, such as [8, 13, 53]. Optimization frameworks offer greater flexibility and improved malleability, so it is a promising research direction.

## 2.2 Long-Tailed Learning

Recently, several correction methods have emerged to address the long-tailed problem in SGG. In Fig. 2, DLFE [4] and RTPB [2] employed different ways to correct the original predictions. DLFE estimated the per-class label frequency  $\mathbf{c}$  and used it to correct the biased probabilities. By dividing the biased probabilities by  $\mathbf{c}$ , DLFE obtained unbiased probabilities. RTPB utilized the resistance bias  $\mathbf{b}$  to enhance the model's focus towards tail relationships. RTPB adjusted the relationship's classification logits by subtracting  $\mathbf{b}$ , thereby correcting the biased relationship predictions.



Fig. 2: Compare our SBG with DLFE [4] and RTPB [2] for corrections. DLFE and RTPB apply the same  $\mathbf{c}$  or  $\mathbf{b}$  to correct all predictions while our SBG predicts the sample-specific bias for each prediction.

In DLFE,  $\mathbf{c}$ , and in RTPB,  $\mathbf{b}$ , are both vectors that capture the overall information of relationships in the dataset. As shown in Fig. 2, DLFE and RTPB apply the same  $\mathbf{c}$  or  $\mathbf{b}$  to correct all predictions, so they do not sufficiently consider the specificities of different object pairs. They are the sample-insensitive bias correction which is also denoted as the dataset-level bias correction. In contrast, our SBG focuses on the characteristics of each object pair to predict the sample-specific bias for refining the original relationship predictions. Our SBG is the sample-sensitive bias correction which is also denoted as the sample-level bias correction.

## 3 Method

#### 3.1 Problem Setting & Overview Framework

**Problem Setting.** The task of SGG is to generate a connected summary graph  $\mathcal{G}$  from an input image *I*. The common SGG methods are the two-stage processes involving object detection followed by relationship prediction. In the first stage, we detect all objects in the image *I*, denotes as  $\mathcal{O} = \{o_i\}$ . In the second stage, we predict relationships  $\mathcal{R} = \{r_{i \to j}\}$  of object pairs  $(o_i, o_j)$  and then generate  $(o_i, r_{i \to j}, o_j)$  triplets. Therefore, the scene graph can be formulated as:

$$\mathcal{G} = \{ (o_i, r_{i \to j}, o_j) | o_i, o_j \in \mathcal{O}, r_{i \to j} \in \mathcal{R} \}.$$

$$\tag{1}$$

**Overall Framework.** Fig. 3 illustrates the overall structure of SBG, and the entire workflow follows the common two-stage SGG pipeline. Consistent with previous methods [2, 40], we utilize Faster R-CNN [35] as the object detection network. To address the long-tailed problem of relationships, we present a novel method SBP to realize the sample-level bias correction. More detailed analysis can be seen in Sec. 3.2.

#### 3.2 Sample-Level Bias Prediction

The unbiased relationship predictions can be recovered from the biased ones [4]. In an ideal scenario, training the SGG model on an unbiased dataset results in



**Fig. 3:** The overall structure of our SBG. After constructing the correction bias set, BGAN learns to predict the constructed correction biases for achieving the sample-level bias correction. The coarse-grained scene graph is generated by the classic SGG model. v is the value at "parked on" in  $\mathbf{b}^{tru}$ .

unbiased relationship predictions which are denoted as  $\mathbf{z}_u = [z_{u1}, z_{u2}, ..., z_{uM}]$ , where M is the number of relationship categories. However, datasets often suffer from the severe long-tailed problem of relationships, thereby leading to biased relationship predictions which are represented by  $\mathbf{z} = [z_1, z_2, ..., z_M]$ . Both  $\mathbf{z}_u$ and  $\mathbf{z}$  are vectors of logits for relationship classification. If we can obtain the differences which are denoted as  $\mathbf{b}_{\mathbf{u}} = [b_{u1}, b_{u2}, ..., b_{uM}]$  between  $\mathbf{z}_u$  and  $\mathbf{z}$ , where an ideal  $\mathbf{b}_{\mathbf{u}}$  equals  $\mathbf{z}_u - \mathbf{z}$ , then we can recover the unbiased predictions  $\mathbf{z}_u$ from the biased predictions  $\mathbf{z}$ . However, in reality, obtaining  $\mathbf{b}_{\mathbf{u}}$  is challenging as acquiring  $\mathbf{z}_u$  is nearly impossible. This poses difficulties in recovering  $\mathbf{z}_u$  from  $\mathbf{z}$ .

An alternative perspective can provide a further insight.  $\mathbf{z}_u$  can accurately predict relationships while other data vectors which are denoted as  $\hat{\mathbf{z}}$  may also achieve accurate relationship predictions. Although  $\mathbf{z}_u$  and  $\hat{\mathbf{z}}$  may have significant differences, their predicted relationships are the same. The similarity between  $\mathbf{z}_u$ and  $\hat{\mathbf{z}}$  here can be expressed as  $argmax(softmax(\mathbf{z}_u)) = argmax(softmax(\hat{\mathbf{z}}))$ , and the difference between them is that  $\mathbf{z}_u$  is ideal and unbiased and  $\hat{\mathbf{z}}$  is realizable and not unbiased. In reality, obtaining  $\hat{\mathbf{z}}$  is relatively easier compared to acquiring  $\mathbf{z}_u$ . Images contain abundant information, offering possibilities for generating  $\hat{\mathbf{z}}$ . The union regions of object pairs contain rich and dedicated contextual information, enabling us to predict the sample-specific biases  $\mathbf{b}_s$ , ensuring  $\hat{\mathbf{z}} = \mathbf{z} + \mathbf{b}_s$ . By applying the sample-specific biases  $\mathbf{b}_s$  to  $\mathbf{z}$ , we achieve the sample-level bias correction for the long-tailed problem.

