

Tackling the Unannotated: Scene Graph Generation with Bias-Reduced Models

Tzu-Jui Julius Wang
tzu-jui.wang@aalto.fi

Selen Pehlivan
selen.pehlivantort@aalto.fi

Jorma Laaksonen
jorma.laaksonen@aalto.fi

Department of Computer Science
School of Science
Aalto University, Espoo, Finland

Abstract

Predicting a scene graph that captures visual entities and their interactions in an image has been considered a crucial step towards full scene comprehension. Recent scene graph generation (SGG) models have shown their capability of capturing the most frequent relations among visual entities. However, the state-of-the-art results are still far from satisfactory, e.g. models can obtain 31% in *overall recall* R@100, whereas the likewise important *mean class-wise recall* mR@100 is only around 8% on Visual Genome (VG). The discrepancy between R and mR results urges to shift the focus from pursuing a high R to a high mR with a still competitive R. We suspect that the observed discrepancy stems from both the annotation bias and sparse annotations in VG, in which many visual entity pairs are either not annotated at all or only with a single relation when multiple ones could be valid. To address this particular issue, we propose a novel SGG training scheme that capitalizes on self-learned knowledge. It involves two relation classifiers, one offering a less biased setting for the other to base on. The proposed scheme can be applied to most of the existing SGG models and is straightforward to implement. We observe significant relative improvements in mR (between +6.6% and +20.4%) and competitive or better R (between -2.4% and +0.3%) across all standard SGG tasks.

1 Introduction

Various deep neural network models have been introduced to interpret visual inputs in varying high-level vision tasks, including object detection and segmentation [20, 58], image/video captioning [8, 32], and referring expression comprehension [23, 34] etc. While describing visual content and identifying visually grounded entities¹ from free-form language have been made possible, inferring about the role of entities (either in isolation or relative to other entities) is a crucial step towards full scene comprehension. This fosters the idea of representing an image with a *scene graph* [16] in which the entity classes and their pairwise relations constitute the nodes and directed edges, respectively. In particular, scene graph is the formulation of the visual subject-relation-object triplets (e.g. mouse-attached to-laptop) over entity and relation classes in an image. The Visual Genome (VG) dataset

© 2020. The copyright of this document resides with its authors.

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

¹We use *entity* here instead of *object* (as in object detector) to prevent mixing it up with the *object* as in subject-object pair.

[16] is introduced to facilitate the advancement in the scene graph generation (SGG) tasks [1, 3, 10, 11, 28, 33, 35, 36]. However, as first studied in [35], the annotations in VG bias towards copiously recurring motifs. A strong baseline can be constructed by predicting the relation between entities solely with the priors made on all subject-relation-object triplets from the training set. Moreover, the scene graph of each image is hardly exhaustively annotated with all possible relations [9] in which the annotated relations are often less informative [29] (e.g. labelling a relation as *girl-near-car* instead of *girl-behind-car*). Hence, the provided scene graphs inherit the imbalanced and long-tailed natures [19]. This incurs an apparent complication: the frequent relations, such as *on*, *has*, or *near* can easily be exploited and more often predicted by a model than the rare ones, such as *walking on*, *wearing*, or *in front of*. Then, one can already attain a decent recall rate by simply making predictions from the most frequent relations.

In addition to the imbalanced distribution of the relations, an SGG model can be biased due to the sparse and incomplete annotations [9]. Specifically, first, due to the large number of entities in most of the images, the related entity pairs are sparsely identified and annotated, leaving most of the entity pairs presumed as not related. Second, a subject-object pair can often be reasonably related by more than one relation as in *on*, *sitting on*, and *riding* relations for the *man-bicycle* pair. We refer to the former case as **annotation False Negative** (aFN), and the latter as **annotation Partially True Positive** (aPTP). It comes to our attention that nearly all of the SGG models do not address these cases. Overlooking these cases can, (1) wrongly suppress the likelihood of a related entity pair in the predictions, or (2) discourage the model from learning a multi-mode distribution that captures more than one valid relation.

This work aims to address the aFN and aPTP cases in the SGG models. Inspired by knowledge distillation (KD) [13, 22, 27, 37] and co-training framework [24], we propose a novel SGG framework that regularizes a model’s predictions with self-learned knowledge in reasoning the missing information from the annotations. In addition, to alleviate the adverse effect from learning on imbalanced data, the proposed relation models are trained on the hyperspherical space [9, 18]. To be more specific, our framework consists of two learners that jointly learn on different assumptions and output spaces. The main learner \mathcal{F} naively learns on all relation classes including treating the unannotated entity pairs as the no-relation class. The other *skeptical* learner \mathcal{G} learns on all the relations except no-relation, i.e. no assumption is made for \mathcal{G} on those not annotated being really not related. Over the course of the training, \mathcal{G} transfers the *knowledge* (i.e. the predictions) on both annotated and unannotated entity pairs to \mathcal{F} , offering a crucial source of regularization that averts \mathcal{F} learning from the biased annotations (see Section 3).

