

Addressing Class Imbalance in Scene Graph Parsing by Learning to Contrast and Score

He Huang¹, Shunta Saito², Yuta Kikuchi², Eiichi Matsumoto², Wei Tang¹, and Philip S. Yu¹

¹ University of Illinois at Chicago, Chicago, USA
 {hehuang, tangw, psyu}@uic.edu


² Preferred Networks Inc., Tokyo, Japan
 {shunta, kikuchi, matsumoto}@preferred.jp

Abstract. Scene graph parsing aims to detect objects in an image scene and recognize their relations. Recent approaches have achieved high *average* scores on some popular benchmarks, but fail in detecting rare relations, as the highly long-tailed distribution of data biases the learning towards frequent labels. Motivated by the fact that detecting these rare relations can be critical in real-world applications, this paper introduces a novel integrated framework of classification and ranking to resolve the class imbalance problem in scene graph parsing. Specifically, we design a new Contrasting Cross-Entropy loss, which promotes the detection of rare relations by suppressing incorrect frequent ones. Furthermore, we propose a novel scoring module, termed as Scorer, which learns to rank the relations based on the image features and relation features to improve the recall of predictions. Our framework is simple and effective, and can be incorporated into current scene graph models. Experimental results show that the proposed approach improves the current state-of-the-art methods, with a clear advantage of detecting rare relations.


1 Introduction

As an extension to object detection [4], scene graph parsing [5, 6] aims to detect not only objects, *e.g.*, *persons* and *bikes*, in an image scene but also recognize their relationships (also called *predicates*), *e.g.*, *ride* and *push*. It is a fundamental tool for several applications such as image captioning [7], image retrieval [8] and image generation [9]. Due to the combinatorially large space of valid relation triplets $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ and the polysemy of a predicate in different contexts, the task of scene graph parsing is challenging.

Recent scene graph parsing systems [5, 3, 2, 10–12] are built on deep neural networks due to their ability to learn robust feature representations for both images and relational contexts directly from data. Xu *et al.* [5] use GRU [13] to approximate a conditional random field to jointly infer objects and predicates. Zellers *et al.* [3] propose to utilize the statistical distribution of relation triplets as external prior knowledge. Yang *et al.* [11] propose to incorporate a GCN [14] for predicting the predicates.



(a) Distribution of the 50 most frequent predicates in VG [1].



(b) Per-class Recall@100 of LinkNet [2] and MotifNet [3] evaluated on the predicate classification task.

Fig. 1. Motivation of this work: (a) the predicate distribution is highly long-tailed, and (b) two current state-of-the-art models [3, 2] perform poorly on rare relations.

While these approaches can achieve high *average* scores on some popular benchmarks, their effectiveness is largely jeopardized in case of rare relationships. For example, Figure 1(a) and Figure 1(b) respectively show the distribution of predicates in the Visual Genome dataset (VG) [1] and the performances of two state-of-the-art methods, *i.e.*, LinkNet [2] and MotifNet [3]. Both approaches perform quite well on frequent relations, *e.g.*, **on** and **has**, but their performances degrade significantly on less common predicates, *e.g.*, **lying on** and **flying in**. The recalls of several rare relations are even zero. Actually, the average recalls of each class (*i.e.*, *macro-averaged Recall@100*) are only 15.3% and 16.9% for MotifNet [3] and LinkNet [2] respectively. By contrast, they can achieve around 67% *micro-averaged Recall@100*, which is calculated in a class-agnostic way. In other words, a high micro-averaged recall is achieved at the cost of sacrificing the performance on rare classes.

This is because the distribution of the predicates is highly long-tailed, and learning therefore tends to bias towards frequent relations so as to boost the overall performance. However, we argue that detecting those rare or *abnormal* relations is critical in practice. For example, consider a common predicate **on** and a rare one **lying on** in the VG dataset. Although these two terms are closely related, it is crucial for a surveillance camera to distinguish between **<person,**

on, street> and <person, lying on, street>, since the latter may imply an emergency that requires immediate reactions.