To realize the aforementioned idea, we propose the SBP to predict the samplespecific bias  $\mathbf{b}_{\mathbf{s}i,j}$  for correcting the biased relationship prediction  $\mathbf{z}_{i,j}$ , where  $\mathbf{z}_{i,j}$  denotes the relationship prediction of the *j*-th object pair in the *i*-th image. The parameters of the classic SGG model and SBP are denoted as  $\Lambda^O$  and  $\Lambda^B$ , respectively. Let  $\mathbf{z}_{i,j} = \Upsilon(\Lambda^O; I_i)$  and  $\mathbf{b}_{\mathbf{s}_{i,j}} = \Upsilon(\Lambda^B; f_{uni}^{i,j}, \mathbf{b}^{glo}, \mathbf{z}_{i,j})$  represent the outputs of the classic SGG model and SBP, respectively. Thus, the sample-level bias correction for the original prediction can be expressed as:

$$\begin{aligned} \hat{\mathbf{z}}_{i,j} &= \Upsilon(\Lambda^O; I_i) + \Upsilon(\Lambda^B; f_{uni}^{i,j}, \mathbf{b}^{glo}, \mathbf{z}_{i,j}) \\ &= \mathbf{z}_{i,j} + \mathbf{b}_{\mathbf{s}_{i,j}}, \end{aligned}$$
(2)

where  $\Upsilon(\cdot; \cdot)$  represents the process of obtaining the neural network output based on the neural network's parameters and input data.  $f_{uni}^{i,j}$  is the union feature of the *j*-th object pair in the *i*-th image. It can be obtained by extracting the visual feature of the object pair's union region using a shared weight feature extractor with Faster R-CNN in the classic SGG model (as shown in Fig. 3), and then fusing it with the spatial feature of object pair's union region.  $\mathbf{b}^{glo}$  is the relationships' global bias, serving as the prior bias of the relationships. It is calculated as  $-log(w^a/\sum_{j\in M}w_j^a + \epsilon)$  [2], where w represents the weights of the relationships, and a and  $\epsilon$  are hyper-parameters.  $\mathbf{z}_{i,j}$  is relationship prediction after the sample-level bias correction. To guide the training of SBP, we construct a set of correction biases  $\mathbf{b}^{tru}$  where the set is denoted as  $\mathcal{S} = {\mathbf{b}^{tru}}$  according to the ground truth labels by Eq. (2).

**Fig. 4:** The construction workflow of Correction Bias Set S.



A	lgorithm 1: The training of									
E	BGAN for one iteration.									
	Input: $f_{uni}$ , $\mathbf{b}^{glo}$ , $\mathbf{z}$ , and $\mathbf{b}^{tru}$									
1	1 for $k \leftarrow 1$ to 5 do									
2	1. clear the grad of D;									
3	2. $\mathbf{b}^{pre} = \Upsilon(\Lambda^{B_G}; f_{uni}, \mathbf{b}^{glo}, \mathbf{z});$									
4	3. calculate $\mathcal{L}_D$ using $\mathbf{b}^{pre}$ and $\mathbf{b}^{tru}$ ;									
5	4. update the parameters $\Lambda^{B_D}$ of $D$ ;									
6	if $k = 5$ then									
7	1. clear the grad of G;									
8	2. $\mathbf{b}^{pre} = \Upsilon(\Lambda^{B_G}; f_{uni}, \mathbf{b}^{glo}, \mathbf{z});$									
9	3. calculate $\mathcal{L}_G$ using $\mathbf{b}^{pre}$ ;									
10	4. update the parameters $\Lambda^{B_G}$									
	of $G$ ;									
11	end									
12	end									
13	update the learning rates of $D$ and $G$ .									

Construction of Correction Bias Set. In each training iteration, S is constructed by calculating the margin between the ground truth label  $r_{tru}$  and the predicted label  $r_{pre}$  from one classic SGG model. The construction workflow of S is summarized in Fig. 4. The fusion of  $f_{uni}$  and  $\mathbf{b}^{glo}$  results in  $\mathbf{b}^{tru} = \phi(f_{uni}) + \mathbf{b}^{glo}$ , where  $\phi$  denotes the mapping of high-dimensional feature to one-dimensional feature using an encoder that includes a single layer of transformer [43]. Subsequently, we assess whether  $\mathbf{b}^{tru}$  meets the condition  $r_{tru} = r_{pre}$ , where  $r_{pre} = argmax(softmax(\hat{\mathbf{z}}))$  and  $\hat{\mathbf{z}} = \mathbf{z} + \mathbf{b}^{tru}$ . If  $\mathbf{b}^{tru}$  satisfies the condition, it is included in S; otherwise, the difference  $d = \hat{\mathbf{z}}[r_{tru}] - \hat{\mathbf{z}}[r_{pre}]$  is calculated. After updating  $\mathbf{b}^{tru}[r_{pre}] = \mathbf{b}^{tru}[r_{pre}] + d + \varepsilon$ , the modified  $\mathbf{b}^{tru}$  is added to S, where  $\varepsilon$  denotes a small non-zero value.

**Training of Bias Prediction.** The constructed S can accurately correct the original predictions. We aim to learn to predict the constructed biases. The closer the predicted biases are to the constructed biases, the better the effectiveness of the bias correction. To achieve this target, we propose the BGAN, which employs one-dimensional convolutions to model the prediction process. The generator (denoted as G) in BGAN takes the object pairs' union features  $f_{uni}$ , the relationships' global bias  $\mathbf{b}^{glo}$ , and the original predictions  $\mathbf{z}$  as input to predict the sample-specific biases  $\mathbf{b}^{pre}$ . The predicted biases  $\mathbf{b}^{pre}$  and the constructed biases  $\mathbf{b}^{tru}$  from S are fed into the discriminator (denoted as D) in BGAN for the training. Through the adversarial training between G and D, G learns to predict the sample-specific biases for correcting the original predictions. Additionally, the cross-entropy loss [6] is computed using the corrected predictions  $\hat{\mathbf{z}}$ to further constrain the generator G. The inputs and outputs of G and D can be represented as:

$$\mathbf{b}^{pre} = \Upsilon(\Lambda^{B_G}; f_{uni}, \mathbf{b}^{glo}, \mathbf{z}), \tag{3}$$

$$\mathcal{T}_G = \Upsilon(\Lambda^{B_D}; \mathbf{b}^{pre}), \mathcal{T}_S = \Upsilon(\Lambda^{B_D}; \mathbf{b}^{tru}), \tag{4}$$

where  $\Lambda^{B_G}$  and  $\Lambda^{B_D}$  represent the parameters of G and D, respectively, and they both belong to the parameters  $\Lambda^B$ . The  $\mathcal{T}_G$  and  $\mathcal{T}_S$  represent the discriminative scores of D for  $\mathbf{b}^{pre}$  and  $\mathbf{b}^{tru}$ . The losses of G and D can be expressed as:

$$\mathcal{L}_G = -mean(\mathcal{T}_G) + \alpha \cdot (-t_m log(\frac{e^{\hat{\mathbf{z}}_m}}{\sum_m^M e^{\hat{\mathbf{z}}_m}})), \tag{5}$$