Our main contributions can be summarized as follows. (1) We propose an SGG framework that addresses biases inherit in the incomplete annotations of the VG dataset. (2) The proposed model is learned on the hyperspherical space, which is shown better at capturing the minority class distribution. (3) The proposed scheme is by design compatible with the most of the existing SGG models and is easily implemented. We demonstrate the effectiveness of the proposed scheme on two state-of-the-art SGGs, Motifs [35] and VCTree [28, 29], and observe significant relative improvements over the baseline counterparts in the mean per-class recall (between +6.6% and +20.4%) while maintaining competitive overall recall (between -2.4% to +0.3%) across all the standard SGG tasks [31, 35] (see Section 4).

2 Related Work

Scene Graph Generation. SGG offers the possibility of modeling the layout of visual entities and holistically describing the image context. Often being a critical modality in most of the SGG works, contextual modeling is usually implemented by a recurrent neural network (RNN) or a graph neural network (GNN) which learns the underlying structure, e.g. a fully connected [2, 3, 51, 53], a chained [35], or a tree-structured graph [28], among the visual entities. Chen et al. [2] formulated SGG as a reinforcement learning problem which aims at directly maximizing the non-differentiable graph-level rewards (e.g. Recall@K). Zhang et al. [57] proposed a new set of contrastive losses to address (1) repetitive occurrence of entities of the same visual class and (2) confusion caused by the close proximity of similar entities. Despite the improved performance achieved by these state-of-the-art models, SGG models are still facing challenges due to the benchmark datasets being biased towards few majority relation classes. Only recently have some works started to address the imbalance issue, shifting the focus of improvement from the majority relation classes to the minority classes [3, 28, 29].

In one of the latest SGG works, Tang et al. [29] proposed a debiased method at the inference phase of SGG models, removing the concentrated probability mass from the majority classes. As one of the few SGG works that address the issues caused by the incomplete annotations, Chen et al. [4] proposed a data-efficient scheme that trains the SGG models in a semi-supervised fashion by using a limited labeled and a large amount of unlabeled samples.

Our proposed method is inspired and closely related to that in [4], but differs in several facets. First, the proposed models co-train two learners, one supervised and another semi-supervised, in one training phase. Second, our supervised learner regularizes itself with the knowledge learned by the semi-supervised learner, significantly improving the predictions on the minority classes. Third, the objective in [4] is to train a generative model that generates informative labels on the unannotated pairs given limited amount of labeled data per relation class. However, our objective is to show that an existing SGG model can be greatly improved by harnessing the soft labels on both the annotated and unannotated pairs.

Learning on Imbalanced Data. Learning on imbalanced distributions is an inevitable subject to consider in most of the vision tasks [1, 3, 14, 19]. Some common techniques, such as undersampling, oversampling, and class re-weighting [6] with a cost-sensitive objective [15], have been proven to be effective. The recent studies suggest that a classification model learned with the cross entropy (XE) loss can easily be biased towards majority classes. That is because the XE loss does not explicitly enforce small intra-class variance [9, 17].

A simple modification, which we adopt in this work, is to respectively L_2 -normalize the input representation and the classifier’s weights before calculating the logits (pre-softmax values). This has been shown effective and usually considered as a strong baseline for making predictions in the low-data regime [3, 3, 14, 19]. Accordingly, we adopt the same normalization scheme and empirically prove its effectiveness in the SGG problem.

Knowledge Distillation. Learning a compact "student" model by distilling knowledge from a more complex "teacher" model can lead to a better model than one trained from scratch [13]. Recently, it has been shown that a neural network model can *self-distill* its knowledge without referring to a teacher model to attain better accuracy [11, 57]. *Knowledge distillation* (KD) approaches indicate that even the imperfect predictions (e.g. the predicted class probabilities) can contain useful signals, e.g. the correlations between classes, that

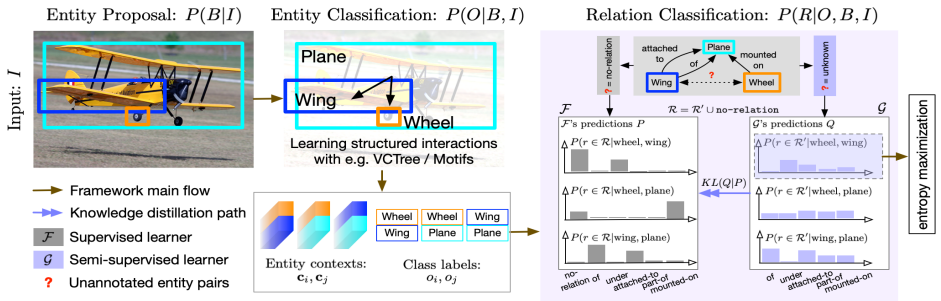


Figure 1: Overview of our proposed two-learner framework. Two learners \mathcal{F} and \mathcal{G} learn on different output spaces \mathcal{R} and \mathcal{R}' , respectively, that make different assumptions on the unannotated entity pairs. Over the course of training, \mathcal{G} transfers its less biased knowledge to \mathcal{F} whose predictions are then regularized to capture more relevant relations among entities. Wheel-Wing refers to the case where the annotation is missing a potential relation under, while Wheel-Plane refers to the case where they could be related by, e.g. part of.

help regularize the model to curb the over-confident predictions [27]. We propose two KD variants to combat against the dataset biases (see Section 1) imposed to the models.