Another limitation of previous approaches is that they model the task as a pure classification problem, by training with two cross-entropy loss functions for **subjects/objects** and **predicates** respectively. This approach is problematic as the evaluation metric is recall@K, which not only depends on classification accuracy but also is sensitive to the ranks of predicted relation triplets. While an object may either exist or not in an image, there is always at least one relationship between any two existing objects, *e.g.*, **left** or **right**. When people annotate the VG [1] dataset, they only label the most important or nontrivial relations but ignore the others. As a result, the detection of an existing but trivial relation will be counted as a *false positive* because it is unlabeled in the dataset. To achieve a high recall for a given number of predictions, we need to not only classify the relations correctly but also rank the nontrivial relations higher than the trivial ones.

To overcome these challenging issues, this paper introduces a novel method which solves predicate classification and relation ranking in a unified framework. We first introduce a loss function, termed as the Contrasting Cross-Entropy (CCE) loss, to handle the class imbalance problem. This is achieved by simultaneously maximizing the predicted probability of the correct class and minimizing that of the hardest negative class. By suppressing incorrect frequent relations, the CCE loss promotes the recall of rare predicates. Furthermore, we propose a novel network module, called Scorer, to tackle the ranking problem. It scores a predicted relation triplet by comparing it with all other triplets as well as exploiting the global context information. Ablation study indicates that (1) the CCE loss significantly improves the detection of rare relationships and (2) the Scorer network helps boost the recall by learning to rank. Experimental results show that the proposed CCE loss and Scorer network can improve the state-of-the-art methods on three tasks. The contributions of this paper are:

- We provide a brand-new perspective on Scene Graph Parsing (SGP). To our knowledge, this is the first attempt to formulate it as a joint task of classification and ranking.
- We propose a simple but effective Contrasting Cross-Entropy (CCE) loss to alleviate the class imbalance problem. By contrasting for each relation the predicted probabilities of the correct label and the hardest incorrect one, it suppresses the incorrect frequent relations and significantly improves the detection performance on rare classes.
- We introduce a novel Scorer network to tackle the ranking problem and improve the recall of baselines by a large margin. It innovatively bridges the point-wise and list-wise ranking approaches to take advantages of both. Specifically, our Scorer exploits a self-attention layer to compare one relation triplet with all the others in a listwise fashion, from which learning to rank is achieved via a point-wise cross-entropy loss. To our knowledge, this unification is unique and novel.

- Our novel framework is general and light-weighted. It can be easily applied to any scene graph models trained with cross-entropy. We conduct extensive experiments and show that our combined CCE-Scorer framework with LinkNet [2] as the base scene graph model achieves new state-of-the-art results on three scene graph detection tasks.

2 Related Work

Scene Graph Parsing. There have been several lines of research on scene graph parsing [5, 15, 16] and visual relationship detection [6, 17] over the past few years. IMP [5] uses GRUs [13] to approximate a CRF and proposes an iterative message passing mechanism. Zellers *et al.* [3] propose a powerful baseline that infers the relation between two objects by simply taking the most frequent relation between the object pair in training set. By incorporating the frequency baseline, MotifNet [3] uses bi-directional LSTMs[18] to decode objects and relations sequentially. Recently, Gu *et al.* [19] propose to extract commonsense knowledge of predicted objects from knowledge bases. Another line of work try to design more sophisticated message passing methods to fine-tune object and predicate features [20, 16, 21]. LinkNet [2] exploits a new Relation Embedding module based on the self-attention to refine object features, which are then used to construct and refine predicate features. Graph RCNN [11] performs reasoning over the scene graph via a Relation Proposal Network, and then applies a graph convolutional network [14] to refine features.

KERN [10] tries to solve the class imbalance problem and uses *mean recall@K* as evaluation metric that can better represent the model’s performance. However, KERN [10] tackles this problem by designing a complex network structure, while our proposed framework does not rely on particular network structure and is general and can be easily applied to any scene graph models that train with a cross-entropy loss, such as KERN [10] itself. Experiments show that our framework can improve the performance of KERN [10] by around 1% to 2%.

