Rethinking the Evaluation of Unbiased Scene Graph Generation

Xingchen Li¹ xingchenl@zju.edu.cn Long Chen^{†2} zjuchenlong@gmail.com Jian Shao¹ jshao@zju.edu.cn Shaoning Xiao¹ shaoningx@zju.edu.cn Songyang Zhang³ szhang83@ur.rochester.edu Jun Xiao¹ junx@zju.edu.cn

- ¹ College of Computer Science Zhejiang University Hangzhou, China
- ² Department of Electrical Engineering Columbia University
 New York, USA
 [†] Corresponding author
- ³ Department of Computer Science University of Rochester Rochester, USA

Abstract

Current Scene Graph Generation (SGG) methods tend to predict frequent predicate categories and fail to recognize rare ones due to the severe imbalanced distribution of predicates. To improve the robustness of SGG models on different predicate categories, recent research has focused on unbiased SGG and adopted mean Recall@K (mR@K) as the main evaluation metric. However, we discovered two overlooked issues about this de facto standard metric, which makes current unbiased SGG evaluation vulnerable and unfair: 1) mR@K neglects the correlations among predicates and unintentionally breaks category independence when ranking all the triplet predictions together regardless of the predicate categories. 2) mR@K neglects the compositional diversity of different predicates and assigns excessively high weights to some oversimple category samples with limited composable relation triplet types. In addition, we investigate the underexplored correlation between objects and predicates, which can serve as a simple but strong baseline for unbiased SGG. In this paper, we refine mR@K and propose two complementary evaluation metrics for unbiased SGG: Independent Mean Recall (IMR) and weighted IMR (wIMR). These two metrics are designed by considering the category independence and diversity of composable relation triplets, respectively. We compare the proposed metrics with the de facto standard metrics through extensive experiments and discuss the solutions to evaluate unbiased SGG in a more trustworthy way¹.

1 Introduction

Scene graphs are prevailing visually-grounded graph structured representations for scene understanding, which consist of a set of visual relation triplets (*i.e.*, {subject, predicate,

^{© 2022.} The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

¹Code will be available at https://github.com/xcppy/Unbiased_SGG.



Figure 1: Illustration of annotation characteristics of the VG dataset. (a) The number distribution of samples for 25 most frequent predicates in VG. (c) The annotated ground-truths of the example image (b) and some missing relation triplets. (d) Some reasonable predicates for unannotated object pairs.

object) [\Box , \Box , \Box]. Given an input image, Scene Graph Generation (SGG) requires models to not only localize and classify object categories accurately but also infer visual relationships (*i.e.*, predicates) between pairwise objects. Due to the inherent interpretability of scene graphs, they have been widely used in many downstream vision-and-language tasks, such as image retrieval [\Box , \Box , \Box], image captioning [\Box , \Box , \Box] and visual grounding [\Box].

Specifically, the implementation of mR@K can be summarized into three steps: 1) For each image, it first ranks top-K triplet predictions as outputs according to the confidence scores of all triplets and calculates the recall scores on each predicate category separately. 2) Then, it calculates the recall on each predicate category by averaging the corresponding recall scores over all images. 3) Finally, the recall scores for all predicate categories are averaged together to obtain the mR@K. *However, there are still two overlooked critical issues that make mR@K vulnerable and unfair:*

The first issue is that the default mR@K unintentionally breaks the *category independence* when ranking across categories to output top-K predictions, *i.e.*, all triplet predictions are ranked together by their confidence scores regardless of their predicate categories. By "category independence", we mean that we hope the evaluation of each category should not be influenced by other categories. The setting of ranking across categories has no problem applying to other conventional single-label classification paradigms. Unfortunately, in SGG

²For conciseness, we use "head" (or "tail") predicates to represent these predicate categories in the "head" (or "tail") part of this long-tailed distributions in the following sections.

dataset, the predicates are highly correlated and there might be multiple reasonable predicates to describe the relationship for one object pair [\square , \square]. For example in Fig. 1 (d), predicate near and behind are both correct for the subject-object pair $\langle lamp, laptop \rangle$. However, SGG is annotated following a single-label paradigm and only near is regarded as the ground-truth³. In such a case, the predictions of those predicates that have high correlations with other predicates will receive lower confidence scores after the normalization for single-label classification (*e.g.*, Softmax). They are then less likely to appear in the top-K when ranking with the predictions of other predicates with low correlations, leading to underestimated recall scores for those categories. Since each predicate category in the dataset has different correlations with other categories, it is unfair to rank their predictions together.

The second issue is that the standard mR@K assigns equal weights to all predicate categories and neglects the *compositional diversity* of predicates in subject-object categories. By "compositional diversity", we mean the range of applicable scenarios for a predicate type, which can be measured by the number of possible compositions of subject-object category pairs. In existing SGG datasets, the compositional diversity of different predicates varies greatly. Specifically, in VG dataset, the predicate on has more than 4K types of subjectobject category pairs while predicate flying-in only has two types. Meanwhile, we observe that the tail predicates with fewer samples in SGG dataset usually possess limited composable relation triplets. Thus, mR@K, by assigning equal weights to each category, actually assigns excessively high weights to the samples of these tail predicates. However, SGG is a task to understand rich visual semantics of images, which encourages models to recognize more types of visual relationship triplets. The simple category-wise averaging strategy can not reflect this goal of the task. And our experiments found that a simple baseline could easily trick mR@K by blindly utilizing the subject-object priors of those predicates with limited compositional diversity.

In addition, to further benchmark unbiased SGG methods, we investigate the intrinsic correlation between objects and predicates from a new perspective. Different from previous frequency bias [13], we turn our attention to the distribution of object categories under each predicate category and find this statistical prior can better reflect the correlation between objects and predicates. Based on this, we devise a simple baseline by directly aggregating this statistical prior into the predictions of SGG models and it greatly improves unbiased SGG performance. We think that this simple but strong method can serve as a baseline for the following unbiased SGG research.

In this paper, we refined the default mR@K and proposed two complementary metrics — Independent Mean Recall (**IMR**) and weighted IMR (**wIMR**) — to help evaluate unbiased SGG more fairly. Specifically, IMR ranks the predictions for each category independently to remove the mutual influence between categories, and wIMR assigns different weights to each predicate category by considering compositional diversity. We evaluate plenty of stateof-the-art unbiased SGG models on both current and new proposed metrics to compare their difference and discuss the solutions to evaluate unbiased SGG in a more trustworthy way.