3 Proposed Framework

A scene graph \mathcal{S} consists of (B, O, R) in which $B = \{B_i | i = 1, 2, \dots, M\}$, $O = \{o_i \in \mathcal{O} | i = 1, 2, \dots, M\}$, and $R = \{r_{i,j} \in \mathcal{R} | j = 1, \dots, M\}$ are the sets of candidate bounding boxes, entity class labels, and relation class labels, respectively. In this work, we consider $|\mathcal{O}| = 150$ most frequent entity classes and $|\mathcal{R}| = 51$ (fifty plus no-relation) most frequent relation classes appearing in the VG dataset. Given an image I , one would like to model $P(\mathcal{S}|I)$ which we decompose into three components, i.e. $P(\mathcal{S}|I) = P(R|O, B, I)P(O|B, I)P(B|I)$ as illustrated in Figure 1. The **entity proposal** component $P(B|I)$ is to generate the bounding box proposals, i.e. the locations and sizes of entities. The **entity classification** component $P(O|B, I)$ models the entity class probabilities of each proposal in B . The **relation classification** component $P(R|O, B, I)$ is to model the relation class probabilities of every two proposals in B along with their entity classes O . Our relation classification component capitalizes on knowledge distillation (KD) and aims at reducing biases in the relation classifiers. We describe the entity proposal and classification components in Section 3.1 and propose our relation classification component and two KD schemes in Section 3.2.

3.1 Entity Proposal and Classification

Entity Proposal Component $P(B|I)$. We model $P(B|I)$ with Faster R-CNN [25] entity detector. We use ResNeXt-101-FPN [60] (following [29]) as the backbone network to extract the bounding boxes B_i and their visual representations $\{\mathbf{f}_i \in \mathbb{R}^{4096} | i = 1, \dots, M\}$. The same backbone network is used to extract the visual representations $\{\mathbf{f}_{i,j} \in \mathbb{R}^{4096} | i, j = 1, \dots, M\}$ from the union region of every two proposals (i.e. the smallest bounding box that covers the two proposals) for the later use (see Section 3.2.2). As the entity proposal component, we use the model pre-trained on entity proposal labels from VG [25] and freeze it when training

two following components.

Entity Classification Component $P(O|B,I)$. This component comprises a *context module* that takes $\{\mathbf{f}_i | i = 1, \dots, M\}$ and outputs the *entity context*² $\{\mathbf{c}_i\}$ ($\mathbf{c}_i \in \mathbb{R}^{512}$) with which $P(O|B,I)$ is modeled. We adopt the context module proposed in either Motifs [85] or VC-Tree [28, 29] (other choices [9, 63] are also reasonable). Given $\{\mathbf{f}_i\}$, Motifs models \mathbf{c}_i with a bi-directional LSTM, while VCtree learns a hierarchical tree structure on which a bi-directional TreeLSTM [26] is constructed. The entity contexts capture the structured (e.g. chained or hierarchical) interactions among the visual entities. The classification component is trained on the standard cross-entropy loss (i.e. L_c) given the entity class labels.

3.2 Relation Classification with Two Learners

Our proposed relation classification component modeling $P(R|O,B,I)$ consists of two relation learners, a main supervised learner \mathcal{F} and an auxiliary semi-supervised learner \mathcal{G} (see Figure 1). \mathcal{F} learns on both the relation labels and \mathcal{G} 's knowledge. This two-learner scheme is introduced to alleviate the biases from the dataset (discussed in Section 1) brought into the models. We elaborate the two learners respectively in the following sections.

3.2.1 Supervised Learner \mathcal{F}

Here we model $P(r \in \mathcal{R} | o_i, o_j)$ for the supervised learner \mathcal{F} , in which o_i, o_j represent the entity classes for the i^{th} and j^{th} proposals. We first separately learn the representations for subjects o_i and objects o_j as $\mathbf{c}_i^{\mathcal{F}} = W_s^{\mathcal{F}} \mathbf{c}_i$ and $\mathbf{c}_j^{\mathcal{F}} = W_o^{\mathcal{F}} \mathbf{c}_j$, respectively with their entity contexts \mathbf{c}_i and \mathbf{c}_j . $W_s^{\mathcal{F}}, W_o^{\mathcal{F}} \in \mathbb{R}^{512 \times 512}$ are the parameter matrices for subjects and objects, and $W^{\mathcal{F}}$ is that for the relations. Now, the relation class probabilities are modeled as

$$P(r \in \mathcal{R} | o_i, o_j) = \mathbf{p}_{i,j} = \text{softmax}(\gamma \cdot \sigma_{L_2}(W^{\mathcal{F}} \sigma_{L_2}(\mathbf{g}_{i,j}^{\mathcal{F}}) + \mathbf{s}_{o_i, o_j})), \quad W^{\mathcal{F}} \in \mathbb{R}^{|\mathcal{R}| \times 4096} \quad (1)$$

$$\mathbf{g}_{i,j}^{\mathcal{F}} = W_c^{\mathcal{F}} [\mathbf{c}_i^{\mathcal{F}}, \mathbf{c}_j^{\mathcal{F}}] \odot \mathbf{f}_{i,j}, \quad W_c^{\mathcal{F}} \in \mathbb{R}^{4096 \times 1024} \quad (2)$$

where $\sigma_{L_2}(\cdot)$ L_2 -normalizes the vector or each row of the matrix, $[,]$ denotes concatenation of two vectors, and \odot denotes element-wise multiplication. \mathbf{s}_{o_i, o_j} is the (constant) frequency prior of $P(r | o_i, o_j)$ pre-calculated from the training set [85]. γ is the radius of the hyper-spherical space [28]. Later in Section 4.3, we will show that L_2 -normalization of both $W^{\mathcal{F}}$ and $\mathbf{g}_{i,j}^{\mathcal{F}}$ consistently improves the predictions on the minority relation classes.