$$\mathcal{L}_D = -mean(\mathcal{T}_S) + mean(\mathcal{T}_G),\tag{6}$$

where  $\alpha$  represents the weight factor for the cross-entropy loss.  $t_m$  takes the value 1 only when m is the ground truth label; otherwise, it is set to 0. The detailed training workflow of BGAN is outlined in Algorithm 1. Thus, the overall loss is defined as follows:

$$\min_{\Lambda^O,\Lambda^B} \mathcal{L}_{total} = \mathcal{L}_{SGG} + \beta \cdot \mathcal{L}_{BGAN},\tag{7}$$

where  $\mathcal{L}_{SGG}$  represents the cross-entropy loss of the classic SGG model, which takes **z** as input.  $\mathcal{L}_{BGAN}$  stands for  $\mathcal{L}_G$  and  $\mathcal{L}_D$ . During the training of the classic SGG model,  $\beta$  is set to 0. After the training of the classic SGG model is completed, its parameters  $\Lambda^O$  are frozen, and the training of BGAN is conducted by setting  $\beta$  to 1. Once the training of BGAN is finished, the classic SGG model and the generator G in BGAN are jointly used for inference. G predicts the sample-specific biases to correct the original predictions, and further obtain the fine-grained relationships.

# 4 Experiments

## 4.1 Experimental Settings

**Datasets.** We conduct experiments on three datasets: VG, GQA, and VG-1800. In SGG, the most commonly used database is VG [17]. From the orig-

inal database, the most frequent 150 object and 50 relationship categories are retained [40,46]. GQA [14] is a vision-and-language dataset similar to VG, consisting of a total of 113k images. Following [1], we split the GQA dataset for the training. We retain images only with the most frequent 160 object and 60 relationship categories. VG-1800 [53] is a large-scale dataset, containing 70k object and 1.8k relationship categories. For both VG and GQA, we use 70% of the images for the training and the remaining 30% for the testing. Similar to [40], we sample a 5k validation set for the parameter tuning. For VG-1800, settings we keep in line with IETrans [53].

**Tasks.** We perform three conventional SGG sub-tasks: (1) Predicate Classification (PredCls): given object bounding boxes and object labels, predicting the relationships between the objects. (2) Scene Graph Classification (SGCls): given object bounding boxes, predicting the object labels and relationships between the objects. (3) Scene Graph Detection (SGDet): performing object detection and then predicting relationships between the detected objects.

**Evaluation Metrics.** We evaluate SGG models on four metrics: (1) Recall@K (R@K) [51]: It indicates the proportion of ground truths that appear among the top-K confident predicted triplets. (2) mean Recall@K (mR@K) [9]: It is the average of R@K scores which are calculated for each relationship category separately. (3) Average@K (A@K) [18]: It calculates the average scores of R@K and mR@K. Since R@K favors head relationships while mR@K favors tail relationships, the A@K can provide a more intuitive and comprehensive reflection of the overall performance of all relationships. (4) F-Acc [53]: It is a metric used for VG-1800 that indicates comprehensive performance of all relationships.

### 4.2 Implementation Details

**Object Detector.** For VG and VG-1800, following [4], we employ a pre-trained Faster R-CNN [35] with ResNeXt-101-FPN [45] backbone. In the training stage, the parameters of the detector are fixed to reduce the computation cost. For GQA, we train a Faster-RCNN with ResNeXt-101-FPN as the object detector, and then the detector reaches 26.0 (%) mAP on the test set.

**Model Settings.** For the classic SGG model, we set the batch size to 16 and employ a SGD optimizer [15] with an initial learning rate of 0.001. The training procedure contains a total of 18,000 iterations. For SBP, BGAN is comprised of one-dimensional convolutions. The number of layers for G and D are 5 and 3, respectively. Both of them employ an RMSProp optimizer [54] with initial learning rates of 0.0001 and 0.0005.  $\varepsilon$  is set to 0.0001. The training process contains a total of 4,000 iterations. All experiments are performed on NVIDIA 24G GeForce RTX 3090 GPUs.

#### 4.3 Experiments on VG

Comparison with the State-of-the-Art Two-Stage Methods on VG. We compare our SBG with the state-of-the-art two-stage methods on three main-stream SGG models: Motif [51], VCtree [41], and Transformer [40], as illustrated

Table 1: Performance comparison with the state-of-the-art two-stage methods on PredCls, SGCls, and SGDet tasks of VG for R@50/100 (%), mR@50/100 (%), and A@50/100 (%). The underlined values represent suboptimal performance. The light blue font belongs to the dataset-level correction.