Zhan *et al.* [22] propose a module that classifies each predicted predicate as either determined or not, which is related to our proposed Scorer network. However, their module is simply a stack of fully-connected layers applied to each prediction individually and does not care about ranking, while ours utilizes information from both image-level and different relations. More importantly, Zhan *et al.* [22] do not consider ranking and only use their module in training, while our Scorer module is used also during inference to generate the ranks of predicted relation triplets.

Class Imbalance. This problem has been widely studied in both image classification [23, 24] and object detection [25]. The most basic approaches are over-sampling minor classes and under-sampling major classes, both of which are prone to overfit. Another common approach is to multiply the loss of each class by its inverse frequency in the dataset. Sampling hard negatives is also widely studied in object detection to address the imbalance problem [26, 27]. Recently, more advanced loss functions have been proposed to address this problem [28–

30]. However, the class imbalance problem in scene graph parsing is rarely studied [10]. To the best of our knowledge, we are among the first to study this problem in depth and propose a general attentive solution that significantly mitigates the problem.

Contrasting Learning. The contrasting mechanism has been widely used for various applications, by pushing the scores of positive samples or matched sample pairs higher than those of negative ones [31–34]. Different from these previous approaches, we contrast for each relation the predicted probabilities of the correct and incorrect classes to suppress the incorrect frequent relations. In other words, they contrast data samples while we contrast classes, and the goals are different. While [35, 36] also contrast classes, our design is different from theirs. They push down all negative classes to maximize the overall discriminativeness while we deliberately inhibit only the hardest negative class, which is critical to suppress the incorrect frequent relations. Besides, the contrasting loss in [35] includes a very complex normalization term while our CCE loss is much simpler and computationally efficient. [36] uses Gaussian as the distance metric while we use cross-entropy. Zhang *et al.* [37] propose a graphical contrastive loss that discriminate positive `<subject, object>` pairs from negative ones, which is different from us since our CCE loss contrast the classes of predicted relations. While all these work focus on one task, to the best of our knowledge, we are the first to apply a contrasting max-margin based loss together with a ranking objective in scene graph parsing.

Learning to Rank. Our work is also related to the field of *learning to rank* [38], which aims to predict the ranks of a set of samples in a non-decreasing order. There are three popular approaches to address this problem, *i.e.*, *point-wise* [39], *pair-wise* [40, 41, 39] and *list-wise* [42, 41]. Different from prior work, our proposed Scorer network innovatively bridges the point-wise and list-wise ranking approaches by exploiting a self-attention module to compare one relation triplet with all the others in a list-wise fashion, from which learning to rank is achieved via a point-wise cross-entropy loss.

3 Proposed Approach

3.1 Problem Definition

The task of scene graph parsing [5, 3] requires a model to take an image as input and then output an ordered set of relation triplets $\mathcal{T} = \{\mathcal{T}_i | i = 1, \dots, M\}$, where $\mathcal{T}_i = \langle \text{subject}, \text{predicate}, \text{object} \rangle$ and M is the number of relation triplets in the given image, each `subject/object` is from a set of object classes \mathcal{C} , and each `predicate` is from a set of predicate classes \mathcal{R} . The triplets in set \mathcal{T} are ranked according to their confidence predicted by the model.

Current scene graph models rely on pretrained object detectors, *e.g.*, Faster R-CNN [4], to obtain the bounding boxes of objects and their features. Given N object proposals from the object detector, the scene graph model predicts their object labels. For each pair of predicted objects, the model generates a relation