The per-image loss for \mathcal{F} is then given by

$$L_{\mathcal{F}} = -\frac{1}{|R|} \sum_{r_{i,j} \in R} \log \mathbf{p}_{i,j}[r_{i,j}], \quad (3)$$

where $r_{i,j}$ is the target relation class for o_i, o_j , and $\mathbf{p}_{i,j}[r_{i,j}] = P(r_{i,j} | o_i, o_j) \in [0, 1]$ denotes the $r_{i,j}^{\text{th}}$ element of $\mathbf{p}_{i,j}$, i.e. the probability of relation $r_{i,j}$. Learning \mathcal{F} purely with Eq. (3) may impose undesirable biases into the classifier. $P(r_{i,j} | o_i, o_j)$ is encouraged to be maximized (and $P(r \in \mathcal{R}, r \neq r_{i,j} | o_i, o_j)$ to be suppressed) despite that (1) (o_i, o_j) contains a valid but not annotated relation (the aFN cases) and (2) $r \neq r_{i,j}$ could be valid relations between o_i, o_j (the aPTP cases). We address the biases stemming from these two cases with the other classifier \mathcal{G} along with the proposed two KD schemes next.

²referred as *object context* in [28, 63] or contextual cues in [9].

3.2.2 Semi-supervised Learner \mathcal{G} and Knowledge Distillation

Semi-supervised Learner \mathcal{G} . We co-train the other classifier \mathcal{G} that learns in a semi-supervised fashion. Instead of learning on \mathcal{R} , \mathcal{G} learns on $\mathcal{R}' = \mathcal{R} \setminus \{\text{no-relation}\}$ while being encouraged to detect as many relations as possible in $\{\text{no-relation}\}$. This is to address the fact that no-relation is an artificial class to accommodate the unannotated entity pairs. Specifically, similar to \mathcal{F} , given $\mathbf{c}_i, \mathbf{c}_j$, \mathcal{G} learns a different set of subject and object representations (for establishing a different perspective from that of \mathcal{F}) as $\mathbf{c}_i^{\mathcal{G}} = W_s^{\mathcal{G}} \mathbf{c}_i$ and $\mathbf{c}_j^{\mathcal{G}} = W_o^{\mathcal{G}} \mathbf{c}_j$, respectively, where $W_s^{\mathcal{G}}, W_o^{\mathcal{G}} \in \mathbb{R}^{512 \times 512}$. Next, the classifier is given by

$$P(r \in \mathcal{R}' | o_i, o_j) = \mathbf{q}_{i,j} = \text{softmax}(\gamma \cdot \sigma_{L_2}(W^{\mathcal{G}}) \sigma_{L_2}(\mathbf{g}_{i,j}^{\mathcal{G}})), \quad W^{\mathcal{G}} \in \mathbb{R}^{|\mathcal{R}'| \times 4096} \quad (4)$$

$$\mathbf{g}_{i,j}^{\mathcal{G}} = W_c^{\mathcal{G}} [\mathbf{c}_i^{\mathcal{G}}, \mathbf{c}_j^{\mathcal{G}}] \odot \mathbf{f}_{i,j}, \quad W_c^{\mathcal{G}} \in \mathbb{R}^{4096 \times 1024} \quad (5)$$

Comparing Eqs. (1) and (4), one can notice that \mathcal{G} learns without the prior \mathbf{s}_{o_i, o_j} . We argue that the prior, on the one hand, offers a strong baseline for the classifier to learn from, but on the other hand, can bias the classifier towards the majority classes as \mathbf{s}_{o_i, o_j} statistically speaks more for them. In addition, adding the same prior as \mathcal{F} to \mathcal{G} can end up learning two classifiers of similar behavior.

The loss function $L_{\mathcal{G}}$ for \mathcal{G} is defined over disjoint sets (1) $R^L = \{(i, j) | r_{i,j} \in \mathcal{R}'\}$, the set of the entity pairs annotated with a relation in \mathcal{R}' and (2) $R^U = \{(i, j) | r_{i,j} \notin \mathcal{R}'\}$, the set of the unannotated entity pairs, in which $|R^L \cup R^U| = |R|$,

$$L_{\mathcal{G}} = -\frac{1}{|R^L|} \sum_{(i,j) \in R^L} \log \mathbf{q}_{i,j}^L[r_{i,j}] - \lambda_{\mathcal{G}} \frac{1}{|R^U|} \sum_{(i,j) \in R^U} \mathcal{H}(\mathbf{q}_{i,j}^U), \quad (6)$$

where $\mathbf{q}_{i,j}^L = \mathbf{q}_{i,j}$, $\forall (i, j) \in R^L$, and $\mathbf{q}_{i,j}^U = \mathbf{q}_{i,j}$, $\forall (i, j) \in R^U$. The role of $\mathcal{H}(\mathbf{q}_{i,j}^U)$, the entropy of the probability distribution $\mathbf{q}_{i,j}^U$ is to be maximized so that \mathcal{G} is encouraged to detect as many relations as possible. $\lambda_{\mathcal{G}}$ is a hyperparameter.