		PredCls			SGCls			SGDet	
Model	R@50/100	mR@50/100	A@50/100	R@50/100	mR@50/100	A@50/100	R@50/100	mR@50/100	A@50/100
Motif [51]	65.4 / 67.2	18.0 / 19.3	41.7 / 43.3	40.0 / 40.8	10.0 / 10.6	25.0 / 25.7	32.5 / 37.0	7.8 / 7.4	20.2 / 23.2
+CogTree [49]	35.6 / 36.8	26.4 / 29.0	31.0 / 32.9	21.6 / 22.2	14.9 / 16.1	18.3 / 19.2	20.0 / 22.1	10.4 / 11.8	15.2 / 17.0
+ DLFE [4]	52.4 / 54.3	26.7 / 28.7	39.6 / 41.5	32.3 / 33.1	15.2 / 15.9	23.8 / 24.5	25.4 / 29.4	11.7 / 13.8	18.6 / 21.6
+GCL [9]	42.7 / 44.4	36.1 / 38.2	39.4 / 41.3	26.1 / 27.1	20.8 / 21.8	23.5 / 24.5	18.4 / 22.0	16.8 / 19.3	17.6 / 20.7
+HML [7]	47.1 / 49.1	36.3 / 38.7	41.7 / 43.9	26.1 / 27.4	20.8 / 22.1	23.5 / 24.8	17.6 / 21.1	14.6 / 17.3	16.1 / 19.2
+IETrans [53]	54.7 / 56.7	30.9 / 33.6	<u>42.8</u> / 45.2	32.5 / 33.4	16.8 / 17.9	24.7 / <u>25.7</u>	26.4 / 30.6	12.4 / 14.9	19.4 / 22.8
+RTPB [2]	40.3 / 42.6	35.4 / 37.4	37.9 / 40.0	26.0 / 26.9	20.0 / 21.0	23.0 / 24.0	19.0 / 22.5	13.1 / 15.5	16.1 / 19.0
+PPDL [22]	47.2 / 47.6	32.2 / 33.3	39.7 / 40.5	28.4 / 29.3	17.5 / 18.2	23.0 / 23.8	21.2 / 23.9	11.4 / 13.5	16.3 / 18.7
+FGPL [32]	51.4 / 55.3	33.6 / 38.6	42.5 / <b>47.0</b>	23.5 / 24.1	21.2 / 22.3	22.4 / 23.2	20.9 / 23.5	15.2 / 18.2	18.1 / 20.9
+CFA [18]	42.3 / 45.1	39.9 / 43.0	41.1 / 44.1	25.7 / 27.4	20.9 / 22.4	23.3 / 24.9	20.7 / 24.4	15.3 / 18.1	18.0 / 21.3
+Inf[1]	51.5 / 55.1	24.7 / 30.7	38.1 / 42.9	32.1 / 33.8	14.5 / 17.4	23.3 / 25.6	27.7 / 30.1	10.4 / 11.9	19.1 / 21.0
Our SBG	55.4 / 57.3	32.1 / 34.4	43.8 / <u>45.9</u>	34.9 / 35.7	17.5 / 18.6	26.2 / 27.1	27.0 / 31.3	13.8 / 16.1	20.4 / 23.7
VCtree [41]	65.4 / 67.1	17.9 / 19.5	41.7 / 43.3	46.6 / 47.6	12.3 / 13.1	29.5 / 30.4	31.3 / 35.6	7.2 / 8.6	<u>19.3 / 22.1</u>
+CogTree [49]	44.0 / 45.4	27.6 / 29.7	35.8 / 37.6	30.9 / 31.7	18.8 / 19.9	24.9 / 25.8	18.2 / 20.4	10.4 / 12.1	14.3 / 16.3
+DLFE [4]	51.8 / 53.5	25.3 / 27.1	38.6 / 40.3	33.5 / 34.6	18.9 / 20.0	26.2 / 27.3	22.7 / 26.3	11.8 / 13.8	17.3 / 20.1
+GCL [9]	40.7 / 42.7	37.1 / 39.1	38.9 / 40.9	27.7 / 28.7	22.5 / 23.5	25.1 / 26.1	17.4 / 20.7	15.2 / 17.5	16.3 / 19.1
+HML [7]	47.0 / 48.8	36.9 / 39.2	42.0 / 44.0	27.0 / 28.4	25.0 / 26.8	26.0 / 27.6	17.6 / 21.0	13.7 / 16.3	15.7 / 18.7
+IETrans [53]	53.0 / 55.0	30.3 / 33.9	41.7 / 44.5	32.9 / 33.8	16.5 / 18.1	24.7 / 26.0	25.4 / 29.3	11.5 / 14.0	18.5 / 21.7
+RTPB [2]	41.2 / 43.3	33.4 / 35.6	37.3 / 39.5	28.7 / 30.0	24.5 / 25.8	26.6 / 27.9	18.1 / 21.3	12.8 / 15.1	15.5 / 18.2
+PPDL [22]	47.6 / 48.0	33.3 / 33.8	40.5 / 40.9	32.1 / 33.0	21.8 / 22.4	27.0 / 27.7	20.1 / 22.9	11.3 / 13.3	15.7 / 18.1
+FPGL [32]	42.3 / 43.8	37.4 / 40.3	39.9 / 42.1	27.3 / 28.0	26.2 / 27.7	26.8 / 27.9	20.8 / 23.4	15.5 / 18.4	18.2 / 20.9
+CFA [18]	41.9 / 45.0	39.2 / 42.5	40.6 / 43.8	32.3 / 33.8	26.3 / 28.3	29.3 / <u>31.1</u>	20.5 / 24.2	15.1 / 17.9	17.8 / 21.1
+Inf[1]	59.5 / 60.1	28.1 / 30.7	<u>43.8</u> / <u>45.4</u>	40.7 / 41.6	17.3 / 19.4	29.0 / 30.5	27.7 / 30.1	10.4 / 11.9	19.1 / 21.0
Our SBG	55.5 / 57.3	32.4 / 34.5	44.0 / 45.9	40.8 / 41.9	21.8 / 23.1	31.3 / 32.5	26.5 / 30.5	12.2 / 14.3	19.4 / 22.4
Transformer [40]	65.5 / 67.3	18.2 / 19.7	41.9 / 43.5	39.8 / 40.6	10.8 / 11.5	<u>25.3</u> / 26.1	32.0 / 36.4	8.2 / 9.8	20.1 / <u>23.1</u>
+CogTree [49]	38.4 / 39.7	28.4 / 31.0	$33.4 \ / \ 35.4$	22.9 / 23.4	15.7 / 16.7	19.3 / 20.1	19.5 / 21.7	11.1 / 12.7	15.3 / 17.2
+HML [7]	45.6 / 47.8	33.3 / 35.9	39.5 / 41.9	22.5 / 23.8	19.1 / 20.4	20.8 / 22.1	15.0 / 17.7	15.0 / 17.7	15.0 / 17.7
+PPDL [22]	46.5 / 47.2	35.7 / 36.0	36.4 / 41.6	28.5 / 29.3	17.9 / 18.8	23.2 / 24.1	20.7 / 24.2	11.5 / 13.2	16.1 / 18.6
+FPGL [32]	50.9 / 54.6	36.3 / 40.1	43.6 / 47.4	21.3 / 22.1	22.5 / 24.1	21.9 / 23.1	19.1 / 21.8	17.5 / 20.5	18.3 / 21.2
+CFA [18]	46.2 / 48.9	38.6 / 41.5	42.4 / 45.2	29.6 / 28.1	20.9 / 22.7	24.5 / 26.2	21.0 / 24.7	15.0 / 17.9	18.0 / 21.3
+IETrans [53]	51.8 / 53.8	30.8 / 34.7	41.3 / 44.3	32.6 / 33.5	17.4 / 19.1	25.0 / <u>26.3</u>	25.5 / 29.6	12.5 / 15.0	<u>19.0</u> / 22.3
Our SBG	55.8 / 57.6	33.3 / 35.7	<b>44.6</b> / <u>46.7</u>	35.6 / 36.5	18.5 / 19.4	27.1 / 28.0	24.5 / 28.5	15.6 / 18.0	20.1 / 23.3

in Tab. 1. For Motif, VCtree, and Transformer, our SBG improves significantly all mR@50/100, demonstrating our SBG's strong detection capability for tail relationships. Compared to the state-of-the-art two-stage methods, our SBG outperforms all other methods in terms of A@50/100, exhibiting a better-balanced improvement between R@50/100 and mR@50/100. Specifically, considering the PredCls of Transformer as an example, other methods sacrifice R@50/100 significantly while improving mR@50/100. By contrast, our SBG minimizes sacrifices on R@50/100 while maintaining a competitive performance on mR@50/100. These results not only demonstrate the effectiveness of our SBG but also show our SBG's capacity to improve tail relationships while avoiding excessive suppression of head relationships. Furthermore, our SBG is mostly highest in terms of A@K on three mainstream models, but the performance of the other methods is inconsistent (i.e., no method has equal competitiveness on three models), which reflects the strong generalization of our SBG.