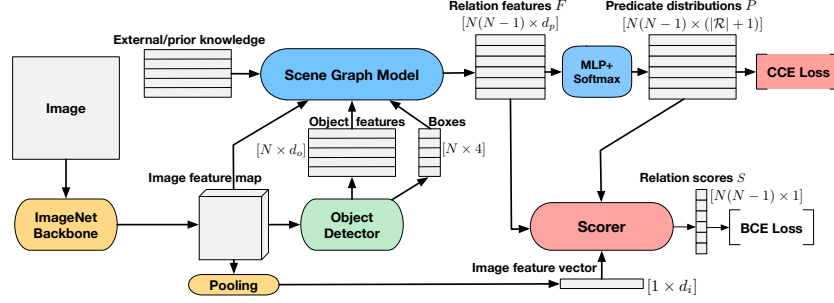


Fig. 2. An overview of the proposed framework. Here N is the number of detected objects in the input image, and (d_o, d_i, d_p) are the feature dimensions of (object, image, predicate). The core components are the Contrasting Cross-Entropy (CCE) loss for predicate classification and the Scorer network for relation ranking. Since these two components do not require modifying the chosen scene graph model, any existing scene graph model that trains with cross-entropy can be easily plugged into our framework. The external/prior knowledge is specific to the chosen base scene graph model. For example, MotifNet [3] and many other models [2, 37] use the probability distribution of predicates given the subjects and objects.

feature for this `<subject, object>` pair, which is then classified among the $|\mathcal{R}|$ predicate classes. Note that there are $N(N - 1)$ relation features in total, excluding self-relations. The classification of objects and predicates are trained using separate cross-entropy loss functions in most existing scene graph models [5, 3, 2, 10]. In practice, to handle un-annotated relations, an **unknown** class is introduced so that classification is performed among $|\mathcal{R}| + 1$ classes. During inference, the predicate’s label is found by ignoring the Softmax score of the **unknown** class and taking the label with the largest probability among the rest $|\mathcal{R}|$ classes.

After obtaining the predicted triplets \mathcal{T} and a given set of ground-truth triplets \mathcal{T}^* , the result is evaluated by calculating *macro-averaged recall@K*, also known as $mR@K$ [10]:

$$mR@K = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \frac{\# \text{ of correctly predicted } r \text{ in } \mathcal{T}[:K]}{\# \text{ of } r \text{ in } \mathcal{T}^*}, \quad (1)$$

where $\mathcal{T}[:K]$ represents the top- K triplets in the predicted \mathcal{T} . This metric is calculated per test image, and then the average $mR@K$ is used as the metric.

3.2 Overview of Framework

The overall framework of our proposed method is illustrated in Figure 2. The main pipeline is general among most scene graph models: an image is first fed into a backbone image classifier pretrained on ImageNet [43], and then the output image feature tensor is fed into an object detector pretrained on the Visual Genome

dataset [1] to obtain a set of N object proposals. By using any existing scene graph models as the base scene graph model, we obtain $N(N-1)$ relation features $F \in \mathbb{R}^{N(N-1) \times d_p}$ and their probability distributions $P \in \mathbb{R}^{N(N-1) \times (|\mathcal{R}|+1)}$ among the $|\mathcal{R}|+1$ predicate classes. These predicate distributions are used in our novel CCE loss for predicate classification, while the Scorer network takes the relation features, predicate distributions as well as the global image feature to generate the significance scores for each of the $N(N-1)$ relations. The scores are then used to rank the relations during inference.

3.3 Contrasting Cross Entropy Loss

As widely studied, classification models trained with a cross-entropy (CE) loss can be highly biased towards frequent classes and thus perform badly on rare ones. Methods such as Focal Loss (FL) [28] have been proposed to address this problem in object detection, but the class imbalance problem in scene graph parsing is less explored. In this section, we propose a simple but effective Contrasting Cross-Entropy (CCE) loss that can significantly improve mean recall@K and can be applied to any existing scene graph model using cross-entropy loss. Recall that the cross-entropy (CE) loss for one relation is defined as:

$$\mathcal{L}_{\text{CE}}(\mathbf{p}, \mathbf{y}) = - \sum_{i=1}^{|\mathcal{R}|+1} \mathbf{y}_i \log \mathbf{p}_i, \quad (2)$$

where \mathbf{p} is the probability distribution over the $|\mathcal{R}|+1$ predicate classes including the **unknown** class, and \mathbf{y} is the one-hot encoded ground-truth label vector.

