

This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

From General to Specific: Informative Scene Graph Generation via Balance Adjustment

Yuyu Guo¹

Lianli Gao^{1*} Xuanhan Wang¹

Heng Tao Shen¹

Yuxuan Hu² Jingkuan Song¹ Xu Lu³

¹Center for Future Media & School of Computer Science and Engineering, University of Electronic Science and Technology of China, China

²Southwest University, China

³Kuaishou, China

Abstract

The scene graph generation (SGG) task aims to detect visual relationship triplets, i.e., subject, predicate, object, in an image, providing a structural vision layout for scene understanding. However, current models are stuck in common predicates, e.g., "on" and "at", rather than informative ones, e.g., "standing on" and "looking at", resulting in the loss of precise information and overall performance. If a model only uses "stone on road" rather than "blocking" to describe an image, it is easy to misunderstand the scene. We argue that this phenomenon is caused by two key imbalances between informative predicates and common ones, i.e., semantic space level imbalance and training sample level imbalance. To tackle this problem, we propose BA-SGG, a simple vet effective SGG framework based on balance adjustment but not the conventional distribution fitting. It integrates two components: Semantic Adjustment (SA) and Balanced Predicate Learning (BPL), respectively for adjusting these imbalances. Benefited from the model-agnostic process, our method is easily applied to the state-of-the-art SGG models and significantly improves the SGG performance. Our method achieves 14.3%, 8.0%, and 6.1% higher Mean Recall (mR) than that of the Transformer model at three scene graph generation sub-tasks on *Visual Genome, respectively. Codes are publicly available*¹.

1. Introduction

Scene Graph Generation (SGG) aims to detect instancesof-interest and their relationships in an image. It provides a structural vision layout, as an auxiliary tool, to bridge the

*Corresponding author.

¹https://github.com/ZhuGeKongKong/SGG-G2S



Xing Xu¹

Figure 1. A confusion matrix of a baseline model (a) and two searching results for "stone on road" in Google (b). The element $c_{i,j}$ means the number of samples labeled as *i* but predicted as *j*. The results of the model concentrate on the left side, *e.g.*, "on" and "near". If the model directly uses the common predicate "on" to describe the first image like the second image in (b), it may cause serious consequences.

gap from computer vision to natural language, supporting many high-level tasks such as visual captioning [35, 1] and visual question answering [27, 15].

A good scene graph can provide sufficient and informa-



Figure 2. Ground-truth annotations from Visual Genome. The first and second examples are marked as "on" because of the huge semantic space of "on". Compared with "person standing on snow" in the third example, "person on snow" in the second example is ambiguous and information-poor.

tive relationships among instance-of-interest. Inspired by remarkable progress in object detection [6, 5, 24], existing solutions [10, 34, 16, 36] mostly follow a common generation paradigm, that is, detecting objects from an image, extracting region features, and then recognizing the predicate categories under the guidance of standard classification objective function. Based on this paradigm, current stateof-the-art methods, however, exhibit the tendency that most of the recognized predicates are general or common (e.g., "on" or "in") and lack specific or informative contents (e.g., "standing on" or "sitting on"), as shown in Fig. 1(a). Thus, it is hard to apply an SGG algorithm in real-world settings like Fig. 1(b), since it provides insufficient clues and may lead to misunderstanding about a scene. In this work, our study reveals that such biased predicates are caused by imbalance issues, preventing the power of the well-designed SGG model from being exploited. Specifically, this issue can further be divided into two sub-problems: semantic space level imbalance and training sample level imbalance. Semantic space level imbalance. Common predicates, such as "on", have large semantic spaces, while informative predicates like "standing on" provide rich content but have small semantic spaces and even may be replaced by common ones in some annotations. Some examples of this semantic space level imbalance are shown in Fig. 2. All examples in Fig. 2 reflect the fact of "on", while only the last two examples represent identical contents of "standing on". Although "standing on" is more specific for the last two examples, humans still label the second image with "on" because of its larger semantic space than "standing on". The semantic space level imbalance also reflects the semantic relationship between predicates, e.g., annotators prefer to label "standing on" as "on" rather than "has". Tagging different labels for identical contents confuses the SGG model and causes poor performance.

Training sample level imbalance. When learning to recognize predicates, informative predicate samples are particularly valuable as they provide precise knowledge to a scene graph. However, predicate samples are dominated by the common predicate categories in the SGG dataset (*e.g.*, Visual Genome). For example, Fig. 3 shows the annotation distribution over predicate categories. We can observe that the number of common predicate samples is much larger



Figure 3. Frequencies of predicates in the Visual Genome dataset. For clarity, we omit some predicates.

than that of informative ones, which leads to the problem of the long-tailed distribution of the classes. Within such an imbalanced sampling space, the prediction of informative predicates is dominated by the common ones.