2 Related Work

Scene Graph Generation. The mainstream SGG methods can be classified into two major categories: 1) Two-stage pipeline [2, 16, 51, 59, 12, 13], which detect the objects first and

³Following the conventions in previous SGG work, for each subject-object pair, only one ground-truth predicate is used for the training and evaluation.

then output the relationship predictions for every object pairs; 2) One-stage pipeline [24, 25, 23], which simultaneously detect objects and relationships. Most SGG studies are based on the two-stage pipeline [1, 21, 23, 53, 53]. Some work focused on learning better contextual features [51, 53, 53]. They encode visual context of images into predicate representations by exploiting various powerful context encoding architectures, such as Bi-LSTM [53], TreeL-STM [53], and graphs [13, 23, 53, 53]. In addition, some studies are interested in improving the training loss [2]. Aside from these two-stage methods, there are also some other studies focusing on one-stage pipeline [12, 24, 25, 25]. Compared to two-stage methods, they are trained in an end-to-end manner. Typically, they usually treat the detection of objects and relationships as a point detection problem. For instance, in PPDM [22] and IPNet [53], they separately localize and classify the object points and interaction points, and then group them. Besides, since the fully supervised SGG requires excessive human labeling efforts, some recent work [21, 53] begins to focus on weakly supervised SGG. They only utilized image-caption pairs to train SGG models. The main challenge of weakly supervised SGG is to generate an alignment between text entities in captions and image regions.

Evaluation Metrics for SGG. SGG evaluation is always challenging due to the intrinsic complexity of the SGG tasks and annotation characteristics. It is continuously discussed in many SGG studies [**B**, **I**, **Z**, **Z**, **Z**, **S**, **S**, **I**]. In early work [**Z**, **Z**, **S**, **S**, **I**, **J**], the most widely used metric is Recall@K (R@K). It measures the fraction of the ground-truths that appear among the top-*K* most confident predictions in an image. In [**Z**, **D**], they allowed each subject-object pair to have multiple predicates, which means all the predicates will be involved in the recall ranking for each subject-object pair not just the one with the highest score. This setting significantly improves the performance of R@K and is named No Graph Constraint Recall@K (ngR@K). In [**II**], they propose Weighted Triplet Recall@K to calculate recall by averaging the recall score on each type of visual relationship triplet. Besides, increasing recent studies have realized the long-tailed issue in SGG datasets, they propose Mean Recall@K (mR@K) to evaluate unbiased SGG. In this paper, we mainly discuss two overlooked issues by current metrics and propose two complementary metrics to evaluate unbiased SGG. Different from mR@K, IMR treats each predicate category independently and WIMR assigns weights to different categories by taking the compositional diversity into consideration.

3 Analyses of Current Unbiased SGG Benchmark

In this section, we first detailed introduce the two overlooked issues in existing metrics and show their influence on evaluation results through data analysis. Then, we discuss the solutions to these issues and provide two complementary metrics for unbiased SGG evaluation.

3.1 Category Independence

The standard mR@K ranks all the predictions together by their estimated confidence scores to output top-K predictions, and then averages recall scores for each predicate category. It originally aims to treat each predicate category independently. However, we argue that ranking across categories in each image inadvertently involves cross-category interaction [**B**]. The current SGG task is formulated as a single-label classification problem and only one predicate is regarded as the ground truth. However, due to the strong label correlation in SGG dataset, there might be multiple reasonable predicates for one object pair. For example in Figure 1(d), looking-at, in-front-of and watching are all reasonable predi-



(a) Biased Outputs. (b) Debiased Outputs. (c) Debiased Outputs after Softmax. Figure 2: Visualization of the output for each predicate (ground truth). Each row of data in the heatmap represents the averaged output distribution of the samples annotated as the predicate left in the figure. For visualization, we normalize the scores over the whole map and only display a part of the predicates.

cates for the subject-object pair $\langle cat, laptop \rangle$. In such a case, the single-label annotation setting will bring unfairness for the predicates with high correlations when ranking across categories. Consider a "perfect" SGG model which is expected to score all these possibly correct predicates high (*e.g.* score 1.0 for all the above three predicates, 0.0 for the rest), the probability scores of these categories will be suppressed relatively after the normalization for single-label classification (*e.g.* softmax). Consequently, the recall score of the watching triplets will be underestimated as these predictions are less likely to appear in top-*K* with low confidence scores. Thus, it is unfair to compare and rank all predicates together.

Dataset Analysis. To observe the influence of predicate correlations on evaluation by a real dataset, we visualize three output distributions for each predicate category based on a stateof-the-art SGG model Motifs [1] in Figure 2, including: (a) the logits distribution of biased Motifs model before softmax; (b) the logits distribution of unbiased Motifs model with reweight $[\square]$ strategy⁴; (c) the probability scores of the unbiased Motifs model after softmax. All visualization outputs are averaged over the test set of VG. From the figure, we can easily observe two phenomena: (1) Comparing Figure 2(a) and Figure 2(b), we can see that reweight strategy can alleviate the bias on the head predicates (e.g., on marked orange in the figure) shown in Figure 2(a) to a certain extent. However, we can find there are some ubiquitous correlations among some predicate categories (marked red in the figure). For example, the samples annotated as wearing normally output high scores on its alternative predicates (i.e., wears, has and in) as well. Similarly, it also occurs between walking on and standing on, belonging to and of. (2) Comparing Figure 2(b) and Figure 2(c), we can see that the normalized probabilities (used as confidence scores to rank) of the predicates (correctly recognized) with high correlations (e.g. wearing and belonging to, marked red in the figure) are suppressed relatively compared to other predicates. In the subsequent ranking process, they are more likely to be sorted behind unfairly. To this end, we suggest removing the setting of ranking across categories and sorting each category independently.

3.2 Compositional Diversity

Unbiased SGG aims to remove the impact of data imbalance in the model training process and encourage the models to recognize the tail predicates which have limited samples. Therefore, it assigns equal weights to each predicate category to reduce the proportion of

⁴The inversed sample fractions were assigned to each predicate category as weights in cross-entropy loss.



Figure 3: The number of composed object pairs for each category in VG (30 predicates are shown).

	#typepair	mR@100	Impro.%
Motifs (N=0)	-	17.2	-
N = 1 (flying in)	2	18.4	1.2
N = 2 (says)	20	19.2	2.0 +0.8
N = 3 (parked on)	29	21.0	3.8 +1.8
N = 4 (playing)	37	20.9	3.7 - 0.1
N = 5 (growing on)	41	22.3	5.1 +1.4
N = 6 (made of)	42	22.9	5.7 +0.6

Table 1: Performance (%) of Motifs [\square] on PredCls by blindly replacing the results of the corresponding subject-object pairs for N predicates with smallest compositional space.

head samples in performance. However, we argue that this category-wise averaging strategy can not reflect the rich structure of predicates on composed relationship triplets. Hence, in this paper, we investigate the predicate categories from a new perspective: compositional diversity, *i.e.*, the number of possible subject-object category pairs for each predicate.