As minimizing cross-entropy loss is equivalent to maximizing the log probability of the correct class, in order to handle the class imbalance problem, we also force the model to minimize the log probability of the hardest negative class, and design our Contrasting Cross-Entropy (CCE) loss as:

$$\mathcal{L}_{\text{CCE}}(\mathbf{p}, \mathbf{y}) = \max(\mathcal{L}_{\text{CE}}(\mathbf{p}, \mathbf{y}) - \mathcal{L}_{\text{CE}}(\mathbf{p}, \hat{\mathbf{y}}) + \alpha, 0), \quad (3)$$

where $\hat{\mathbf{y}}$ is the sampled hardest negative label, and α is a hyper-parameter to control the margin. Our intuition is that the hardest negative classes usually correspond to frequent relations. By suppressing the Softmax scores of incorrect frequent classes, this loss function can help the model perform better on rare classes. Here we choose the hardest negative label by first ignoring the **unknown** class (*i.e.* \mathbf{p}_1) and taking the incorrect class with the highest probability, *i.e.*, $c = \arg \max_{j>1, \mathbf{y}_j=0} \mathbf{p}_j$, and $\hat{\mathbf{y}} = \{\hat{\mathbf{y}}_i | i \in [1, |\mathcal{R}|+1]\}$, where:

$$\hat{\mathbf{y}}_i = \begin{cases} 1, & \text{if } i = c \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

3.4 Learning to Score Relations

In this section, we introduce how we design our Scorer network, which produces a score of each predicted relation triplet by utilizing the relation feature and




Fig. 3. The proposed Scorer network that learns to score each predicted relation triplet.

the predicate distribution produced by a base scene graph model. By exploiting a self-attention module to compare each relation triplet with all the others, it learns to rank the nontrivial relations higher than the trivial ones and thus reduces *false positives*.

As illustrated in Figure 3, the relation features and predicate probability distributions from the base scene graph model are fed into two different multi-layer perceptron (MLP) networks, where the two MLPs serve as non-linear transformation functions. The output of two MLPs are added along the feature dimension, and then the input image feature extracted by the ImageNet backbone is concatenated to each relation feature to let the relation feature maintain global information of the image.

As the significance of a relation is also dependent on the existence of other relations, it is necessary to compare all detected relations before assigning their significance scores. In order to perform such reasoning, the combined relation features are fed into a dot-product self-attention (SA) module [44] to pass information among different relations:

$$\text{SA}(H) = \text{Softmax}(f_{\text{query}}(H)f_{\text{key}}(H)^{\top})f_{\text{value}}(H), \quad (5)$$

where $H \in \mathbb{R}^{N(N-1) \times (d_h + d_i)}$ is a matrix containing relation features concatenated with image features, while $f_{\text{query}}(\cdot)$, $f_{\text{key}}(\cdot)$ and $f_{\text{value}}(\cdot)$ are non-linear transformation functions applied to H . Then a final MLP with Sigmoid as its last activation function is applied to obtain the scores $S \in \mathbb{R}^{N(N-1) \times 1}$ which contains the significance score s for each predicted relation triplet.

Although the goal of Scorer is to predict the significance scores of relations, we do not have ground-truth scores. Remember that we only need to rank the annotated relations higher than the un-annotated ones (whose predicate labels are given as **unknown**), so we treat all annotated relations equally and assign 1 as their significance scores, while all un-annotated relations are assigned with 0 scores. In this case, the Scorer module can be trained with a binary cross-entropy (BCE) loss for each predicted relation triplet:

$$\mathcal{L}_{\text{BCE}}(s, q) = -q \log(s) - (1 - q) \log(1 - s), \quad (6)$$

Table 1. Quantitative comparison of different methods. *Note that we directly used the KERN [10] checkpoint trained on SGCls and test on all three tasks without fine-tuning for SGGen, while the KERN paper [10] finetunes for SGGen. We do not finetune for SGGen since we treat scene graph parsing as an extension of object detection, thus the bounding box regression should be solved by the object detector.