The above challenges motivate us to study two problems: 1) how to revise common predicates as informative ones based on the semantic relation. And 2) how to make training sample space balanced. To tackle these problems, we believe that the key is to explore the predicate relation in semantic space and adjust sampling space to be balanced. Motivated by this, we propose a simple yet effective pipeline, namely, *Scene Graph Generation with Balance Adjustment* (BA-SGG), to learn informative scene graph generation. This pipeline integrates two novel components: 1) Semantic Adjustment (SA), which casts common predictions as informative ones with rich information contents by fully exploiting semantic relation among predicates. 2) Balanced Predicate Learning (BPL), which mines informative predicates according to their information contents.

To sum up, the main contributions of our work contain: 1). We systematically review the process of scene graph generation and reveal the problem of insufficient information contents that limit SGG's overall performance and practical applicability, i.e., semantic space level imbalance and training sample level imbalance; 2). We propose the BA-SGG for informative scene graph generation, a novel framework that adjusts biased predicate predictions by two components: semantic adjustment (SA) and balanced predicate learning (BPL); 3). The proposed BA-SGG is tested on Visual Genome, obtaining significant improvements over state-of-the-art SGG approaches. Without bells and whistles, BA-SGG achieves 14.3%, 8.0%, and 6.1% higher Mean Recall (mR) than Transformer [31, 28] on three scene graph generation sub-tasks, respectively. Besides, we propose a new metric (mRIC@K) to measure the information content of results.

2. Related Work

Scene Graph Generation. Many methods [10, 34, 16, 36] were proposed to handle the scene graph generation task in recent years. These methods deal with the task from differ-



Figure 4. The main process of previous methods from the input image to the scene graph.

ent perspectives, such as extracting context information by message passing [34, 38, 22, 20], constructing visual embedding in the semantic space [39, 11, 23], improving robustness by external knowledge [7, 8, 21]. Except for these models trapped in common predicates, Liang [18] reconstructed a dataset that focuses on visually-relevant relationships. The unbiased SGG is proposed in [29] to remove the vision-agnostic bias with counterfactual causality. Different from these works, our method explores the imbalance in the semantic and learning space adequately to learn scene graphs with precise and rich information.

Informative Prediction. In order to get out of common prediction, many works have been proposed in computer vision, such as fine-grained recognition [3, 14, 32] and informative image captioning [4, 40]. Most of them focus on the informative object categories, which are either precisely represented by samples or constructed with distinct hierarchies. Unlike object categories, it is difficult to define a clear hierarchical or coarse-to-fine structure for relationship predicates. Instead of seeking a way to define a structure among predicates manually, we estimate a semantic relation from a baseline model and define the predicate information content to learn informative predicates.

3. Approach

3.1. Approach Overview

Problem Formulation. As shown in Fig. 4, conventional methods formulate scene graph generation as a two-stage process, where it firstly detects all instances-of-interest and then recognizes relationship predicates between pair-wise instances. Given an image X, its corresponding scene graph is generated from the complete graph $\mathbf{G} = \langle \mathbf{O}, \mathbf{R} \rangle$, where **O** is the set of instance nodes. **R** is the full set of edges, each of which connects two nodes and encodes a relationship between them. More specifically, each edge is represented in the form of a triplet (o_i, y_{ij}, o_j) , where o_i denotes the subject, o_j is the object, and y_{ij} represents the predicate of the edge/triplet. We denote by **Y** the list



Figure 5. The whole training process of our method, *i.e.*, Semantic Adjustment (SA) and Balanced Predicate Learning (BPL).

of predicate labels for the edges **R**, hence $|\mathbf{Y}| = |\mathbf{R}| = |\mathbf{O} \times \mathbf{O}|$, where $|\cdot|$ represents the length/size of a list/set. The pair of (o_i, o_j) is embodied as region bounding boxes with object categories, which are usually obtained by an object detector such as Faster R-CNN [25]. Thus, SGG models omit the process of object detection and only focus on predicting relationship predicates **Y** between instance pairs, by maximizing the probability $\Pr(\mathbf{Y}|\mathbf{O} \times \mathbf{O}; \theta)$, where θ is the learnable parameter of the scene graph generation model.

Overview. We follow this pipeline approach and focus on the second stage, *i.e.*, predicate recognition, after instance regions are detected by an object detection system. In particular, the proposed method consists of two processes, as shown in Fig. 5. 1) Semantic Adjustment (SA) is applied to casting common predictions generated by an SGG model as informative ones, where it performs a relation modeling among predicates. 2) Balanced Predicate Learning (BPL) is designed for extending the sampling space for informative predicates. In the following subsections, we present details of each component.