Superiority of Sample-Level Bias Correction. To demonstrate the superiority of sample-level bias correction, we compare our SBG with dataset-level bias correction methods (light blue in Tab. 1) DLFE and RTPB on VG. Compared to DLFE, our method outperforms on R@50/100, mR@50/100, and A@50/100, indicating its superior performance. Although our mR@50/100 metrics (exclud-

**Table 2:** Improvement comparison of dataset-level bias correction methods DLFE [4], RTPB [2] on PredCls, SGCls, and SGDet tasks of Motif and VCtree models for R@100 (%), mR@100 (%), and A@100 (%). Imp (%) indicates the degree of improvement for Motif and VCtree models. The underlined values have the same meaning as in Tab. 1.

	PredCls				SGCls		SGDet			
Model	R@100/Imp	mR@100/Imp	A@100/Imp	R@100/Imp	mR@100/Imp	A@100/Imp	R@100/Imp	mR@100/Imp	A@100/Imp	
Motif [51]	67.2	19.3	43.3	40.8	10.6	25.7	37.0	7.4	23.2	
+ DLFE [4]	54.3 / -12.9	28.7 / +9.4	41.5 / -1.8	33.1 / - <u>7.7</u>	15.9 / +5.3	24.5 / -1.2	29.4 / - <u>7.6</u>	13.8 / +6.4	21.6 / -1.6	
+RTPB [2]	42.6 / -24.6	37.4 / +18.1	40.0 / -3.3	26.9 / -13.9	21.0 / +10.4	24.0 / -1.7	22.5 / -14.5	15.5 / + 8.1	19.0 / -4.2	
Our SBG	57.3 / -9.9	$34.4 / + \underline{15.1}$	45.9 / +2.6	35.7 / -5.1	$18.6 / + \underline{8.0}$	27.1 / +1.4	31.3 / -5.7	16.1 / +8.7	23.7 / +0.4	
VCtree [41]	67.1	19.5	43.3	47.6	13.1	30.4	35.6	8.6	22.1	
+ DLFE [4]	53.5 / - <u>13.6</u>	27.1 / +7.6	40.3 / - <u>3.0</u>	34.6 / -13.0	20.0 / +6.9	27.3 / -3.1	26.3 / - <u>9.3</u>	13.8 / +5.2	20.1 / -2.0	
+RTPB [2]	43.3 / -17.7	35.6 / +16.1	39.5 / -3.8	30.0 / -17.6	25.8 / +12.7	27.9 / - <u>2.5</u>	21.3 / -14.3	15.1 / +6.5	18.2 / -3.9	
Our SBG	57.3 / -9.8	$34.5 / + \underline{15.0}$	$45.9 \; / \; + \textbf{2.6}$	41.9 / - <b>5.7</b>	$23.1 / + \underline{10.0}$	32.5 / +2.1	30.5 / -5.1	14.3 / + 5.7	$22.4 \; / \; + \textbf{0.3}$	

ing SGDet task of Motif) are slightly lower than RTPB, we achieve a higher R@50/100 and A@50/100. These demonstrate that our SBG exhibits a better overall performance. To provide a comprehensive analysis of the balance between R@50/100 and mR@50/100, we compare the improvements on Motif and VCtree with RTPB and DLFE, as depicted in Tab. 2. We have the smallest decrease on all R@100 metrics while realizing a competitive improvement on all mR@100 metrics, indicating a more balanced improvement across the R@50/100 and mR@50/100 metrics. Additionally, we compare our SBG on the overall performance (i.e., the mean of A@50 and A@100) with DLFE and RTPB across PredCls, SGCls, and SGDet tasks on Motif and VCtree models, as illustrated in Fig. 5. The results indicate that our SBG achieves the improvements of 5.6%, 3.9%, and 3.2% across the three tasks, further demonstrating the superiority of our SBG (i.e., the sample-level bias correction).



**Fig. 5:** Comparison with DLFE and RTPB on the overall performance.

Table	3:	Effic	ciency	com	parison	with
DLFE	[4]	and	RTPE	3 [2].	$\uparrow$ and	↓ de-
note th	le p	ercer	ntage c	hange	e of our	SBG
compai	red	to D	LFE a	nd R	TPB.	

11

Method	A@100	Training Time	Inference Speed	Parameters
DLFE [4]	41.5	22.9	0.1049	253.41
Our SBG	$45.9 \uparrow 10.6\%$	$12.6 \downarrow 45.0\%$	$0.1062 \uparrow 1.2\%$	$254.75 \uparrow 0.5\%$
RTPB [2]	40.1	12.1	0.1039	254.05
Our SBG	45.9 \14.5%	$12.6 \uparrow 4.1\%$	$0.1062 \uparrow 2.2\%$	254.75 \cdot 0.3%

Tab. 3 presents the efficiency comparison on the PredCls task of Motif model for Training Time (h), Inference Speed (s/img), and Parameters (M). Compared to DLFE, we achieve a 10.6% improvement in A@100, while reducing training time by 45.0%. However, the inference speed and parameters only increase by 1.2% and 0.5%, respectively. Compared to RTPB, we achieve a 14.5% improvement in A@100, with only marginal increases of 4.1% in training time, 2.2%

in inference speed, and 0.3% in parameters. These results are deemed acceptable, indicating that our SBG achieves a significant A@100 enhancement with minimal efficiency losses.

Extensibility to One-Stage SGG Methods. We validate our SBG on one-stage SGG methods on SGDet task which they support of VG for R@50/100, mR@50/100, and A@50/100 to show the expansibility in Tab. 4. From the table, we find that our SBG has a significant improvement on mR@50/100, demonstrating our SBG's strong detection capability for tail relationships. While improving mR@50/100, there is a lesser sacrifice on R@50/100, thereby demonstrating the superior comprehensive performance. The results validate the extensibility of our SBG.