Task	SGGen			SGCls			PredCls		
Macro-averaged Recall@K	K=20	K=50	K=100	K=20	K=50	K=100	K=20	K=50	K=100
MotifNet [3]	2.09	3.37	4.52	6.22	7.62	8.06	10.87	14.18	15.31
LinkNet [2]	2.58	4.03	5.43	6.64	8.11	8.69	12.02	15.64	16.98
MotifNet+Focal [28]	1.68	3.16	4.50	5.79	7.56	8.41	10.13	14.11	16.12
LinkNet+Focal [28]	1.51	2.99	4.62	6.1	8.41	9.52	11.02	15.85	18.35
KERN [10] w/o fine-tuning*	2.24	3.96	5.39	7.68	9.36	10.00	13.83	17.72	19.17
KERN+CCE+Scorer	2.95	4.85	6.06	9.39	11.28	11.94	16.61	20.57	22.14
MotifNet+CCE+Scorer	3.06	4.72	6.11	7.89	9.91	10.61	15.03	19.62	21.49
LinkNet+CCE+Scorer	3.14	4.96	6.23	9.32	11.53	12.48	17.53	22.23	24.22
KERN [10] paper (fine-tuned)	N/A	6.4	7.3	N/A	9.4	10.0	N/A	17.7	19.2

where s is an output score of Scorer and $q = 0$ indicates the ground-truth label of this predicate belongs to the **unknown** class, otherwise $q = 1$. We have also tried other loss functions for regression, such as mean-square error and hinge loss, but we found that BCE works best.

During inference, the output scores of the Scorer network are used to determine the ranks of predicted relation triplets (the higher the score, the higher the rank), which are then used for calculating $mR@K$.

The whole framework is trained by combining \mathcal{L}_{CCE} and \mathcal{L}_{BCE} and other loss functions $\mathcal{L}_{\text{other}}$ specific to the chosen base scene graph model:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{CCE}} + \lambda_2 \mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{other}}, \quad (7)$$

where λ_1 and λ_2 are two hyper-parameters to balance the two loss terms.

4 Experiments

4.1 Experiment Settings

Dataset. Visual Genome [1] (VG) is the largest scene graph dataset containing 108,077 images with an average of 38 objects and 22 relation triplets per image. As there exist different splits for the VG dataset, we adopt the most widely used one [5], which contains 75,651 images for training (including 5,000 for validation) and 32,422 images for testing. There are numerous objects and predicates in the original VG dataset, but many of them have very low frequencies and quality. We follow Xu *et al.* [5] and only use the most frequent 150 object classes and 50 most frequent predicate classes in VG.

Tasks. We adopt the most widely studied tasks [5, 3, 2, 22] for evaluation: (1) **Scene Graph Generation** (SGGen) aims to simultaneously localize and classify objects in an image, and predict the potential predicate between each pair of

detected objects. An object is considered to be correctly detected if its intersection over union (IoU) with a ground-truth object is over 0.5. (2) **Scene Graph Classification** (SGCls) provides the model with a set of bounding boxes of objects, and requires the model to predict the object labels as well as the pairwise relations between the objects. (3) **Predicate Classification** (PredCls) requires the model to predict the potential predicates between each pair of objects, where the object locations and labels are given as groundtruth.

Metric. As discussed in previous sections, the widely used *micro-averaged Recall@K* ($R@K$) can be highly biased towards frequent classes and cannot tell whether a model is performing well on both frequent and rare classes. Instead, we use the *macro-averaged Recall@K* ($mR@K$) (Equation 1), *i.e.*, “mean recall@K” [10], as our evaluation metric, which treats major and minor classes alike. We adopt the widely used protocol for determining whether a predicted relation triplet matches groundtruth as in [5, 3].