3.2. Semantic Adjustment

As we discussed earlier, most of the current SGG models bias toward the common predicate classes. The Semantic Adjustment component is designed to restore informative predictions from common ones, based on the semantic relationship between predicates. Intuitively, "standing on" is more likely to be predicted as "on" than "has" because the meanings of "standing on" and "on" are closer than between "standing on" and "has". Therefore, we exploit such relationships to adjust the prediction results. In particular, we formalize this adjustment as Eq. 1:

$$\Pr(y_{ij}^s|o_i, o_j; \theta) = \Pr(y_{ij}^s|y_{ij}^g) \Pr(y_{ij}^g|o_i, o_j; \theta) , \qquad (1)$$

where $\Pr(y_{ij}^g | o_i, o_j; \theta) \in \mathbb{R}^K$ is the prediction of an SGG model for all predicate categories between subject o_i and object o_j , $\Pr(y_{ij}^s | o_i, o_j; \theta) \in \mathbb{R}^K$ is produced by the SGG model after semantic adjustment, and K is the total number of predicate categories. The superscripts g and s are used to

distinguish the prediction before and after SA. $\Pr(y_{ij}^s | y_{ij}^g)$ represents the semantic adjustment, which measures how confident a common result is transformed into an informative one.

The semantic adjustment $\Pr(y_{ij}^s | y_{ij}^g)$ reflects the latent semantic relationship between predicates. It can be considered as a transition matrix between pairs of predicate labels, thus is of dimension $K \times K$. For simplicity, we denote $\Pr(y_{ij}^s | y_{ij}^g)$ as a transition matrix $C^* \in \mathbb{R}^{K \times K}$, and it can be derived from prediction priors. We denote by $C \in \mathbb{R}^{K \times K}$ the prediction confusion matrix, where each element c_{kl} denotes the number of samples labeled as the k-th category but being predicted as l. In this paper, we use the frequency model [38] to generate predictions on the training dataset for C, and fix it in our framework. Next, a normalized transition matrix C', in which each element c'_{kl} is defined in Eq. 2 as follows:

$$c_{kl}' = \frac{c_{kl}}{\sum\limits_{m=1}^{K} c_{km}} \,. \tag{2}$$

The visualization of the row-normalized confusion matrix is shown in the heatmap in Fig. 1(a). We can observe that large values of the matrix concentrate on the left side. This implies that many informative predicates are classified as a few common ones. The matrix also implies the semantic relationship between predicates to some extent, *e.g.*, "parked on" is more likely to be predicted "on" than "has". However, if an informative predicate has been predicted, its probability should not be adjusted significantly. Directly multiplying C' may greatly decrease the score of an informative predicate because the diagonal tail values are small. Therefore, the diagonal elements' values should be increased to alleviate this impact. To achieve this goal, we add a weighted identity matrix to C', resulting in the final transition matrix C^* as formed by Eq. 3:

$$C^{\star} = Row_Normalize(C' + \alpha I_K), \qquad (3)$$

where $I_K \in \mathbb{R}^{K \times K}$ is the identity matrix and α is a hyperparameter. Normalizing the matrix by row ensures that the sum of the row is 1. We freeze the transition probability matrix during training to avoid semantic drifting.

3.3. Balanced Predicate Learning

To address imbalanced training, a new data source with a balanced sample space is needed. Formally, we view a data source with an imbalanced sample space as a source domain, and a data source with a relatively balanced sample space as a target domain. Unlike existing methods that train the SGG model in the source domain only, balanced predicate learning (BPL) divides the learning process into three stages: 1). Training SGG model on a source domain; 2). Creating a target domain; 3). Transfer learning on the target domain. Fig. 5 shows the overview of this learning process.

In the first stage, we take a common training strategy that is identical to that of previous works [38, 34] to train an SSG model on the source domain.

In the second stage, a target domain with a relatively balanced sample space is constructed. However, it is hard to decide whether a predicate is informative or not, since annotations related to the predicate property are not provided. Intuitively, common predicates are semantically general, while informative predicates are context-specific. Therefore, that discriminative problem can be transformed into a process of measurement, where it measures how informative an event is, in other words, how much information content does an event involves. According to Shannon information theory, the information content I of an event zcan be assessed by Eq. 4:

$$I(z) = -\log_b[\Pr(z)], \qquad (4)$$

where Pr(z) denotes the probability of the event z. It reflects the fact that an event with a small probability provides more information content. It is noteworthy that other ways of calculating 'informativeness', e.g., semantic context based method [17] or statistical information in corpora by TF-IDF [26], can also be used here. We calculate 'informativeness' based on Shannon information theory mainly for simplicity and general cases. Inspired by this, we measure informativeness for each predicate category by estimating its information content, using Eq. 4. The probability of a predicate is estimated by its frequency in an information source. An example of this can be shown in Fig. 3, where predicate "on" has higher occurrence frequency but fewer information contents when compared with "standing on". Based on a specific information source (e.g., Visual Genome), the frequency for each predicate category is used to calculate the information content. Next, we sort all predicates by their information content, resulting in ascending order. Top-M predicates are chosen as common ones, and the rest are informative predicates. Next, we create a target domain by sampling labeled data from the source domain. A direct way is the balance sampling strategy that increases the sampling frequency of rare categories. However, such upsampling strategy is ineffective for the scene graph generation task, as demonstrated in previous work [29]. Therefore, we take a separation undersampling strategy to create a target domain. In particular, we move all samples that belong to the informative predicate category in the source domain to the target domain. As for categories belonging to Mcommon predicates, we randomly sample N labeled relationship triplet samples from source domains for each predicate, resulting in a balanced sample space between common predicates and informative ones. Besides balancing the dataset, randomly removing samples of common predicates may also have another benefit: because there are many ambiguous samples in common predicates that disturb informative predicates, as shown in Fig. 2, reducing samples of common predicates may help to reduce the impact of these samples to a certain extent. To eliminate the impact of ambiguous samples further, we also introduce another undersampling method, which uses the pre-trained SGG model in the first stage to find ambiguous samples of common predicates and remove them in the target domain. Details and Experiments can be found in Sec. 4.4.