Knowledge Distillation. \mathcal{G} transfers its knowledge to \mathcal{F} through a knowledge distillation (KD) loss $L_{\mathcal{G} \rightarrow \mathcal{F}}$,

$$L_{\mathcal{G} \rightarrow \mathcal{F}} = \lambda_{\mathcal{G} \rightarrow \mathcal{F}} \left(\sum_{(i,j) \in R^L} w_{i,j}^L KL(\tilde{\mathbf{q}}_{i,j}^L || \mathbf{p}_{i,j}^L) + \sum_{(i,j) \in R^U} w_{i,j}^U KL(\tilde{\mathbf{q}}_{i,j}^U || \mathbf{p}_{i,j}^U) \right), \quad (7)$$

where $\lambda_{\mathcal{G} \rightarrow \mathcal{F}}$ is a hyperparameter and $KL(\cdot || \cdot)$ denotes the Kullback-Leibler divergence between two distributions. $\mathbf{p}_{i,j}^*$ denotes the first $|\mathcal{R}'|$ elements of $\mathbf{p}_{i,j}^*$ calculated via Eq. (1) for $(i, j) \in R^*$, where $*$ $\in \{L, U\}$. $\tilde{\mathbf{q}}_{i,j}^*$ considers *temperature scaling* before transferring the knowledge [□], i.e.

$$\tilde{\mathbf{q}}_{i,j}^* = \text{softmax}(\gamma \cdot \sigma_{L_2}(W^{\mathcal{G}}) \sigma_{L_2}(\mathbf{g}_{i,j}^{\mathcal{G}}) / T), \quad \forall (i, j) \in R^*, \quad * \in \{L, U\}, \quad (8)$$

where T is the temperature hyperparameter. Note that Eq. (8) reduces to Eq. (4) when $T = 1$. In Eq. (7), $w_{i,j}^*$ serves as a normalization factor for which we propose two different forms: (1) the *uniform KD* (uKD), i.e. $w_{i,j}^* = 1/|R^*|$ and (2) the *certainty-based KD* (cKD), i.e. $w_{i,j}^* = \sum_{(m,n) \in R^*} \mathcal{H}(\mathbf{q}_{m,n}^*) / \mathcal{H}(\mathbf{q}_{i,j}^*)$. One can see that a more certain $\mathbf{q}_{i,j}^*$ (smaller $\mathcal{H}(\mathbf{q}_{i,j}^*)$) yields a larger $w_{i,j}^*$, i.e. \mathcal{F} is forced to absorb more knowledge from the relatively confident predictions of \mathcal{G} . We do not propagate gradients from \mathcal{G} to \mathcal{F} to avert \mathcal{F} affecting \mathcal{G} because, otherwise, one would end up obtaining two classifiers that make similar predictions.

The combined loss L from entity classification (see Section 3.1) and the three relation classification losses is then $L = L_c + L_{\mathcal{F}} + L_{\mathcal{G}} + L_{\mathcal{G} \rightarrow \mathcal{F}}$.

4 Experiments

4.1 Dataset, Tasks, Metrics

Dataset. We evaluate the SGG models on the VG dataset [16]. We follow the same training, validation and test splits as in [2, 3, 28, 29, 35], where only the most frequent $|\mathcal{O}| = 150$ entity classes and $|\mathcal{R}| = 50$ relation classes are considered.

Tasks. The SGG tasks, from the easiest to the hardest setups, are (1) *predicate classification* (PrdCls), (2) *scene graph classification* (SGCls), and (3) *scene graph detection* (SGDet) [35]. In the PrdCls task, the entity proposals and class labels are provided to the model for predicting the relations among entity pairs. In the SGCls task, only the entity proposals are given and models are used to predict both the entity classes and relations. In the SGDet task, models are used to generate entity proposals and predict their classes and relations. Among all three tasks, the models are evaluated either *with* or *without the graph constraint* [2, 3, 35, 36]. Models with the constraint predict a single relation with the highest possibility for each entity pair. Those without the constraint are allowed to predict multiple relations for each entity pair.

Metrics. For the tasks with the constraint, we consider recall@K (R@K) and the mean class-wise recall@K (mR@K), where $K = \{20, 50, 100\}$. For those without the constraint, we consider $R@\{50, 100\}$ and $mR@\{50, 100\}$. mR@K has been emphasized more in the most recent SGG works [2, 28, 29] to address the imbalanced relation class distributions. R@K accounts for the proportion of the top K confident predicted relation triplets that are in the ground-truth relation triplets. mR@K calculates the average of all R@Ks, each of which is computed separately for the triplets consisting of each relation class. Other metrics, e.g. precision@K, false positive and negative rates etc., are not adopted and neither in other SGG works [2, 3, 35, 36] because they are considered less fair in evaluating the methods performed on the sparsely annotated dataset such as VG.