Studies of Compositional Diversity. To observe the compositional diversity of different predicates, we count the numbers of all possible composed subject-object category pairs for each predicate category in VG dataset. As shown in Figure 3, we can see that the compositional diversity of different predicates varies greatly, *i.e.*, the head categories in the figure typically have more generic semantics and can be applied to plenty of subject-object pairs, while the tail ones in Figure 3 are more specific and can be only applied to a few subject-object category pairs. For example, predicate parked on mainly occurs between various vehicles and places, like (car, parked-on, road).

Meanwhile, we can observe the tail predicates (rare ones) in the SGG datasets usually possess limited compositional diversity. Thus, mR@K unconsciously assigns high weights to the samples of those predicates with limited compositional diversity. However, our experiments found that the performance of these tail predicates can be easily boosted by subjectobject priors. Here, we devised a simple experiment and recorded mR@K scores: 1) We firstly evaluate a well-trained state-of-the-art SGG model Motifs [1] on the PredCls⁵. 2) For the last N tail predicate categories in Figure 3, we find out all their possible composable subject-object pairs in the training dataset, then for each composable subject-object pair of predicate category c, we replace their predictions as the predicate c^6 at test. For example, when N = 1, we directly predict the predicates of all subject-object pairs with categories (kites, snow) and (planes, planes) as flying-in. We set N from 1 to 6 and record mR@K performance in Table 1. From the table, we can observe that the results on mR@K are continuously improved with the increase of N. Specifically, mR@100 is improved by 1.2% points only by blindly replacing the predictions for flying-in. Then, by further replacing the results of all corresponding subject-object pairs for predicate says, mR@100 is improved by 2.0% points accumulatively. We can observe that the predicates with limited compositional diversity have a stronger correlation with subject-object priors, which can be simply improved even without visual information.

⁵There are three typical SGG tasks: 1) *Predicate Classification* (**PredCls**): Given the ground-truth of object bounding boxes (bboxes) and class labels, the model is required to predict the predicate class of pairwise objects. 2) *Scene Graph Classification* (**SGCls**): Given the ground-truth of object bboxes, the model is required to predict both object classes and predicate classes of pairwise objects. 3) *Scene Graph Detection* (**SGDet**): Given an image, the model is required to detect object bboxes, and predict object classes and predicate classes of pairwise objects.

⁶For the same subject-object category pair with multiple tail predicate choices, we select the one with more limited compositional diversity.

3.3 Suggestion on Evaluation Metrics

To handle the above-mentioned issues, we discuss and provide two complementary metrics to improve the current unbiased SGG evaluation benchmarking. In addition, we devise a simple baseline and improve unbiased SGG performance greatly.

Independent Mean Recall (IMR). Since there are actually many multiple reasonable predicates for some subject-object pairs, it is unfair to rank all triplet predictions across categories under the single-label classification problem formulation. One intuitive and straightforward solution to solve this problem is to build a well-annotated test dataset and redefine the task as a multi-label classification problem. However, it costs expensive to annotate exhaustively all the reasonable predicates for the subject-object pairs. Hence, we provide one solution to treat each predicate category independently and avoid their mutual influence.

For each image, we independently rank and output top-K (K = 10/20/50) predictions for each predicate category to calculate their own recall scores on this image. In this way, we remove the influence of different categories and the predicates with high correlations will not be suppressed by the other predicates. Then we obtain the recall on each predicate category by averaging the corresponding scores over all images. The final score is the averaged value over categories. We call this metric Independent Mean Recall (IMR).

Weighted Independent Mean Recall (wIMR). The standard mR@K treats each predicate category equally and assigns equal weights to each predicate category. However, the target of the SGG task is to recognize more types of composed relation triplets rather than purely more kinds of predicates. Because the predicate categorical labels do not reflect the rich structures in the object relationships. The composable diversity varies widely between different predicates. Recent studies report the mean value of R@K and mR@K to expect models pay more attention to head predicates when realizing debiasing targets. However, R@K is not reliable because the variety of the number of head and tail predicates is too great. Besides, some predicates (*e.g.*, wearing) with large distribution does not have rich semantics and have limited compositional diversity as shown in Figure 5.

Therefore, we suggest reassigning weights to each predicate category c according to the complexity of their compositional space. We count the number of composed subject-object pairs n_c for each predicate category, and reassign weights to each predicate category c:

$$w(n_c) = \frac{n_c^{\tau}}{\sum_{k \in \mathcal{C}} n_k^{\tau}}, \qquad wIMR = \sum_{c \in \mathcal{C}} w(n_c) \times IMR(c), \tag{1}$$

where C is the set of predicate categories, $\tau \in [0, 1]$ controls the softness of weight distribution. When $\tau = 0$, wIMR equals IMR, which assigns equal weights to each predicate.

4 Intrinsic Correlation between Predicates and Objects

Different from other conventional scene understanding tasks, SGG predicts the predicates with corresponding subject-object pair information, which will bring in plenty of inductive information. Zellers *et al.* [1] finds the prediction of predicates is highly correlated with subject-object priors. They count the co-occurrence of subjects, objects and predicates in relation triplets and obtain a statistics matrix from training dataset $\mathbf{A} \in \mathbb{R}^{N_s \times N_o \times N_p}$ (where $N_s = N_o$ is the number of object categories, and N_p is the number of predicate categories). They aggregate the dimensions of subjects and objects to find the frequent predicate under



Figure 4: Performance of each predicate category on mR@50 and IMR@20 for Motifs+Reweight (Predcls). Categories are displayed in descending order according to the number of samples in VG.



Figure 5: Weights assigned to each predicate category on different metric settings. Categories are displayed in descending order according to the number of samples in VG.

the given pair. For example, on is the most frequent predicate under $\langle car, street \rangle$. However, this statistical prior actually collects the frequency bias and can not reflect the intrinsic correlation between predicates and objects. In this paper, we find that it is more reasonable to observe the distribution of objects over each predicate to find the correlation between predicates and objects. For example, if we count the most frequent objects given parked on, we can find its most relevant object category is car, which is in line with common sense. We call this statistical prior *Predicate Knowledge on Objects (PKO)*. Reasonably exploiting this part of statistical knowledge PKO can improve unbiased SGG models' performance.