After target domain construction, we can train an SGG model on it. Instead of training an SGG model from scratch on the target domain, we first initialize an SGG model from the one pretrained in the first stage. And then, only the last recognition layer of the SGG model is finetuned on the target domain. There are two reasons behind this: First, finetuning the whole model requires high computation costs. Second, training the whole model on the target domain with fewer labeled samples increases the risk of the overfitting problem, leading to the failure of recognizing common predicates. More details can be seen in Sec. 4.4.

4. Experiments

4.1. Experimental Settings

Datasets. Following previous work [2, 38, 29], we conduct experiments on a widely used SGG dataset, namely Visual Genome (VG) [13]. It is composed of 108k images across 75k object categories together with 37k predicate classes. Since 92% of the predicates have few samples, we followed previous works [34] and adopted a new VG split, containing the most frequent 150 object categories and 50 predicate classes. Moreover, the VG dataset is split into a training set (70%) and a test set (30%), and we further sample a validation set (5k) from the training set for model validation.

Model Configuration. In this work, we evaluate our method based on three baselines: **MotifNet** [38], **VC-Tree** [30] and **Transformer** [31, 28, 12, 20]. The numbers of inner layers in the object encoder and the edge encoder are set to 4 and 2, respectively. Other hyperparameters are identical to the setting in **Model Zoo** [28]. All models share the same settings and the pre-trained detector as well.

Metrics. Following previous works [29, 30, 2], we evaluate the SGG method on three subtasks: 1) Predicate Classification (PredCls), 2) Scene Graph Classification (SGCls), and 3) Scene Graph Detection (SGDet). The PredCls takes ground truth bounding boxes and labels as inputs. The SG-Cls takes ground truth object bounding boxes as inputs, but without labels. The SGDet requires predicting relationships from scratch. In this work, we use three evaluation metrics: R@K, mR@K, and mRIC@K. In detail, R@K averages the recall for all samples, while mR@K averages the re-

call across predicate categories. R@K underestimates informative predicate categories but focuses on common predicates with rich samples. Due to the defect existing in R@K measurement reported in [29], we mainly report the mR@K metric, where it averages R@K for all predicate categories.

Moreover, we propose a new metric, mRIC@K, that measures the mean recall with information content. Specifically, for each predicate, its recall with information content, denoted RIC, is calculated by multiplying its recall with the predicates information content, which is calculated by Eq. 4. mRIC@K, finally, can be obtained by averaging the RIC values across all predicate categories. This metric reflects how much information a scene graph can provide. A higher score implies that the generated scene graph involves more informative content and vice versa. Moreover, we use two information sources to estimate information content for each predicate, *i.e.*, Visual Genome and Wikipedia. We separately denote mRIC (VG) and mRIC (Wiki) as the information content from Visual Genome and Wikipedia.

4.2. Implementation Details

We utilize Faster R-CNN [25] with ResNeXt-101-FPN [9, 19, 33] pre-trained by [29] to detect instances in images and freeze the weights during learning scene graph generation. For scene graph generation, we firstly train all SGG models on the source domain according to the recommended configuration [28] for all tasks, including the learning rate and batch size. Then each model is finetuned on the target domain with the same configuration. We take the logits before the softmax layer as $\Pr(y_{ij}^g|o_i, o_j; \theta)$ in Eq. 1 for the sake of simplicity. For constructing the target domain, M predicates with the smallest information content are set as common predicates, where M is 15 in this paper. Moreover, the sampling number N is 2k. We use the information content only from Visual Genome to distinguish common predicates and informative ones unless otherwise stated.

4.3. Exploration Study

In this section, we investigate the superiority of the proposed method from three aspects: 1). performance of scene graph generation; 2). information content of predicate prediction; and 3). accuracy of informative predicates.