4.2 Training Settings

We separately train different models for SGDet, SGCls, and PrdCls on top of the pre-trained Faster R-CNN with ResNeXt-101-FPN backbone (provided by the authors in [29]). The SGG models is subject to the performance of entity detector, therefore, we choose a more complex backbone network to align with one of the most recent SGG works [29], however, a direct comparison to other works [2, 3, 28, 35] seems difficult as they use an earlier VGG16 backbone. Motifs [35] and VCTree [28] are the two baselines for which we adapt our proposed relation component.

Hyperparameters. The radius hyperparameter $\gamma = 12$ in Eqs. (1,4,8), temperature $T = 1.5$ in Eq. (8), $\lambda_G = 0.1$ in Eq. (6), and $\lambda_{G \rightarrow \mathcal{F}} = 0.1$ in Eq. (7) are selected according to the evaluations on the validation set.

Training Details. All models are trained up to 50,000 batch iterations of 12 images using SGD with 0.9 momentum. KD starts only after 10,000 iterations ($\lambda_{G \rightarrow \mathcal{F}} = 0$ from 0–10,000 iterations). Learning rate starts from 0.01 and is linearly increased every batch iteration up to 0.12 until 500 iterations (i.e. the warm-up period). After the warm-up, the learning rate is decayed by 0.2 (0.1) for Motifs (VCTree) models if no improvement in R@100 is observed within two successive validation rounds (2,000 iterations / round). The training is early terminated once the learning rate is decayed for three times.

w/ constraint Model	SGDet R@{20,50,100}			mR@{20,50,100}			SGCls R@{20,50,100}			mR@{20,50,100}			PrdCls R@{20,50,100}			mR@{20,50,100}		
Motifs† [▣]	21.4	27.2	30.3	4.2	5.7	6.6	32.9	35.8	36.5	6.3	7.7	8.2	58.5	65.2	67.1	10.8	14.0	15.3
VCtree† [▣]	22.0	27.9	31.3	5.2	6.9	8.0	35.2	38.1	38.8	8.2	10.1	10.8	60.1	66.4	68.1	14.0	17.9	19.4
Routing† [▣]	-	27.1	29.8	-	6.4	7.3	-	36.7	37.4	-	9.4	10.0	-	65.8	67.6	-	17.7	19.2
Critic† [▣]	22.1	27.9	31.2	-	-	-	35.9	39.0	39.8	-	-	-	60.2	66.4	68.1	-	-	-
Motifs* [▣]	25.48	32.78	37.16	4.98	6.75	7.90	35.63	38.92	39.77	6.68	8.28	8.81	58.46	65.18	67.01	11.67	14.79	16.08
Motifs baseline	25.78	33.08	37.57	5.37	7.36	8.62	36.13	39.35	40.14	7.32	8.94	9.50	59.15	65.57	67.29	13.01	16.63	17.89
Motifs L2	25.86	33.22	37.69	5.68	7.77	9.15	36.06	39.21	39.98	8.01	9.80	10.43	59.33	65.76	67.51	13.78	17.55	18.99
Motifs L2+uKD	24.80	32.24	36.75	5.68	7.88	9.53	35.11	38.50	39.32	8.57	10.90	11.81	57.37	64.14	65.95	14.19	18.59	20.32
Motifs L2+cKD	25.20	32.50	37.08	5.81	8.05	9.58	35.58	38.93	39.75	8.66	10.71	11.39	57.73	64.58	66.44	14.36	18.49	20.22
VCtree* [▣]	24.53	31.93	36.21	5.38	7.44	8.66	42.77	46.67	47.64	9.59	11.81	12.52	59.02	65.42	67.18	13.12	16.74	18.16
VCtree baseline	25.15	32.23	36.32	5.31	7.29	8.46	41.48	45.10	46.04	8.54	10.48	11.17	59.80	65.87	67.49	13.33	16.89	18.23
VCtree L2	25.13	32.20	36.36	5.29	7.22	8.43	41.86	45.54	46.48	8.85	10.90	11.61	59.76	65.90	67.51	13.68	17.25	18.65
VCtree L2+uKD	24.44	31.60	35.90	5.68	7.69	9.21	40.93	44.66	45.60	9.90	12.43	13.40	58.49	65.01	66.71	14.20	18.24	19.92
VCtree L2+cKD	24.84	32.02	36.12	5.66	7.73	9.06	41.41	45.17	46.11	9.65	12.14	13.11	59.04	65.42	67.07	14.41	18.42	20.03

w/o constraint Model	SGDet R@{50,100}		mR@{50,100}		SGCls R@{50,100}		mR@{50,100}		PrdCls R@{50,100}		mR@{50,100}	
Motifs† [▣]	30.5	35.8	-	-	44.5	47.7	-	-	81.1	88.3	-	-
VCtree† [▣]	-	-	-	-	-	-	-	-	-	-	-	-
Routing† [▣]	30.9	35.8	11.7	16.0	45.9	49.0	19.8	26.20	81.9	88.9	36.3	49.0
Critic† [▣]	31.6	36.8	-	-	48.6	52.0	-	-	83.2	90.1	-	-
Motifs* [▣]	36.58	43.43	-	-	48.48	51.98	-	-	81.02	88.24	-	-
Motifs baseline	37.12	43.97	12.63	17.07	49.03	52.61	18.66	25.01	81.95	89.05	33.28	45.53
Motifs L2	37.23	44.05	13.54	18.34	48.97	52.52	20.70	27.66	82.22	89.29	36.61	49.52
Motifs L2+uKD	35.98	42.98	13.95	19.47	48.53	52.12	22.73	30.09	80.02	87.18	36.92	50.89
Motifs L2+cKD	36.27	43.15	14.22	19.78	48.63	52.17	22.07	29.64	80.29	87.49	37.15	50.83
VCtree* [▣]	35.73	42.34	-	-	58.36	62.70	-	-	81.63	88.83	-	-
VCtree baseline	36.25	42.90	12.42	16.98	56.46	60.59	22.26	29.84	82.56	89.43	34.45	46.82
VCtree L2	36.22	42.71	12.53	16.94	56.94	61.12	22.82	30.80	82.69	89.54	35.12	47.61
VCtree L2+uKD	35.48	42.05	13.82	19.13	55.97	59.97	26.83	35.19	81.43	88.46	37.69	51.72
VCtree L2+cKD	35.88	42.38	13.93	19.01	56.61	60.67	26.80	35.82	81.94	88.83	38.35	52.42