**Table 4:** Performance on one-stage SGGmethods.

		SGDet	
Model	R@50/100	mR@50/100	A@50/100
ISG [16]	30.8 / 35.6	19.5 / 23.4	25.2 / 29.5
Our SBG	29.3 / 33.6	24.1 / 28.2	26.7 / 30.9
SGTR [20]	24.6 / 28.4	12.0 / 15.2	18.3 / 21.8
+FGPL-A [33]	24.6 / 28.2	12.5 / 19.4	18.6 / 23.8
Our SBG	22.3 / 26.3	18.1 / 20.9	20.2 / 23.6
SS R-CNN [42]	33.5 / 38.4	8.6 / 10.3	21.1 / 24.4
+FGPL-A [33]	31.9 / 36.6	9.5 / 11.9	20.7 / 24.3
Our SBG	$27.3 \ / \ 31.1$	20.0 / 23.5	23.7 / 27.3

**Table 5:** Comparison on Motif for thePredCls task.

	GG	QA	V	G	VG-1800				
Model	A@50	A@100	A@50	A@100	F-Acc(Top-1)	F-Acc(Top-5)	F-Acc(Top-10)		
Motif	45.4	46.7	41.7	43.3	1.21	5.20	8.33		
+CFA	47.2	48.6	41.1	44.1	5.33	18.25	24.51		
Our SBG	$47.4 ^{\uparrow}0.2 \%$	$48.7{\uparrow}0.1\%$	$43.8^{+}2.7\%$	$45.9{\uparrow}1.8\%$	$13.45 \uparrow 8.12\%$	26.74 + 8.49%	$34.99^{10.48\%}$		

**Table 6:** Performance comparison with RTPB [2], CFA [18] on PredCls, SGCls, and SGDet tasks of GQA for R@50/100 (%), mR@50/100 (%), and A@50/100 (%). The underlined values has the same meaning as in Tab. 1.

		PredCls			SGCls			SGDet	
Model	R@50/100	$\rm mR@50/100$	A@50/100	R@50/100	$\mathrm{mR}@50/100$	A@50/100	R@50/100	mR@50/100	A@50/100
Motif [51]	66.2 / 67.6	24.5 / 25.8	45.4 / 46.7	33.4 / 34.0	12.0 / 12.5	22.7 / 23.3	30.3 / 34.4	10.5 / 12.2	20.4 / 23.3
+RTPB [2]	50.1 / 51.5	43.9 / 45.2	47.0 / 48.4	25.7 / 26.3	21.3 / 21.7	<u>23.5</u> / 24.0	21.9 / 25.6	21.0 / 23.2	21.5 / <u>24.4</u>
+CFA [18]	50.1 / 51.2	44.3 / 46.0	47.2 / 48.6	23.6 / 24	23.4 / 24.2	<u>23.5</u> / <u>24.1</u>	19.8 / 23.9	22.8 / 24.7	<u>21.3</u> / 24.3
Our SBG	56.6 / 58.0	38.1 / 39.3	47.4 / 48.7	28.6 / 29.2	18.8 / 19.4	23.7 / 24.3	25.2 / 29.4	17.8 / 19.6	21.5 / 24.5
VCtree [41]	65.5 / 67.1	24.8 / 26.3	45.5 / 46.7	33.5 / 34.1	12.5 / 13.0	23.0 / 23.6	28.4 / 32.1	9.9 / 11.7	19.2 / 21.9
+RTPB [2]	49.9 / 51.4	42.6 / 44.0	46.3 / 47.7	26.0 / 26.6	22.2 / 22.6	24.1 / 24.6	20.6 / 24.0	20.0 / 22.3	20.3 / <u>23.2</u>
+CFA [18]	48.7 / 49.3	45.1 / 47.5	<u>46.9</u> / <u>48.4</u>	26.1 / 26.7	22.9 / 23.3	<u>24.5</u> / <u>25.0</u>	18.9 / 23.0	21.9 / 23.4	<u>20.4</u> / <u>23.2</u>
Our SBG	57.9 / 59.4	37.2 / 38.5	47.6 / 49.0	29.7 / 30.3	19.5 / 20.0	24.6 / 25.2	24.3 / 27.8	16.6 / 18.8	20.5 / 23.3
Transformer [40]	67.5 / 68.9	26.8 / 28.2	47.2 / 48.6	34.6 / 35.1	14.7 / 15.2	<u>24.7</u> / <u>25.2</u>	30.6 / 34.7	12.2 / 14.1	21.4 / 24.4
+RTPB [2]	50.8 / 52.3	44.6 / 45.8	47.7 / 49.1	26.0 / 26.6	21.9 / 22.3	24.0 / 24.5	21.1 / 24.8	21.3 / 23.6	21.2 / 24.2
+CFA [18]	50.5 / 52.8	46.1 / 47.2	<u>48.3</u> / <u>50.0</u>	24.9 / 25.5	23.5 / 24.1	24.2 / 24.8	22.0 / 26.0	21.4 / 23.6	<u>21.7</u> / <u>24.8</u>
Our SBG	58.6 / 60.0	41.6 / 42.9	50.1 / 51.5	28.6 / 29.2	21.0 / 21.6	24.8 / 25.4	24.3 / 28.0	20.3 / 22.6	22.3 / 25.3

Comparison with the State-of-the-Art Two-Stage Methods on GQA and VG-1800. Following VG, we validate the effectiveness of our SBG on GQA as shown in Tab. 6. We compare the best dataset-level bias correction method RTPB [40] and the latest method CFA [18]. From the results, it can be seen that our SBG achieves a significant improvement for all baselines on mR@50/100 and outperforms RTPB and CFA on almost all A@50/100. Furthermore, we compare the gains of our SBG with CFA in Tab. 5 on three datasets GQA, VG, and VG-1800, where their data volume increase and the long-tailed effect is more severe sequentially. The results show that the gains of our SBG are more significant on larger datasets with more severe long-tailed effects. The trivial gains on GQA in Tab. 6 are due to that compared to VG and VG-1800, the data volume of GQA is smaller and the long-tailed effect of GQA is weaker.

## Ablation Studies.

i) The Effect of Region Scope for Sample-Level Bias Prediction: In our SBG, the sample-specific biases are predicted by utilizing the contextual information from object pairs' union regions which is abbreviated as *union*. We expand the region scope of contextual information to the entire image which is abbreviated as *entire* for comparison on PredCls task using of Transformer model, as shown in Tab. 7. The A@50/100 decreases when utilizing the contextual information of the entire image, indicating that the contextual information of the entire image introduces additional interference. This is because, in reality, some object pairs' union regions have already a sufficiently large scope that can provide enough contextual information.