In this paper, we devise a simple baseline to directly aggregate PKO into the inference results of SGG models (without training) following the frequency prior in Motifs [1]. From **A**, we can easily obtain the predicate-subject co-occurrence matrix $\mathbf{A}^s \in \mathbb{R}^{N_p \times N_s}$ and predicateobject co-occurrence matrix $\mathbf{A}^o \in \mathbb{R}^{N_p \times N_o}$, separately. Then, we normalize the matrices on the dimension of subject (object) and obtain the distribution of subject (object) under all predicates: $\mathbf{\tilde{A}}^s \in \mathbb{R}^{N_p \times N_s}$ and $\mathbf{\tilde{A}}^o \in \mathbb{R}^{N_p \times N_o}$. Finally, given a subject-object pair (i, j), an SGG model will output the predicate logits $\mathbf{\hat{z}}_{i,j} \in \mathbb{R}^{N_p}$, we aggregate the statistical prior by:

$$\mathbf{z}_{i,j} = \hat{\mathbf{z}}_{i,j} + \mathbf{b}_{i,j}, \qquad b_{i,j,k} = -\log \frac{\tilde{\mathbf{A}}_{k,i}^s}{\sum_{c \in \mathcal{C}} \tilde{\mathbf{A}}_{c,i}^s} - \log \frac{\tilde{\mathbf{A}}_{k,j}^o}{\sum_{c \in \mathcal{C}} \tilde{\mathbf{A}}_{c,j}^o}, \tag{2}$$

where C is the set of predicate categories and $\mathbf{z}_{i,j}$ is the final output of the SGG model. We can find this simple baseline improves unbiased performance greatly (in Sec. 5.2). The main improvements are on the predicates with limited compositional diversity which verifies the statements in Sec. 3.2, *i.e.*, these predicates are more highly correlated with subject-object priors and can be simply improved even without visual information.

5 Experiments

In this section, we develop experiments to discuss the proposed metrics and compare them with current metrics. We also evaluate different unbiased SGG methods on these metrics.

Dataset. We evaluated all results on the challenging Visual Genome [1], which is a largescale widely-used benchmark in SGG tasks. We adopt the popular split [1], 1], which includes the most frequent 150 object categories and 50 predicate classes. After preprocessing, each image has 11.5 objects and 6.2 relationships on average. The split uses 70% of images for training and 30% for the test.

	Mathad	PredCls			
	Method	R@20/50/100	mR@20 / 50 / 100	IMR@10/20/50	wIMR@10/20/50
	РКО	10.10 / 15.51 / 19.44	17.86 / 27.38 / 33.93	38.03 / 39.02 / 40.90	29.75 / 30.79 / 31.38
Motifs	Baseline	59.19 / 65.59 / 67.30	12.62 / 15.96 / 17.23	14.53 / 16.11 / 17.45	24.80 / 27.99 / 30.97
	TDE	33.35 / 45.89 / 51.24	17.85 / 24.77 / 28.72	27.33 / 29.52 / 31.06	32.64 / 35.59 / 37.72
	DLFE	44.28 / 50.33 / 51.99	21.86 / 26.81 / 28.53	26.37 / 27.92 / 29.03	31.49 / 34.02 / 35.73
	NICE	48.15 / 55.14 / 57.15	23.67 / 29.83 / 32.24	28.35 / 31.31 / 33.12	31.61 / 35.26 / 37.84
	Reweight	26.57 / 36.08 / 40.39	23.98 / 30.79 / 34.48	35.39 / 36.58 / 36.98	36.13 / 37.52 / 38.05
	RTPB(CB)	36.58 / 42.64 / 44.36	27.61 / 32.78 / 34.57	30.55 / 33.30 / 35.02	33.63 / 37.16 / 39.32
	РКО	49.08 / 55.95 / 58.18	24.98 / 31.44 / 33.98	29.73 / 32.46 / 34.21	31.77 / 35.51 / 38.51
	Baseline	59.72 / 65.86 / 67.50	13.26 / 16.82 / 18.12	15.36 / 16.99 / 18.30	25.31/28.58/31.58
	TDE	34.48 / 44.89 / 49.20	19.07 / 25.61 / 29.13	27.24 / 29.46 / 31.03	32.59 / 35.70 / 38.05
VCTree	DLFE	45.35 / 51.21 / 52.75	22.53 / 27.36 / 28.86	26.46 / 28.28 / 29.30	31.49 / 34.27 / 36.01
	NICE	48.38 / 55.03 / 56.92	24.42 / 30.74 / 33.01	29.03 / 32.01 / 33.93	31.72/35.41/38.12
	Reweight	28.66 / 35.62 / 37.90	28.64 / 34.93 / 37.28	34.70 / 36.95 / 38.30	34.41 / 36.94 / 38.42
	RTPB(CB)	36.65 / 42.39 / 43.95	28.64 / 33.41 / 35.11	30.57 / 33.47 / 33.50	33.22 / 36.99 / 39.50
	РКО	49.39 / 56.06 / 58.18	26.06 / 32.20 / 34.61	30.61 / 33.41 / 35.29	31.98 / 35.73 / 38.88

Table 2: Performance comparison of different unbiased SGG methods (*w.r.t.* Predcls) on metrics: R@K, mR@K, IMR@K, wIMR@K. VCTree [1] and Motifs [1] are two popular compared baseline models in unbiased SGG. All experimental results are re-implemented using official released codes. The **best** and second best results are marked according to formats.

5.1 Comparison between Existing and New Metrics

IMR@K vs. mR@K. In Figure 4, we display the recall scores of each predicate category on mR@50 and IMR@20 for Motifs-Reweight *w.r.t.* PredCls, respectively. We can see that the recall scores of the predicates in Figure 2 underestimated by mR@K are more fairly evaluated on IMR@K. For example, the recall scores of the high correlated predicates, like wearing, of and watching, ascend significantly on IMR@20 compared to mR@50. While the predicate parked on has no obvious improvement on IMR@20 as it is slightly influenced by the confidence sharing shown in Figure 2(c). Compared to mR@K, IMR@K provides a more fair score for these over-estimated predicate categories.

wIMR@K vs. mR@K/R@K. In Figure 5, we display the weights of different predicates (on VG dataset) in R@K, mR@K/IMR ($\tau = 0.0$) and wIMR ($\tau = 0.5$). We can observe that the classic R@K mainly focuses on head predicates although some head predicates have limited compositional diversity (*e.g.* wearing). While mR@K assigns equal weights to all predicates despite their difference in compositional diversity, and the samples of some predicates with limited compositional diversity like flying-in are assigned excessively high weights. In contrast, wIMR, by considering compositional diversity, manages to assign high weights to predicates with rich semantics (*e.g.*,of and under) and low weights to predicates with simple semantics (*e.g.*,flying-in and wearing). The value of τ controls the softness of the weight distribution and can be adjusted according to task target, we set $\tau = 0.5$ in the default implementation of wIMR.