For scene graph generation, we investigate the generalization capability of the proposed method by adding it to three baseline models as mentioned above. The experimental results are summarized in Tab. 1. As can be seen, all baseline models are improved by a large margin for all metrics. In particular, Transformer trained with BPL learning strategy outperforms the Transformer trained in source domain by 12.1. It is further improved (24.5 *vs.* 26.7) when applying SA to cast common predicates as informative ones. Similar gains are obtained when applying our method to MotifNet and VCTree, indicating the generaliz-

Mathad		PredCls		SGCls			SGDet		
Metilod	mR@20	mR@50	mR@100	mR@20	mR@50	mR@100	mR@20	mR@50	mR@100
Transformer	12.4	16.0	17.5	7.7	9.6	10.2	5.3	7.3	8.8
Transformer+BPL	$24.5^{+12.1}$	$29.4^{+13.4}$	31.7 ^{+14.2}	$14.1^{+6.4}$	16.8 ^{+7.2}	$17.8^{+7.6}$	$10.2^{+4.9}$	$13.2^{+5.9}$	$15.4^{+6.6}$
Transformer+BPL+SA	$26.7^{+14.3}$	$31.9^{+15.9}$	$34.2^{+16.7}$	$15.7^{+8.0}$	18.5 ^{+8.9}	$19.4^{+9.2}$	$11.4^{+6.1}$	$14.8^{+7.5}$	$17.1^{+8.3}$
MotifNet	11.5	14.6	15.8	6.5	8.0	8.5	4.1	5.5	6.8
MotifNet+BPL	22.6 ^{+11.1}	$27.1^{+12.5}$	$29.1^{+13.3}$	$13.0^{+6.5}$	$15.3^{+7.3}$	$16.2^{+7.7}$	$9.7^{+5.6}$	$12.4^{+6.9}$	$14.4^{+7.6}$
MotifNet+BPL+SA	$24.8^{+13.3}$	$29.7^{+15.1}$	$31.7^{+15.9}$	$14.0^{+7.5}$	$16.5^{+8.5}$	$17.5^{+9.0}$	$10.7^{+6.6}$	13.5 ^{+8.0}	15.6 ^{+8.8}
VCTree	11.7	14.9	16.1	6.2	7.5	7.9	4.2	5.7	6.9
VCTree+BPL	$23.8^{+12.1}$	$28.4^{+13.5}$	$30.4^{+14.3}$	15.6 ^{+9.4}	$18.4^{+10.9}$	$19.5^{+11.6}$	$9.9^{+5.7}$	$12.5^{+6.8}$	$14.4^{+7.5}$
VCTree+BPL+SA	$26.2^{+14.5}$	$30.6^{+15.7}$	$32.6^{+16.5}$	17.2 ^{+11.0}	$20.1^{+12.6}$	$21.2^{+13.3}$	$10.6^{+6.4}$	13.5 ^{+7.8}	15.7 ^{+8.8}

Table 1. Ablation studies on generalizability and effectiveness of our proposed components, *i.e.*, Balanced Predicate Learning (BPL) and Semantic Adjustment (SA). The proposed method is applied on three baseline models, where the MotifNet and VCTree are reimplemented by [29].

Metric	Method	mRIC@20	mRIC@50	mRIC@100
	Transformer	43.5	59.2	65.8
mRIC (VG)	Transformer+BPL	118.2	142.6	154.5
	Transformer+BPL+SA	134.5	160.3	172.5
	Transformer	74.1	96.6	105.9
mRIC (Wiki)	Transformer+BPL	182.1	216.3	231.3
	Transformer+BPL+SA	204.6	239.2	254.4

Table 2. Information content assessment on the PredCls task. Models are evaluated on two information sources (Visual Genome & Wikipedia).



Figure 6. Confusion matrices of Transformer and Transformer (BPL+SA). The coordinates of these matrices are arranged in the increasing order of predicate information content from left to right and top to bottom.

ability of the proposed method.

To evaluate whether the BA-SGG algorithm can provide more informative content, we respectively report evaluation results on mRIC@K for Transformer and the proposed method. The evaluation results based on two information sources are summarized in Tab. 2. From the results, we can observe that the proposed methods obtain higher scores when compared with Transformer, which implies that the proposed method generates more informative scene graphs.

In order to intuitively show the accuracy of informative predicates, we visualize confusion matrices across predicate categories. In particular, we compare two confusion matrices generated from Transformer and our method in Fig. 6. Not surprisingly, the proposed method accurately detects more informative predicates than Transformer's since the confusion matrix generated from the proposed method is brighter on the diagonal than that of Transformer. Besides, Transformer (BPL+SA) gets rid of common predicates because its bright elements no longer concentrate on the left, compared with the confusion matrix of Transformer.

4.4. Ablation Study on Details of Proposed Method

Except for the effectiveness of our method, we also investigate variants of our method for more insights. Next, we mainly study two variants of the proposed method: variants to SA and variants to BPL. Furthermore, we use Transformer as the baseline in the following experiments.

Variants to SA. For SA, we investigate variants of the transition matrix, *i.e.*, C^* in Eq. 3. In particular, we compare four settings: 1) Transformer model with BPL but without SA, where it is the baseline. 2) C^* is implemented by a random matrix and updated during training. 3) C^* is calculated by Eq. 3 and updated during training. 4) C^* is obtained by Eq. 3 while it is fixed during training. The experimental results are shown in Tab. 4. From the results, we have the following observations: first, the transition matrix initialized from a random matrix and updated during training brings minor gains in terms of mean recall score. Second, the fixed transition matrix obtained from prediction priors (*i.e.*, Eq. 3) performs the best. This implies that updating the transition matrix during the training hurts the quality of the common-to-informative process.