Table 1: Comparisons with the state of the art. Motifs [▣] and VCtree [▣] baselines are established to be compared. Models with † refer to those with VGG16 as the backbone network while the other moare with ResNeXt-101-FPN. Models with * refer to the results reported by the authors of [▣]. They are excluded from the comparison due to some modifications to the Motifs/VCtree baseline on which our proposed models are based.

4.3 Experimental Results and Comparison with the State of the Art

We conduct extensive quantitative comparisons between our proposed models and the state-of-the-art in Table 1 and Figure 2 to see if the proposed scheme can be used to improve the existing SGG models. Motifs/VCtree baseline and L2 differ only in that baseline (which is the re-implementation of the original Motifs/VCtree provided in [▣]) does not L_2 -normalize the row vectors in W^F and the input features $\mathbf{g}_{i,j}^F$ in Eq. (1) while the latter does. Neither baseline nor L2 involves learning \mathcal{G} . L2+uKD and L2+cKD adopt uniform and certainty-based KD schemes, respectively (see Section 3.2.2). We inspect the relative improvement (%), i.e. $(R_A - R_B)/R_B \times 100\%$, when comparing the recall R_A of a proposed model A to R_B of a baseline model B , and similarly for mean class-wise recall mR.

L2 vs. baseline. We observe consistent improvements in mR@K across the three SGG tasks in almost all metrics, e.g. +6.15% in mR@100 with Motifs L2 in SGDet with the graph constraint. The changes in R@K across all tasks (with and without the constraint) are nearly negligible (-1% to 1%).

L2+cKD vs. L2+uKD. L2+cKD performs better than L2+uKD in R@K across all tasks with and without the constraint. However, L2+uKD can sometimes slightly outperforms L2+cKD in mR. For instance, VCtree L2+uKD leads L2+cKD by 0.49% in mR@K but is 1.19% behind in R@K on average in SGDet with the constraint. Although cKD falls short in mR@K in some cases when evaluated with the graph constraint (first row, Figure 2(a)), the difference in mR@K obtained by cKD and uKD usually becomes almost negligible without the graph constraint (second row, Figure 2(a)). Overall, we favor L2+cKD over L2+uKD models here since the former usually produce more balanced results in R@K and mR@K (see Figure 2). Hence, we mainly compare L2+cKD against other models from here on.

L2+cKD vs. L2. L2+cKD significantly improves mR@K at a negligible cost of R@K, e.g.




Figure 2: (a) Relative improvements (%) of models against their baselines, Motifs [65] and VCTree [28] in R@100 and mR@100. Two baselines, black hollow \triangle and \circ , serve as reference points for two groups of experiments, respectively. (b) Class-wise recall comparison of VCTree baseline and L2+cKD shows clear improvement in the minority classes.

VCTree L2+cKD outperforms VCTree L2 by 7.18% in mR@K with only -0.79% loss in R@K on average in SGDet with the constraint. The relative improvements in mR@K to R@K are even higher in SGCLS. This indicates that L2+cKD is more effective in reducing the biases from the imbalanced distributions in VG than L2.

Improvements in mR@K vs. R@K. The relative improvements of L2+cKD against baseline in mR@K (6.03% – 20.39%) are much higher than the losses (at worst, -1.21%) in R@K across all tasks (see Figure 2(a)) with or without the constraint. The main losses in the majority classes are usually negligible compared to the gains in the minority classes (see Figure 2(b)). This indicates that our proposed scheme is more capable of modeling rare relation classes, e.g. on back of, lying on, and playing etc.

Improvements in models w/ constraint vs. w/o constraint. In SGDet, our proposed models show much more improvements in mR@50 and mR@100 when not under the graph

constraint than under it. This suggests that, when not constrained, our models are more able to retrieve more relevant relations (and ignore the less relevant) than baseline.