5.2 Evaluation on Current Methods

Benchmarking SOTA Unbiased SGG Methods. For the convenience to further compare with current unbiased SGG methods, we report the performance of five SOTA unbiased SGG methods on current and new proposed metrics, including *TDE* [\Box], *DLFE* [\Box], *NICE* [\Box], *RTPB*(*CB*) [\Box] and *Reweight* [\Box]. The details are reported in Table 2 and Table 3.

Effectiveness of Predicate Knowledge on Objects (PKO). We also report the performance of models with our proposed statistical prior in Table 2 and Table 3. Wherein, the simple

	Mathad		SGCls			
	Method	R@20/50/100	mR@20/50/100	IMR@10/20/50	wIMR@10/20/50	
	РКО	7.26 / 10.48 / 12.46	11.23 / 16.77 / 20.13	20.51 / 22.47 / 23.11	16.04 / 17.13 / 17.68	
	Baseline	36.39 / 39.59 / 40.35	7.44 / 9.09 / 9.63	7.69 / 8.65 / 9.57	13.65 / 15.70 / 17.76	
	TDE	20.46 / 26.31 / 28.78	9.78 / 13.21 / 15.07	13.98 / 15.22 / 16.29	16.95 / 18.69 / 20.16	
Motifs	DLFE	26.63 / 29.79 / 30.61	13.23 / 15.66 / 16.48	14.75 / 15.82 / 16.57	17.96 / 19.65 / 20.85	
	NICE	29.48 / 33.06 / 34.05	13.63 / 16.67 / 17.88	14.82 / 16.76 / 18.02	17.00 / 19.35 / 21.07	
	Reweight	18.67 / 23.49 / 25.47	13.49 / 16.75 / 18.34	17.82 / 18.69 / 19.23	19.13 / 20.38 / 21.14	
	RTPB(CB)	22.66 / 25.85 / 26.65	15.77 / 18.16 / 18.97	16.08 / 17.68 / 18.81	18.51 / 20.61 / 21.99	
	РКО	30.41 / 33.99 / 35.05	14.06 / 17.59 / 19.12	15.87 / 17.49 / 18.66	17.43 / 19.72 / 21.72	
	Baseline	42.09 / 45.80 / 46.73	9.09 / 11.28 / 12.04	9.72 / 10.89 / 11.94	16.52 / 18.90 / 21.26	
	TDE	23.48 / 31.17 / 34.59	10.36 / 14.47 / 16.72	15.81 / 17.32 / 18.48	19.71 / 21.73 / 23.35	
VCTree	DLFE	30.09 / 33.85 / 34.80	16.17 / 19.23 / 20.20	18.81 / 20.04 / 20.62	21.75 / 23.58 / 24.65	
	NICE	33.77 / 37.84 / 38.99	16.14 / 20.03 / 21.29	17.89 / 19.91 / 21.59	20.16 / 22.77 / 24.80	
	Reweight	20.44 / 24.66 / 25.97	18.72 / 22.88 / 24.19	22.14 / 23.66 / 24.58	22.04 / 23.93 / 25.07	
	RTPB(CB)	26.06 / 29.63 / 30.61	18.67 / 21.41 / 22.52	19.64 / 21.41 / 22.63	22.04 / 24.41 / 25.93	
	РКО	35.01 / 39.10 / 40.40	18.43 / 22.27 / 23.74	19.88 / 21.72 / 23.01	21.14 / 23.77 / 26.09	
		SGDet				
	Mathad		SG	Det		
	Method	R@20/50/100	SG mR@20 / 50 / 100	Det IMR@10/20/50	wIMR@10/20/50	
	Method PKO	R@20/50/100 5.53/7.77/9.46	SG mR@20 / 50 / 100 7.48 / 10.97 / 13.85	Det IMR@10 / 20 / 50 19.37 / 19.87 / 20.76	wIMR@10/20/50 15.17/15.48/16.14	
	Method PKO Baseline	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49	SG mR@20 / 50 / 100 7.48 / 10.97 / 13.85 5.37 / 7.37 / 8.61	Det IMR@10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85	wIMR@10/20/50 15.17/15.48/16.14 14.98/15.26/15.80	
	Method PKO Baseline TDE	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15	SG mR@20 / 50 / 100 7.48 / 10.97 / 13.85 5.37 / 7.37 / 8.61 6.54 / 8.93 / 10.95	Det IMR@10/20/50 19.37/19.87/20.76 7.99/8.33/8.85 13.67/13.89/14.55	wIMR@10/20/50 15.17/15.48/16.14 14.98/15.26/15.80 17.58/17.72/18.18	
Motifs	Method PKO Baseline TDE DLFE	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45	SG mR@20 / 50 / 100 7.48 / 10.97 / 13.85 5.37 / 7.37 / 8.61 6.54 / 8.93 / 10.95 8.50 / 11.29 / 13.24	Det IMR@10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85 13.67 / 13.89 / 14.55 12.15 / 13.18 / 14.58	wIMR @ 10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09	
Motifs	Method PKO Baseline TDE DLFE NICE	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75	SG mR@20/50/100 7.48/10.97/13.85 5.37/7.37/8.61 6.54/8.93/10.95 8.50/11.29/13.24 8.82/12.23/14.40	Det IMR@10/20/50 19.37/19.87/20.76 7.99/8.33/8.85 13.67/13.89/14.55 12.15/13.18/14.58 12.89/13.42/14.41	wIMR @ 10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75 12.63/17.26/20.59	SG mR@20 / 50 / 100 7.48 / 10.97 / 13.85 5.37 / 7.37 / 8.61 6.54 / 8.93 / 10.95 8.50 / 11.29 / 13.24 8.82 / 12.23 / 14.40 8.88 / 11.78 / 14.09	Det IMR@10/20/50 19.37/19.87/20.76 7.99/8.33/8.85 13.67/13.89/14.55 12.15/13.18/14.58 12.89/13.42/14.41 15.36/16.12/17.75	wIMR @ 10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight RTPB(CB)	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75 12.63/17.26/20.59 14.72/20.13/23.75	SG mR@20 / 50 / 100 7.48 / 10.97 / 13.85 5.37 / 7.37 / 8.61 6.54 / 8.93 / 10.95 8.50 / 11.29 / 13.24 8.82 / 12.23 / 14.40 8.88 / 11.78 / 14.09 11.35 / 14.45 / 16.74	Det IMR@10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85 13.67 / 13.89 / 14.55 12.15 / 13.18 / 14.58 12.89 / 13.42 / 14.41 15.36 / 16.12 / 17.75 13.60 / 14.35 / 15.91	wIMR @ 10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31 16.65 / 17.15 / 18.31	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight RTPB(CB) PKO	R@20 / 50 / 100 5.53 / 7.77 / 9.46 25.79 / 33.04 / 37.49 11.94 / 16.57 / 20.15 18.29 / 24.52 / 28.45 21.30 / 27.83 / 31.75 12.63 / 17.26 / 20.59 14.72 / 20.13 / 23.75 20.38 / 26.95 / 31.14	SG mR@20 / 50 / 100 7.48 / 10.97 / 13.85 5.37 / 7.37 / 8.61 6.54 / 8.93 / 10.95 8.50 / 11.29 / 13.24 8.82 / 12.23 / 14.40 8.88 / 11.78 / 14.09 11.35 / 14.45 / 16.74 9.64 / 13.44 / 16.06	Det IMR@10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85 13.67 / 13.89 / 14.55 12.15 / 13.18 / 14.58 12.89 / 13.42 / 14.41 15.36 / 16.12 / 17.75 13.60 / 14.35 / 15.91 14.66 / 15.39 / 16.82	wIMR@10/20/50 15.17/15.48/16.14 14.98/15.26/15.80 17.58/17.72/18.18 16.12/16.88/18.09 16.54/16.96/17.78 17.71/18.20/19.31 16.65/17.15/18.31 17.48/17.96/18.97	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight RTPB(CB) PKO Baseline	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75 12.63/17.26/20.59 14.72/20.13/23.75 20.38/26.95/31.14 24.96/32.15/36.36	SG mR@20/50/100 7.48/10.97/13.85 