**Table 7:** The effect of region scope for sample-level bias prediction.

**Table 8:** The effect of global bias  $\mathbf{b}^{glo}$ .  $\checkmark$  indicates the use of  $\mathbf{b}^{glo}$  or  $f_{uni}$ .

13

								PredCls	
		PredCls		-	$f_{uni}$	$\mathbf{b}^{glo}$	R@50/100	$\mathrm{mR}@50/100$	A@50/100
	R@50/100	$\mathrm{mR}@50/100$	A@50/100	-			$65.5 \ / \ 67.3$	18.2 / 19.7	41.9 / 43.5
entire	48.0 / 48.8	$32.6 \ / \ 35.1$	40.3 / 42.0			$\checkmark$	$65.3 \ / \ 67.0$	18.9 / 20.1	$42.1 \ / \ 43.5$
union	55.8 / 57.6	$33.3 \ / \ 35.7$	<b>44.6</b> / <b>46.7</b>		$\checkmark$		$54.7 \ / \ 55.9$	$30.4 \ / \ 33.7$	$42.6 \ / \ 44.8$
					$\checkmark$	$\checkmark$	55.8 / 57.6	$33.3 \ / \ 35.7$	<b>44.6</b> / <b>46.7</b>

ii) The Effect of Global Bias  $\mathbf{b}^{glo}$ : In Sec. 3.2, we utilize the global bias  $\mathbf{b}^{glo}$  in the prediction of sample-specific biases, and the impact of it on the Predcls task of Transformer model is analyzed in Tab. 8,  $f_{uni}$  denotes object pairs' union features. The introduction of global bias  $\mathbf{b}^{glo}$  results in an improvement, as indicated by rows 3 and 4 in the table, showing its effect on improving the accuracy of bias correction. However, we can observe that using  $\mathbf{b}^{glo}$  only is not satisfactory. As it is a set of fixed values, the training has minimal effectiveness. Additionally, the union features  $f_{uni}$  have a crucial impact, as our SBG relies heavily on the object pairs' union features to predict the sample-specific biases. More ablation experiments are in Supplementary Materials.

**Visualization Results.** Fig. 6 presents the results of SGG from Transformer model and our SBG. The left side displays the scene graph generated by Transformer model, while the right side displays the scene graph generated by SBG. The differences in relationships between the two scene graphs are highlighted in the orange font. It is evident that the relationships obtained by Transformer model consist mostly of the coarse-grained relationships, such as "on", "has", and "near". By contrast, our SBG predicts more fine-grained relationships, such as "sitting on", "walking on", and "carrying".

Generalization for Object Detection. We extend our SBP to other tasks with the long-tailed problem, such as object detection in Tab. 9. We experiment



Fig. 6: Comparison of the generated scene graphs between Transformer and SBG.

on the COCO dataset [26], which contains 80 object categories with the severe long-tailed effect. After applying our SBP, the mAP achieves 1.2% improvements, with the tail categories achieving a larger 2.8% improvements.

**Table 9:** Generalization for object detection.  $mAP_{tail}$  is mAP for the last 50 categories of COCO [26] ordered by long tail.

Model (backbone, training schedule)	mAP	$mAP_{tail}$
Faster R-CNN [35] (R-50-FPN, 1x)	36.4	41.6
Faster R-CNN [35] (R-50-FPN, $1x$ ) + SBP	$37.6 \uparrow 1.2\%$	$44.4 \uparrow 2.8\%$

# 5 Conclusion & Future Work

In this paper, to tackle the long-tailed problem, we propose one novel method SBP to conduct sample-level bias correction, and further generate the finegrained scene graph. Specifically, we design a BGAN to predict the samplespecific bias. Extensive experiments on VG, GQA, and VG-1800 datasets validate the effectiveness and generalizability of our SBG. We believe that this work contributes to the advancement of research in this field and offers insights into tackling the long-tailed issue. Our future work aims to further enhance the performance of the proposed method and extend its applicability to other tasks.

## Acknowledgements

This work was partly supported by the National Natural Science Foundation of China (42371321 and 42030102), Natural Science Foundation of HuBei Province (Grant No. 2024AFB283), and Science Foundation of China Three Gorges University (Grant No. 2023RCKJ0022).

15

# References

- Biswas, B.A., Ji, Q.: Probabilistic debiasing of scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10429– 10438 (2023)
- Chen, C., Zhan, Y., Yu, B., Liu, L., Luo, Y., Du, B.: Resistance training using prior bias: toward unbiased scene graph generation. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 212–220 (2022)
- Chen, T., Yu, W., Chen, R., Lin, L.: Knowledge-embedded routing network for scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6163–6171 (2019)
- Chiou, M.J., Ding, H., Yan, H., Wang, C., Zimmermann, R., Feng, J.: Recovering the unbiased scene graphs from the biased ones. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 1581–1590 (2021)
- Dai, B., Zhang, Y., Lin, D.: Detecting visual relationships with deep relational networks. In: Proceedings of the IEEE conference on computer vision and Pattern recognition. pp. 3076–3086 (2017)
- De Boer, P.T., Kroese, D.P., Mannor, S., Rubinstein, R.Y.: A tutorial on the crossentropy method. Annals of operations research 134, 19–67 (2005)
- Deng, Y., Li, Y., Zhang, Y., Xiang, X., Wang, J., Chen, J., Ma, J.: Hierarchical memory learning for fine-grained scene graph generation. In: European Conference on Computer Vision. pp. 266–283. Springer (2022)
- Desai, A., Wu, T.Y., Tripathi, S., Vasconcelos, N.: Learning of visual relations: The devil is in the tails. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15404–15413 (2021)
- Dong, X., Gan, T., Song, X., Wu, J., Cheng, Y., Nie, L.: Stacked hybrid-attention and group collaborative learning for unbiased scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19427–19436 (2022)
- Ghosh, S., Burachas, G., Ray, A., Ziskind, A.: Generating natural language explanations for visual question answering using scene graphs and visual attention. arXiv preprint arXiv:1902.05715 (2019)
- Gkanatsios, N., Pitsikalis, V., Koutras, P., Maragos, P.: Attention-translationrelation network for scalable scene graph generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
- Gu, J., Joty, S., Cai, J., Zhao, H., Yang, X., Wang, G.: Unpaired image captioning via scene graph alignments. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10323–10332 (2019)
- Guo, Y., Gao, L., Wang, X., Hu, Y., Xu, X., Lu, X., Shen, H.T., Song, J.: From general to specific: Informative scene graph generation via balance adjustment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16383–16392 (2021)
- Hudson, D.A., Manning, C.D.: Gqa: A new dataset for real-world visual reasoning and compositional question answering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6700–6709 (2019)
- 15. Keskar, N.S., Socher, R.: Improving generalization performance by switching from adam to sgd. arXiv preprint arXiv:1712.07628 (2017)
- Khandelwal, S., Sigal, L.: Iterative scene graph generation. Advances in Neural Information Processing Systems 35, 24295–24308 (2022)