4.2 Baselines

We use three of the current state-of-the-art methods **MotifNet** [3], **LinkNet** [2] and **KERN** [10] as the base scene graph models in our framework and study how our proposed method can help them improve their performance on rare predicate classes. KERN [10] is a recent approach that also tries to solve the class imbalance problem in scene graph parsing by designing a more complex network structure, and also uses the same evaluation metric as we do. As **Focal Loss** [28] is a popular method that addresses the class imbalance problem in object detection, we also try applying it to LinkNet [2] and MotifNet [3] in the scene graph parsing setting to see if it helps alleviate class imbalance in scene graph parsing.

4.3 Implementation Details

We use pretrained VGG16 [45] as our ImageNet backbone and Faster R-CNN [4] pretrained on VG [1] as the object detector. For MotifNet [3] and KERN [10], we use the official code provided by the authors^{3,4}, and use the same set of hyperparameters for the models as they are provided in their code base. Although there is no official implementation for LinkNet [2], we found an unofficial one⁵ that produces results close to those in the original paper. We apply our proposed CCE loss and Scorer to KERN [10], LinkNet [2] and MotifNet [3] without modifying their original model architectures, and only take the relation features before and after their final MLPs together with the global-average pooled image feature extracted from VGG16 pretrained on ImageNet [43] as input to our Scorer module. The two MLPs in our Scorer are both implemented as one fully

³ <https://github.com/rowanz/neural-motifs>

⁴ <https://github.com/HCLab-SYSU/KERN>

⁵ <https://github.com/jiayan97/linknet-pytorch>




Fig. 4. Per-class Recall@100 for KERN [10], LinkNet [2] and MotifNet [3] with CCE loss and Scorer, evaluated on the PredCls task. The number above each bar indicates the absolute increase (+) or decrease (−) in each predicate compared with the corresponding base scene graph model trained with cross-entropy. Predicates are sorted according to their frequencies in Figure 1(a).

connected layer of 512 units and a ReLU activation. The non-linear transformation functions $f_{\text{query}}(\cdot)$, $f_{\text{key}}(\cdot)$ and $f_{\text{value}}(\cdot)$ in Equation 5 are implemented as a fully connected layer with 256 units, and the last MLP in Scorer compresses the 256-dim features to scalars followed by a Sigmoid function. *We only tune λ_1 and λ_2 for loss weights in Equation 7 and keep other hyper-parameters fixed, while we found that setting $\alpha = 0$ generally works well for all models.* For fair comparison, we train all models on the SGCls tasks, and then directly evaluate on all three tasks without fine-tuning for each task. The random seed is fixed as 42 for all experiments. Our code is available online⁶.

4.4 Results

We present the quantitative results for our methods together with those of baselines in Table 1. As can be seen from the table, although MotifNet [3] and LinkNet [2] are among the current state-of-the-art methods that have high *micro-averaged Recall@K* (R@K) of around 67%, they have relatively low *macro-averaged Recall@K* (mR@K), where their mR@100 are 15.31% and 16.98% for the easiest PredCls task. As the tasks become more challenging, their performance for SGCls and SGGen are much worse, which shows that they work badly on detecting rare predicate classes. Among all methods without fine-tuning for SGGen, our proposed framework with LinkNet [2] as the base scene graph model achieves the highest mR@K in all three tasks and all K s. In the PredCls task, the LinkNet+CCE+Scorer method outperforms KERN [10] by a significant margin of around 4% to 5%. For the other two more challenging tasks SGCl and SGGen, LinkNet+CCE+Scorer method still outperforms KERN w/o fine-tuning by 2% and 1% respectively, which show the effectiveness of our proposed framework. Our method with MotifNet [3] and KERN [10] as the base scene graph models also outperform all other methods without our proposed framework, which

⁶ https://github.com/stevehuanghe/scene_graph

Table 2. Ablation study for CCE loss and Scorer.