5.37/7.37/8.61 6.54/8.93/10.95 8.50/11.29/13.24 8.82/12.23/14.40 8.88/11.78/14.09 11.35/14.45/16.74 9.64/13.44/16.06 5.20/7.05/8.28	Det IMR@10/20/50 19.37/19.87/20.76 7.99/8.33/8.85 13.67/13.89/14.55 12.15/13.18/14.58 12.89/13.42/14.41 15.36/16.12/17.75 13.60/14.35/15.91 14.66/15.39/16.82 7.44/7.68/8.11	wIMR@10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31 16.65 / 17.15 / 18.31 17.48 / 17.96 / 18.97 14.34 / 14.53 / 14.98	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight RTPB(CB) PKO Baseline TDE	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75 12.63/17.26/20.59 14.72/20.13/23.75 20.38/26.95/31.14 24.96/32.15/36.36 12.11/17.07/20.72	SG mR@20/50/100 7.48/10.97/13.85 5.37/7.37/8.61 6.54/8.93/10.95 8.50/11.29/13.24 8.82/12.23/14.40 8.88/11.78/14.09 11.35/14.45/16.74 9.64/13.44/16.06 5.20/7.05/8.28 6.73/9.32/11.28	Det IMR @ 10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85 13.67 / 13.89 / 14.55 12.15 / 13.18 / 14.58 12.89 / 13.42 / 14.41 15.36 / 16.12 / 17.75 13.60 / 14.35 / 15.91 14.66 / 15.39 / 16.82 7.44 / 7.68 / 8.11 13.57 / 13.73 / 14.28	wIMR@10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31 16.65 / 17.15 / 18.31 17.48 / 17.96 / 18.97 14.34 / 14.53 / 14.98 17.69 / 17.81 / 18.17	
Motifs	Method PKO Baseline DLFE NICE Reweight RTPB(CB) PKO Baseline DLFE	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75 12.63/17.26/20.59 14.72/20.13/23.75 20.38/26.95/31.14 24.96/32.15/36.36 2.11/17.07/20.72 17.61/23.58/27.21	SG mR@20/50/100 7.48/10.97/13.85 5.37/7.37/8.61 6.54/8.93/10.95 8.50/11.29/13.24 8.82/12.23/14.40 8.88/11.78/14.09 11.35/14.45/16.74 9.64/13.44/16.06 5.20/7.05/8.28 6.73/9.32/11.28 8.76/11.63/13.32	Det IMR @ 10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85 13.67 / 13.89 / 14.55 12.15 / 13.18 / 14.58 12.89 / 13.42 / 14.41 15.36 / 16.12 / 17.75 13.60 / 14.35 / 15.91 14.66 / 15.39 / 16.82 7.44 / 7.68 / 8.11 13.57 / 13.73 / 14.28 11.80 / 12.73 / 14.24	wIMR @10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31 16.65 / 17.15 / 18.31 17.48 / 17.96 / 18.97 14.34 / 14.53 / 14.98 17.69 / 17.81 / 18.17 15.03 / 15.68 / 16.87	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight RTPB(CB) PKO Baseline TDE DLFE NICE	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75 12.63/17.26/20.59 14.72/20.13/23.75 20.38/26.95/31.14 24.96/32.15/36.36 12.11/17.07/20.72 17.61/23.58/27.21 20.70/26.99/30.78	SG mR@20/50/100 7.48/10.97/13.85 5.37/7.37/8.61 6.54/8.93/10.95 8.50/11.29/13.24 8.82/12.23/14.40 8.88/11.78/14.09 11.35/14.45/16.74 9.64/13.44/16.06 5.20/7.05/8.28 6.73/9.32/11.28 8.76/11.63/13.32 8.39/11.96/14.08	Det IMR@10/20/50 19.37/19.87/20.76 7.99/8.33/8.85 13.67/13.89/14.55 12.15/13.18/14.58 12.89/13.42/14.41 15.36/16.12/17.75 13.60/14.35/15.91 14.66/15.39/16.82 7.44/7.68/8.11 13.57/13.73/14.28 11.80/12.73/14.24 12.84/13.28/14.27	wIMR @10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31 16.65 / 17.15 / 18.31 17.48 / 17.96 / 18.97 14.34 / 14.53 / 14.98 17.69 / 17.81 / 18.17 15.03 / 15.68 / 16.87 16.13 / 16.49 / 17.30	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight RTPB(CB) PKO Baseline TDE DLFE NICE Reweight	R@20/50/100 5.53/7.77/9.46 25.79/33.04/37.49 11.94/16.57/20.15 18.29/24.52/28.45 21.30/27.83/31.75 12.63/17.26/20.59 14.72/20.13/23.75 20.38/26.95/31.14 24.96/32.15/36.36 12.11/17.07/20.72 17.61/23.58/27.21 20.70/26.99/30.78 15.35/20.30/23.72	SG mR@20/50/100 7.48/10.97/13.85 5.37/7.37/8.61 6.54/8.93/10.95 8.50/11.29/13.24 8.82/12.23/14.40 8.88/11.78/14.09 11.35/14.45/16.74 9.64/13.44/16.06 5.20/7.05/8.28 6.73/9.32/11.28 8.76/11.63/13.32 8.39/11.96/14.08 9.90/12.69/15.00	Det IMR@10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85 13.67 / 13.89 / 14.55 12.15 / 13.18 / 14.58 12.89 / 13.42 / 14.41 15.36 / 16.12 / 17.75 13.60 / 14.35 / 15.91 14.66 / 15.39 / 16.82 7.44 / 7.68 / 8.11 13.57 / 13.73 / 14.28 11.80 / 12.73 / 14.24 12.84 / 13.28 / 14.27 13.23 / 14.08 / 15.80	wIMR @ 10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31 16.65 / 17.15 / 18.31 17.48 / 17.96 / 18.97 14.34 / 14.53 / 14.98 17.69 / 17.81 / 18.17 15.03 / 15.68 / 16.87 16.13 / 16.49 / 17.30 15.98 / 16.50 / 17.73	
Motifs	Method PKO Baseline TDE DLFE NICE Reweight RTPB(CB) Baseline TDE DLFE NICE Reweight RTPB(CB)	$\begin{array}{c} R@20/50/100\\ \hline 5.53/7.77/9.46\\ \hline 25.79/33.04/37.49\\ \hline 11.94/16.57/20.15\\ \hline 18.29/24.52/28.45\\ \hline 21.30/27.83/31.75\\ \hline 12.63/17.26/20.59\\ \hline 14.72/20.13/23.75\\ \hline 20.38/26.95/31.14\\ \hline 24.96/32.15/36.36\\ \hline 12.11/17.07/20.72\\ \hline 17.61/23.58/27.21\\ \hline 20.70/26.99/30.78\\ \hline 15.35/20.30/23.72\\ \hline 14.52/19.70/23.17\\ \hline \end{array}$	SG mR@20/50/100 7.48/10.97/13.85 5.37/7.37/8.61 6.54/8.93/10.95 8.50/11.29/13.24 8.82/12.23/14.40 8.88/11.78/14.09 11.35/14.45/16.74 9.64/13.44/16.06 5.20/7.05/8.28 6.73/9.32/11.28 8.76/11.63/13.32 8.39/11.96/14.08 9.90/12.69/15.00 10.79/13.68/15.85	Det IMR @ 10 / 20 / 50 19.37 / 19.87 / 20.76 7.99 / 8.33 / 8.85 13.67 / 13.89 / 14.55 12.15 / 13.18 / 14.58 12.89 / 13.42 / 14.41 15.36 / 16.12 / 17.75 13.60 / 14.35 / 15.91 14.66 / 15.39 / 16.82 7.44 / 7.68 / 8.11 13.57 / 13.73 / 14.28 11.80 / 12.73 / 14.24 12.84 / 13.28 / 14.27 13.23 / 14.08 / 15.80 12.89 / 13.51 / 15.12	wIMR @10 / 20 / 50 15.17 / 15.48 / 16.14 14.98 / 15.26 / 15.80 17.58 / 17.72 / 18.18 16.12 / 16.88 / 18.09 16.54 / 16.96 / 17.78 17.71 / 18.20 / 19.31 16.65 / 17.15 / 18.31 17.48 / 17.96 / 18.97 14.34 / 14.53 / 14.98 17.69 / 17.81 / 18.17 15.03 / 15.68 / 16.87 16.13 / 16.49 / 17.30 15.98 / 16.50 / 17.73 15.95 / 16.38 / 17.59	