- 16 Y. Li et al.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. International journal of computer vision 123, 32–73 (2017)
- Li, L., Chen, G., Xiao, J., Yang, Y., Wang, C., Chen, L.: Compositional feature augmentation for unbiased scene graph generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21685–21695 (2023)
- Li, P., Zhang, D., Wulamu, A., Liu, X., Chen, P.: Semantic relation model and dataset for remote sensing scene understanding. ISPRS International Journal of Geo-Information 10(7), 488 (2021)
- Li, R., Zhang, S., He, X.: Sgtr: End-to-end scene graph generation with transformer. In: proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 19486–19496 (2022)
- Li, R., Zhang, S., Wan, B., He, X.: Bipartite graph network with adaptive message passing for unbiased scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11109–11119 (2021)
- Li, W., Zhang, H., Bai, Q., Zhao, G., Jiang, N., Yuan, X.: Ppdl: Predicate probability distribution based loss for unbiased scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19447–19456 (2022)
- 23. Li, Y., Wang, L., Wang, T., Yang, X., Luo, J., Wang, Q., Deng, Y., Wang, W., Sun, X., Li, H., et al.: Scene graph generation in large-size vhr satellite imagery: A large-scale dataset and a context-aware approach. arXiv preprint arXiv:2406.09410 (2024)
- Li, Y., Yang, X., Shang, X., Chua, T.S.: Interventional video relation detection. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 4091–4099 (2021)
- Liang, K., Guo, Y., Chang, H., Chen, X.: Visual relationship detection with deep structural ranking. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision– ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014)
- Lin, X., Ding, C., Zeng, J., Tao, D.: Gps-net: Graph property sensing network for scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3746–3753 (2020)
- Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14. pp. 852–869. Springer (2016)
- Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14. pp. 852–869. Springer (2016)
- Lu, Y., Rai, H., Chang, J., Knyazev, B., Yu, G., Shekhar, S., Taylor, G.W., Volkovs, M.: Context-aware scene graph generation with seq2seq transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 15931– 15941 (2021)

Fine-Grained Scene Graph Generation via Sample-Level Bias Prediction

- 31. Luo, J., Pang, Z., Zhang, Y., Wang, T., Wang, L., Dang, B., Lao, J., Wang, J., Chen, J., Tan, Y., et al.: Skysensegpt: A fine-grained instruction tuning dataset and model for remote sensing vision-language understanding. arXiv preprint arXiv:2406.10100 (2024)
- 32. Lyu, X., Gao, L., Guo, Y., Zhao, Z., Huang, H., Shen, H.T., Song, J.: Fine-grained predicates learning for scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19467–19475 (2022)
- Lyu, X., Gao, L., Zeng, P., Shen, H.T., Song, J.: Adaptive fine-grained predicates learning for scene graph generation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
- Malek, S., Melgani, F., Bazi, Y.: One-dimensional convolutional neural networks for spectroscopic signal regression. Journal of Chemometrics 32(5), e2977 (2018)
- 35. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada. pp. 91-99 (2015), https://proceedings.neurips.cc/paper/2015/hash/ 14bfa6bb14875e45bba028a21ed38046-Abstract.html
- Schroeder, B., Tripathi, S.: Structured query-based image retrieval using scene graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 178–179 (2020)
- Schwing, A.G., Urtasun, R.: Fully connected deep structured networks. arXiv preprint arXiv:1503.02351 (2015)
- Sharifzadeh, S., Baharlou, S.M., Tresp, V.: Classification by attention: Scene graph classification with prior knowledge. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 5025–5033 (2021)
- Suhail, M., Mittal, A., Siddiquie, B., Broaddus, C., Eledath, J., Medioni, G., Sigal, L.: Energy-based learning for scene graph generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13936– 13945 (2021)
- 40. Tang, K., Niu, Y., Huang, J., Shi, J., Zhang, H.: Unbiased scene graph generation from biased training. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3716–3725 (2020)
- 41. Tang, K., Zhang, H., Wu, B., Luo, W., Liu, W.: Learning to compose dynamic tree structures for visual contexts. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6619–6628 (2019)
- 42. Teng, Y., Wang, L.: Structured sparse r-cnn for direct scene graph generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19437–19446 (2022)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems 30 (2017)
- Wu, B., Yu, S., Chen, Z., Tenenbaum, J.B., Gan, C.: Star: A benchmark for situated reasoning in real-world videos. In: Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2) (2021)
- Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1492–1500 (2017)

- 18 Y. Li et al.
- Xu, D., Zhu, Y., Choy, C.B., Fei-Fei, L.: Scene graph generation by iterative message passing. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5410–5419 (2017)
- 47. Yan, S., Shen, C., Jin, Z., Huang, J., Jiang, R., Chen, Y., Hua, X.S.: Pcpl: Predicate-correlation perception learning for unbiased scene graph generation. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 265– 273 (2020)
- Yang, X., Tang, K., Zhang, H., Cai, J.: Auto-encoding scene graphs for image captioning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10685–10694 (2019)
- 49. Yu, J., Chai, Y., Wang, Y., Hu, Y., Wu, Q.: Cogtree: Cognition tree loss for unbiased scene graph generation. arXiv preprint arXiv:2009.07526 (2020)
- Zareian, A., Karaman, S., Chang, S.F.: Bridging knowledge graphs to generate scene graphs. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16. pp. 606–623. Springer (2020)
- Zellers, R., Yatskar, M., Thomson, S., Choi, Y.: Neural motifs: Scene graph parsing with global context. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5831–5840 (2018)
- Zeng, P., Gao, L., Lyu, X., Jing, S., Song, J.: Conceptual and syntactical crossmodal alignment with cross-level consistency for image-text matching. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 2205–2213 (2021)
- 53. Zhang, A., Yao, Y., Chen, Q., Ji, W., Liu, Z., Sun, M., Chua, T.S.: Fine-grained scene graph generation with data transfer. In: European conference on computer vision. pp. 409–424. Springer (2022)
- Zou, F., Shen, L., Jie, Z., Zhang, W., Liu, W.: A sufficient condition for convergences of adam and rmsprop. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition. pp. 11127–11135 (2019)