In Eq. 3, we add an identity matrix and introduce a hyper-parameter α to investigate the importance of the identity matrix. The adjustment of α is shown in Tab. 5. From Tab. 5, we can find that the model not only needs to add the identity matrix but also needs to set α equal to 1 to get better results.

Variants to BPL. As illustrated in Sec. 3.3, the proposed BPL learning process consists of three steps, where the construction of the target domain is the essential part. To create the target domain, an information source is needed to calculate information content (IC) for each predicate category. Based on their IC values, predicates are further categorized as common or informative. Here, we use two different corpus sources for IC calculation: 1) Visual Genome, where it provides triplets for visual contents; and 2) Wikipedia, where it provides the latest articles (about 110k). Notably, Wikipedia is a language-based corpus source without visual

Target Domain	SGG (PredCls)			mRIC (VG)			mRIC (Wiki)			
Target Domain	mR@20	mR@50	mR@100	mRIC@20	mRIC@50	mRIC@100	mRIC@20	mRIC@50	mRIC@100	
General	12.4	16.0	17.5	43.5	59.2	65.8	74.1	96.6	105.9	
Wikipedia	21.3	26.0	28.0	99.5	121.9	131.5	166.1	198.4	211.7	
Visual Genome	26.7	31.9	34.2	134.5	160.3	172.5	204.6	239.2	254.4	

Table 3. Ablation studies of target domain construction in BPL.

Settings	C^{\star}	finetune	mR@20	mR@50	mR@100
1	-	-	24.5	29.4	31.7
2	Random Init	√	25.1	29.5	31.3
3	SA	√	25.8	30.2	32.1
4	SA	-	26.7	31.9	34.2

Table 4. Ablation studies of transition matrix on PredCls.

α	mR@20	mR@50	mR@100
$\alpha = 0.0$	20.5	25.4	27.3
$\alpha = 0.3$	25.8	30.5	32.8
$\alpha = 0.6$	26.4	31.6	33.9
$\alpha = 1.0$	26.7	31.9	34.2

Table 5. Ablation studies of α in Eq. 3 on PredCls.

Undersampling	PredCls	SGCls	SGDet
Random	26.7	15.7	11.4
Pre-trained Model	28.5	16.4	11.4

Table 6. Ablation studies of undersampling methods on mR@20 of the three subtasks.

Settings	pretrained	f-backbone	R@20	R@50	R@100
1	-	-	28.4	34.9	37.2
2	\checkmark	\checkmark	30.6	37.3	39.5
3	\checkmark	-	49.0	55.7	57.6

Table 7. Ablation studies of training approaches on PredCls.

contents. Next, two target domains can be generated based on these two information sources. Then, we perform diagnostic experiments to investigate whether the SSG model can benefit from the target domain with a balanced sample space. In particular, we compare three settings: 1) Transformer model trained on the source domain. 2) Transformer model trained on target domain based on Wikipedia. 3) Transformer model trained on target domain based on Visual Genome. The experimental results are summarized in Tab. 3. We have the following findings: first, SGG models trained on the target domain, either language-based or visual-based, outperform the one trained on the source domain by a large margin (at least 8.9 gains for mR@20). This indicates that the balanced sample space plays an important role in training the SGG model. Besides, the SGG model is improved after being trained on the language-based target domain, demonstrating the effectiveness and generalizability of the proposed IC-based predicate partition. As for the assessment of information contents, models trained on the target domain achieve higher scores than the model trained on source domain, either based on VG information source (134.5 vs. 43.5 and 99.5 vs. 43.5) or Wiki information source (204.6 vs. 74.1, and 166.1 vs. 74.1). This demonstrates that the scene graph generated from the models trained on the target domain provides more informative content for better scene understanding.

To further eliminate the impact of ambiguous samples in common predicates, as shown in Fig. 2, we design another undersampling method. Intuitively, the ambiguous sample in Fig. 2 can be annotated as either a common predicate "on" or an informative predicate "standing on". In other words, the confidence of the annotation "on" should be less. Therefore, we use the confidence of the pre-trained SGG model in the first stage to remove ambiguous samples. Specifically, we input all training samples to the model and keep the top-2k samples with the highest confidence for each common predicate in the target domain. As shown in Tab. 6, although using the pre-trained model for undersampling improves the performances on PredCls, this improvement does not achieve the desired effect, especially on SGDet. There may be two reasons for the insignificant improvement: 1). Random undersampling by a large margin (e.g., "on" from 70k to 2k) can reduce the influence of ambiguous samples to a great extent. 2). Using the pre-trained model to find ambiguous samples is not very reliable, and another effective method should be introduced. Therefore, this problem should still be valued in future work, and we use the random undersampling method in this paper for simplicity and effectiveness.