Table 3: Performance on SGCls and SGDet. All experimental results are re-implemented using official released codes. The **best** and second best results are marked according to formats.

baseline, *PKO*, means that we only utilize the statistical prior to make predictions. Amazingly, we can find it achieves excellent performance on IMR@K and mR@K. However, as it mainly improves the performance of the predicates with limited compositional diversity, it did not perform well on wIMR@K and R@K. We aggregate PKO into other SGG models, and it greatly improves their unbiased performance (*e.g. Motifs-PKO* and *VCTree-PKO*).

6 Conclusions

In this paper, we focused on two overlooked issues which make the current evaluation benchmark vulnerable and unfair: 1) mR@K unintentionally breaks the category independence when ranking across categories; 2) mR@K assigns equal weights for each predicate but neglects the compositional diversity of subject-object pairs. We discussed the influence of these issues through statistical data analysis and provided two complementary metrics for unbiased SGG evaluation to address the two above-mentioned issues. We additionally investigated a statistical prior which reflects the intrinsic correlation between predicates and objects. Based on this, we devised a simple but strong baseline for unbiased SGG research. We hope that our work is able to contribute some insights to this area and help benchmark unbiased SGG methods in a more trustworthy way. Furthermore, our discussion is general which may be extended to video SGG scenarios $[\Box, [S]]$.

Acknowledgements

This work was supported by the National Key Research & Development Project of China (2021ZD0110700), the National Natural Science Foundation of China (U19B2043, 61976185), Zhejiang Natural Science Foundation (LR19F020002), Zhejiang Innovation Foundation (2019R52002), and the Fundamental Research Funds for the Central Universities (226-2022-00087).

References

- [1] Chao Chen, Yibing Zhan, Baosheng Yu, Liu Liu, Yong Luo, and Bo Du. Resistance training using prior bias: toward unbiased scene graph generation. *arXiv preprint arXiv:2201.06794*, 2022.
- [2] Long Chen, Hanwang Zhang, Jun Xiao, Xiangnan He, Shiliang Pu, and Shih-Fu Chang. Counterfactual critic multi-agent training for scene graph generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4613–4623, 2019.
- [3] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 6163–6171, 2019.
- [4] Vincent S Chen, Paroma Varma, Ranjay Krishna, Michael Bernstein, Christopher Re, and Li Fei-Fei. Scene graph prediction with limited labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2580–2590, 2019.
- [5] Meng-Jiun Chiou, Henghui Ding, Hanshu Yan, Changhu Wang, Roger Zimmermann, and Jiashi Feng. Recovering the unbiased scene graphs from the biased ones. *arXiv* preprint arXiv:2107.02112, 2021.
- [6] Achal Dave, Piotr Dollár, Deva Ramanan, Alexander Kirillov, and Ross Girshick. Evaluating large-vocabulary object detectors: The devil is in the details. *arXiv preprint arXiv:2102.01066*, 2021.
- [7] Kaifeng Gao, Long Chen, Yifeng Huang, and Jun Xiao. Video relation detection via tracklet based visual transformer. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4833–4837, 2021.
- [8] Kaifeng Gao, Long Chen, Yulei Niu, Jian Shao, and Jun Xiao. Classification-thengrounding: Reformulating video scene graphs as temporal bipartite graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19497–19506, 2022.
- [9] Jiuxiang Gu, Shafiq Joty, Jianfei Cai, Handong Zhao, Xu Yang, and Gang Wang. Unpaired image captioning via scene graph alignments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10323–10332, 2019.
- [10] Chenchen Jing, Yuwei Wu, Mingtao Pei, Yao Hu, Yunde Jia, and Qi Wu. Visualsemantic graph matching for visual grounding. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4041–4050, 2020.