Ablation study on temperature T in Eq. (8). From Table 2, we can see that one can attain a higher mR@K with a larger T in the PrdCls task with VCTree L2+cKD. However, the drop in R@K can become deeper in exchange of lower improvements in mR@K. This can be explained by the fact that the learner \mathcal{G} can provide less informative (i.e. more uniform) predictions with a larger T that overly regularizes \mathcal{F} 's predictions.

	with constraint, PrdCls			w/o constraint, PrdCls						
	R@{20,50,100}		mR@{20,50,100}	R@{50,100}		mR@{50,100}				
$T=1.0$	58.74	65.68	67.54	13.66	17.75	19.37	81.81	89.02	36.25	49.52
$T=1.25$	58.76	65.55	67.39	13.6	18.15	19.94	81.74	88.98	37.23	51.31
$T=1.5$	59.04	65.42	67.07	14.41	18.42	20.03	81.94	88.83	38.35	52.42
$T=1.75$	58.86	65.32	66.96	14.36	18.5	20.18	81.52	88.31	38.88	52.48

Table 2: Ablation study on temperature hyperparameter T in Eq. (8) on VCTree L2+cKD.

5 Conclusion

We proposed a novel SGG framework consisting of two relation learners in which one learner regularizes the other learner’s predictions with the less biased knowledge. This particularly addresses the biases incurred by the incomplete annotations in VG. The quantitative analyses showed that our proposed models are less biased towards the majority relation classes and more capable of capturing more than one relevant relations. Empirically we improved two state-of-the-arts, Motifs [65] and VCTree [28] with (and without) the graph constraint, respectively by 11.14% (14.23%) and 7.09% (12.06%) in SGDet, by 19.89% (18.39%) and 17.37% (20.22%) in SGCls, and by 13.02% (11.64%) and 9.87% (11.96%) in PrdCls in mR@100. R@100 results are -1.3% (-1.86%) at worst across all tasks with (without) the constraint. We believe that the improvements come from the fact that our proposed models alleviate the impact of the biases from the partially annotated and unannotated samples in the popular VG dataset. One of the future directions is to investigate the generalization ability of our models on the rarely seen (few-shot [4]) or unseen (zero-shot [21]) subject-relation-object triplets.

Acknowledgements

This work has been funded by the Academy of Finland project number 313988 (DeepGraph), and the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 780069 (MeMAD). We also thank the Aalto University’s Aalto Science IT project and CSC – IT Center for Science Ltd. for providing computer resources and NVIDIA Corporation for donation of GPU for this research.

References

- [1] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems*, pages 1565–1576, 2019.
- [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 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 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 International Conference on Computer Vision*, pages 2580–2590, 2019.
- [5] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019.
- [6] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019.
- [7] Apoorva Dornadula, Austin Narcomey, Ranjay Krishna, Michael Bernstein, and Fei-Fei Li. Visual relationships as functions: Enabling few-shot scene graph prediction. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [8] Lianli Gao, Zhao Guo, Hanwang Zhang, Xing Xu, and Heng Tao Shen. Video captioning with attention-based lstm and semantic consistency. *IEEE Transactions on Multimedia*, 19(9):2045–2055, 2017.
- [9] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018.
- [10] Jiuxiang Gu, Handong Zhao, Zhe Lin, Sheng Li, Jianfei Cai, and Mingyang Ling. Scene graph generation with external knowledge and image reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2019.
- [11] Sangchul Hahn and Heeyoul Choi. Self-knowledge distillation in natural language processing. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 423–430, Varna, Bulgaria, September 2019. INCOMA Ltd. doi: 10.26615/978-954-452-056-4_050. URL <https://www.aclweb.org/anthology/R19-1050>.
- [12] Munawar Hayat, Salman Khan, Syed Waqas Zamir, Jianbing Shen, and Ling Shao. Gaussian affinity for max-margin class imbalanced learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6469–6479, 2019.

- [13] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. URL <http://arxiv.org/abs/1503.02531>.
- [14] Salman Khan, Munawar Hayat, Syed Waqas Zamir, Jianbing Shen, and Ling Shao. Striking the right balance with uncertainty. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 103–112, 2019.
- [15] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems*, 29(8): 3573–3587, 2017.
- [16] 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.
- [17] 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.
- [18] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheraface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017.
- [19] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019.
- [20] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [21] 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.
- [22] Hossein Mobahi, Mehrdad Farajtabar, and Peter L Bartlett. Self-distillation amplifies regularization in hilbert space. *arXiv preprint arXiv:2002.05715*, 2020.
- [23] Varun K Nagaraja, Vlad I Morariu, and Larry S Davis. Modeling context between objects for referring expression understanding. In *European Conference on Computer Vision*, pages 792–807. Springer, 2016.
- [24] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan Yuille. Deep co-training for semi-supervised image recognition. In *Proceedings of the european conference on computer vision (eccv)*, pages 135–152, 2018.
- [25] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.

- [26] Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.
- [27] Jiayi Tang, Rakesh Shivanna, Zhe Zhao, Dong Lin, Anima Singh, Ed H Chi, and Sagar Jain. Understanding and improving knowledge distillation. *arXiv preprint arXiv:2002.03532*, 2020.
- [28] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6619–6628, 2019.
- [29] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiayin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. *arXiv preprint arXiv:2002.11949*, 2020.
- [30] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [31] 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.
- [32] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015.
- [33] 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.
- [34] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. MATTNet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1307–1315, 2018.
- [35] 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.
- [36] Ji Zhang, Kevin J Shih, Ahmed Elgammal, Andrew Tao, and Bryan Catanzaro. Graphical contrastive losses for scene graph parsing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11535–11543, 2019.
- [37] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3713–3722, 2019.
- [38] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems*, 30(11):3212–3232, 2019.