After the target domain construction, a good training approach is needed. To explore this aspect, we perform an additional experiment as illustrated in Tab. 7, where it compares three different training settings: 1) Training an SGG model from scratch on the target domain. 2) Finetuning a whole SGG model on the target domain, where it is pretrained on the source domain. 3) Finetuning the classifier layer of an SGG model on the target domain, where it is pretrained on the source domain. From Tab. 7, we can observe that the third setting provides the best result, whilst directly training an SGG model from scratch on the adjusted domain gives the worst result. Furthermore, the model trained with the third setting outperforms the one trained with the second setting. One possible reason is that the scale of sample space in the target domain (about 55k) is much smaller than the source domain's (about 400k), posing high risks of overfitting. Since this experiment focuses on the training approach, we remove SA in this experiment.

4.5. Comparison with State-of-the-art Methods

After verifying the effect of the proposed method, we compare it with state-of-the-art methods. The comparison results are shown in Tab. 8. As shown in the results, the proposed method achieves the best performance in all

Method		PredCls		SGCls			SGDet		
Wiethou	mR@20	mR@50	mR@100	mR@20	mR@50	mR@100	mR@20	mR@50	mR@100
IMP+ [34, 2]	-	9.8	10.5	-	5.8	6.0	-	3.8	4.8
FREQ [38, 30]	8.3	13.0	16.0	5.1	7.2	8.5	4.5	6.1	7.1
KERN [2]	-	17.7	19.2	-	9.4	10.0	-	6.4	7.3
GPS-Net [20]	-	-	22.8	-	-	12.6	-	-	9.8
GB-Net [37]	-	22.1	24.0	-	12.7	13.4	-	7.1	8.5
MotifNet (Focal) [38, 29]	10.9	13.9	15.0	6.3	7.7	8.3	3.9	5.3	6.6
MotifNet (Reweight) [38, 29]	16.0	20.0	21.9	8.4	10.1	10.9	6.5	8.4	9.8
MotifNet (Resample) [38, 29]	14.7	18.5	20.0	9.1	11.0	11.8	5.9	8.4	9.8
MotifNet (TDE) [38, 29]	18.5	24.9	28.3	11.1	13.9	15.2	6.6	8.5	9.9
VCTree (TDE) [30, 29]	18.4	25.4	28.7	8.9	12.2	14.0	6.9	9.3	11.1
VTransE (TDE) [39, 29]	18.9	25.3	28.4	9.8	13.1	14.7	6.0	8.5	10.2
Transformer (BA-SGG)	26.7	31.9	34.2	15.7	18.5	19.4	11.4	14.8	17.1
MotifNet (BA-SGG)	24.8	29.7	31.7	14.0	16.5	17.5	10.7	13.5	15.6
VCTree (BA-SGG)	26.2	30.6	32.6	17.2	20.1	21.2	10.6	13.5	15.7

Table 8. Comparison between our method (BA-SGG) and previous methods.



Figure 7. Visualization results of Transformer and Transformer (BA-SGG) on the PredCls task. The generated scene graph from the Transformer (BA-SGG) is more informative than the one from Transformer. Only top 30% relationships on each image are shown for clarity.

the measure metrics among all the comparison methods, reaching 26.7 mR@20 for PredCls, 15.7 mR@20 for SG-Cls, and 11.4 mR@20 for SGDet, respectively. For Pred-Cls, it outperforms the best competitor TDE by 6.3 for MotifNet and 7.8 for VCTree on mR@20. Similar gains also are observed for another two tasks. Compared with the resample or reweight methods, *e.g.*, MotifNet (Focal), MotifNet (Reweight), and MotifNet (Resample), our method also outperforms them because they do not simultaneously consider the two imbalances, *i.e.*, semantic space level imbalance and training sample level imbalance. This clearly indicates the superiority of our method and the importance of the two imbalances.

4.6. Visualization Results

Qualitative results across various scenes are shown in Fig. 7. Here, we compare Transformer with Transformer equipped with the proposed method (BA-SGG). From Fig. 7, Transformer (BA-SGG) generates more informative scene graphs than vanilla Transformer, such as ({motorcycle, parked on, street} vs. {motorcycle, on, street}) in the first example, ({tree, growing on, hill} vs. {tree, on, hill}) in the second example, and ({person, riding, skateboard} vs. {person, on, skateboard}) in the third example. This clearly demonstrates the success of predicates adjustment by the proposed method. However, the fourth example shows both models fail to recognize the predicate between two zebras since both zebras can be viewed as subjects. This implies that recognition of the subject-subject interaction relationship is still a bottleneck for SGG.