12 LI, CHEN, SHAO ET AL.: RETHINKING THE EVALUATION OF UNBIASED SGG

- [11] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *CVPR*, pages 3668– 3678, 2015.
- [12] Bumsoo Kim, Taeho Choi, Jaewoo Kang, and Hyunwoo J Kim. Uniondet: Unionlevel detector towards real-time human-object interaction detection. In *ECCV*, pages 498–514, 2020.
- [13] Boris Knyazev, Harm de Vries, Cătălina Cangea, Graham W Taylor, Aaron Courville, and Eugene Belilovsky. Graph density-aware losses for novel compositions in scene graph generation. *arXiv preprint arXiv:2005.08230*, 2020.
- [14] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.
- [15] Lin Li, Long Chen, Yifeng Huang, Zhimeng Zhang, Songyang Zhang, and Jun Xiao. The devil is in the labels: Noisy label correction for robust scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18869–18878, 2022.
- [16] Lin Li, Long Chen, Hanrong Shi, Wenxiao Wang, Jian Shao, Yi Yang, and Jun Xiao. Label semantic knowledge distillation for unbiased scene graph generation. arXiv preprint arXiv:2208.03763, 2022.
- [17] Lin Li, Long Chen, Hanrong Shi, Hanwang Zhang, Yi Yang, Wei Liu, and Jun Xiao. Nicest: Noisy label correction and training for robust scene graph generation. arXiv preprint arXiv:2207.13316, 2022.
- [18] Rongjie Li, Songyang Zhang, Bo Wan, and Xuming He. 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*, pages 11109–11119, 2021.
- [19] Xiangyang Li and Shuqiang Jiang. Know more say less: Image captioning based on scene graphs. *IEEE Transactions on Multimedia*, 21(8):2117–2130, 2019.
- [20] Xingchen Li, Long Chen, Wenbo Ma, Yi Yang, and Jun Xiao. Integrating object-aware and interaction-aware knowledge for weakly supervised scene graph generation. *arXiv* preprint arXiv:2208.01834, 2022.
- [21] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1261–1270, 2017.
- [22] Yikang Li, Wanli Ouyang, Bolei Zhou, Jianping Shi, Chao Zhang, and Xiaogang Wang. Factorizable net: an efficient subgraph-based framework for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 335–351, 2018.

- [23] Yuanzhi Liang, Yalong Bai, Wei Zhang, Xueming Qian, Li Zhu, and Tao Mei. Vrr-vg: Refocusing visually-relevant relationships. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10403–10412, 2019.
- [24] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jiashi Feng. Ppdm: Parallel point detection and matching for real-time human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 482–490, 2020.
- [25] Hengyue Liu, Ning Yan, Masood S Mortazavi, and Bir Bhanu. Fully convolutional scene graph generation. *arXiv preprint arXiv:2103.16083*, 2021.
- [26] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *European conference on computer vision*, pages 852–869. Springer, 2016.
- [27] Li Mi and Zhenzhong Chen. Hierarchical graph attention network for visual relationship detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13886–13895, 2020.
- [28] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. *arXiv* preprint arXiv:1706.07365, 2017.
- [29] Mengshi Qi, Weijian Li, Zhengyuan Yang, Yunhong Wang, and Jiebo Luo. Attentive relational networks for mapping images to scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3957–3966, 2019.
- [30] Brigit Schroeder and Subarna Tripathi. Structured query-based image retrieval using scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 178–179, 2020.
- [31] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6619–6628, 2019.
- [32] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3716–3725, 2020.
- [33] Hongshuo Tian, Ning Xu, An-An Liu, and Yongdong Zhang. Part-aware interactive learning for scene graph generation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3155–3163, 2020.
- [34] Sijin Wang, Ruiping Wang, Ziwei Yao, Shiguang Shan, and Xilin Chen. Cross-modal scene graph matching for relationship-aware image-text retrieval. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1508–1517, 2020.
- [35] Tiancai Wang, Tong Yang, Martin Danelljan, Fahad Shahbaz Khan, Xiangyu Zhang, and Jian Sun. Learning human-object interaction detection using interaction points. In *CVPR*, pages 4116–4125, 2020.

14 LI, CHEN, SHAO ET AL.: RETHINKING THE EVALUATION OF UNBIASED SGG

- [36] Weitao Wang, Ruyang Liu, Meng Wang, Sen Wang, Xiaojun Chang, and Yang Chen. Memory-based network for scene graph with unbalanced relations. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2400–2408, 2020.
- [37] Wenbin Wang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Exploring context and visual pattern of relationship for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2019.
- [38] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5410–5419, 2017.
- [39] Shaotian Yan, Chen Shen, Zhongming Jin, Jianqiang Huang, Rongxin Jiang, Yaowu Chen, and Xian-Sheng Hua. Pcpl: Predicate-correlation perception learning for unbiased scene graph generation. In *Proceedings of the 28th ACM International Conference* on Multimedia, pages 265–273, 2020.
- [40] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *Proceedings of the European conference on computer vision* (ECCV), pages 670–685, 2018.
- [41] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision* and Pattern Recognition, pages 10685–10694, 2019.
- [42] Jing Yu, Yuan Chai, Yue Hu, and Qi Wu. Cogtree: Cognition tree loss for unbiased scene graph generation. *arXiv preprint arXiv:2009.07526*, 2020.
- [43] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018.
- [44] Yiwu Zhong, Jing Shi, Jianwei Yang, Chenliang Xu, and Yin Li. Learning to generate scene graph from natural language supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1823–1834, 2021.
- [45] Hao Zhou, Chongyang Zhang, and Chuanping Hu. Visual relationship detection with relative location mining. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 30–38, 2019.