5. Conclusion

In this work, balance adjustment scene graph generation (BA-SGG) is proposed for learning scene graphs from general to specific. We first reveal the imbalance between common predicates and informative predicates. Then two strategies are designed: Semantic Adjustment and Balanced Predicate Learning, in view of the imbalance. The semantic adjustment strategy explores the semantic relation between predicates, and the balanced predicate learning strategy transfers the knowledge learned from general predicates to informative ones. Finally, numerous experiments show that our method significantly improves the performance of scene graph generation.

Acknowledgement. This work is funded by Kuaishou.

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- [2] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *CVPR*, 2019.
- [3] Yue Chen, Yalong Bai, Wei Zhang, and Tao Mei. Destruction and construction learning for fine-grained image recognition. In *CVPR*, 2019.
- [4] Reuben Cohn-Gordon, Noah D. Goodman, and Christopher Potts. Pragmatically informative image captioning with character-level inference. In NAACL, 2018.
- [5] Ross B. Girshick. Fast R-CNN. In ICCV, 2015.
- [6] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Region-based convolutional networks for accurate object detection and segmentation. *PAMI*, 38(1), 2016.
- [7] Jiuxiang Gu, Handong Zhao, Zhe Lin, Sheng Li, Jianfei Cai, and Mingyang Ling. Scene graph generation with external knowledge and image reconstruction. In *CVPR*, 2019.
- [8] Yuyu Guo, Jingkuan Song, Lianli Gao, and Heng Tao Shen. One-shot scene graph generation. In ACMMM, 2020.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [10] Roei Herzig, Moshiko Raboh, Gal Chechik, Jonathan Berant, and Amir Globerson. Mapping images to scene graphs with permutation-invariant structured prediction. In *NeurIPS*, 2018.
- [11] Zih-Siou Hung, Arun Mallya, and Svetlana Lazebnik. Contextual translation embedding for visual relationship detection and scene graph generation. *PAMI*, 2020.
- [12] Rajat Koner, Poulami Sinhamahapatra, and Volker Tresp. Relation transformer network. *CoRR*, abs/2004.06193, 2020.
- [13] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123(1), 2017.
- [14] Michael Lam, Behrooz Mahasseni, and Sinisa Todorovic. Fine-grained recognition as hsnet search for informative image parts. In CVPR, 2017.
- [15] Xiangpeng Li, Jingkuan Song, Lianli Gao, Xianglong Liu, Wenbing Huang, Xiangnan He, and Chuang Gan. Beyond rnns: Positional self-attention with co-attention for video question answering. In AAAI, volume 33, 2019.
- [16] 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 *ECCV*, 2018.
- [17] Chen Liang, Zhaohui Wu, Wenyi Huang, and C. Lee Giles. Measuring prerequisite relations among concepts. In *EMNLP*, 2015.
- [18] Yuanzhi Liang, Yalong Bai, Wei Zhang, Xueming Qian, Li Zhu, and Tao Mei. Vrr-vg: Refocusing visually-relevant relationships. In *ICCV*, 2019.

- [19] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In CVPR, 2017.
- [20] Xin Lin, Changxing Ding, Jinquan Zeng, and Dacheng Tao. Gps-net: Graph property sensing network for scene graph generation. In *CVPR*, 2020.
- [21] Cewu Lu, Ranjay Krishna, Michael S. Bernstein, and Fei-Fei Li. Visual relationship detection with language priors. In *ECCV*, 2016.
- [22] Li Mi and Zhenzhong Chen. Hierarchical graph attention network for visual relationship detection. In CVPR, 2020.
- [23] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. In *NeurIPS*, 2017.
- [24] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.
- [25] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *PAMI*, 39(6), 2017.
- [26] Gerard Salton and Chris Buckley. Term-weighting approaches in automatic text retrieval. *Inf. Process. Manag.*, 24(5), 1988.
- [27] Jingkuan Song, Pengpeng Zeng, Lianli Gao, and Heng Tao Shen. From pixels to objects: Cubic visual attention for visual question answering. In *IJCAI*, 2018.
- [28] Kaihua Tang. A scene graph generation codebase in pytorch, 2020. https://github.com/KaihuaTang/ Scene-Graph-Benchmark.pytorch.
- [29] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In CVPR, 2020.
- [30] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *CVPR*, 2019.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [32] Yaming Wang, Vlad I. Morariu, and Larry S. Davis. Learning a discriminative filter bank within a CNN for fine-grained recognition. In CVPR, 2018.
- [33] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017.
- [34] Danfei Xu, Yuke Zhu, Christopher B. Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *CVPR*, 2017.
- [35] Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, and Tao Mei. Boosting image captioning with attributes. In *ICCV*, 2017.
- [36] 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 *ECCV*, 2018.
- [37] Alireza Zareian, Svebor Karaman, and Shih-Fu Chang. Bridging knowledge graphs to generate scene graphs. In *ECCV*, 2020.

- [38] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *CVPR*, 2018.
- [39] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual

relation detection. In CVPR, 2017.

[40] Sanqiang Zhao, Piyush Sharma, Tomer Levinboim, and Radu Soricut. Informative image captioning with external sources of information. In *ACL*, 2019.