Task	SGGen			SGCls			PredCls		
	K=20	K=50	K=all	K=20	K=50	K=all	K=20	K=50	K=all
MotifNet	2.09	3.37	6.81	6.22	7.62	8.31	10.87	14.18	16.19
MotifNet+CCE	1.43	2.97	7.94	4.63	6.68	9.17	7.87	12.91	18.93
MotifNet+Scorer	2.37	3.56	7.23	6.36	7.65	8.34	11.79	14.78	16.71
MotifNet+CCE+Scorer	3.06	4.72	9.21	7.89	9.91	10.93	15.03	19.62	22.57
LinkNet	2.58	4.03	7.87	6.64	8.11	9.02	12.02	15.64	17.86
LinkNet+CCE	1.32	2.73	8.63	5.38	8.04	11.28	9.07	14.30	22.06
LinkNet+Scorer	2.71	4.17	7.65	6.79	8.18	8.94	12.41	15.62	17.77
LinkNet+CCE+Scorer	3.14	4.96	10.32	9.32	11.53	12.94	17.53	22.23	25.30
LinkNet+Ours w/o SA	2.68	4.21	8.99	7.86	10.07	11.57	14.96	19.79	23.15

demonstrates that our framework is not specifically designed for a single model but can help improve the performance of different existing scene graph models trained with cross-entropy loss. It can also be noticed that, our framework with LinkNet [2] performs better than the one with MotifNet [3], which is reasonable as LinkNet [2] itself outperforms MotifNet [3] when they are both trained with cross-entropy loss. Our framework with KERN [10] achieves a performance that lies between the other two base models trained with our framework. The reason may be that the architecture of KERN [10] is more complex and requires extra hyper-parameter search and more delicate training techniques.

As for Focal Loss [28], we can see that it helps both LinkNet [2] and MotifNet [3] improve the results when $K = 100$ for all three tasks, but it does not work well with a smaller K such as $K = 20$. This result is not surprising, since the Focal Loss does not have an explicit mechanism to rank the predicted relations, and we have discussed in previous sections that the scene graph parsing task requires the model to not only classify predicates correctly but also know how to discriminate more important relations from trivial ones to achieve higher (mean) recall@K. The KERN model fine-tuned for SGGen achieves the highest scores on the hardest task SGGen, while the KERN without fine-tuning performs much worse than the fine-tuned version, which shows that fine-tuning is really important for this task. However, for fair comparison, we only compare the no-fine-tuned version with others.

We also investigate the changes in each predicate’s recall@100 on the PredCls task, as illustrated in Figure 4. We can see that our framework can generally improve the performance of base scene graph model in many rare classes starting from **standing on**, which occupies only 1.1% of all training samples. More remarkably, our framework with LinkNet as the base model increases the recall for a very rare predicate **parked on** (0.13% coverage in the whole dataset) from 0% to 56.08%. Among the top-5 most frequent predicates (**on**, **has**, **in**, **of**, and **wearing**) covering 74.5% of all relations in the dataset, our methods sacrifice some performance in **on**, **has** and **wearing**, which is reasonable and acceptable as we try to focus more on fine-grained predicates instead of these coarse-grained ones. Nonetheless, we also have an increase in a highly frequent class **in**. Among




Fig. 5. Confusion matrix evaluated on VG’s training set: (a) LinkNet [2] with cross-entropy (CE) loss and (b) LinkNet with our proposed contrasting cross-entropy (CCE) loss. The darker the color on diagonal line the better the model. The x-axis shows predicted labels, while the y-axis indicates true labels. Predicates are sorted according to their frequencies in Figure 1(a).

all 50 predicates, our framework helps increase the per-class recall@100 of 29, 37, and 42 classes when trained with KERN [10], MotifNet [3] and LinkNet [2] respectively, which again demonstrates that our framework can help increase the per-class recall of most predicates, whether they are frequent or not. More importantly, our framework helps improve the per-class recall of 18 of the rarest 20 predicates with MotifNet [3] and LinkNet [2], while the performance of 14 of the rarest 20 classes is also improved for KERN [10] trained with our framework. Since KERN [10] is already designed to address the class imbalance problem, it makes sense that our framework does not improve KERN’s performance as much as it does with MotifNet [3] and LinkNet [2].

4.5 Ablation Study

CCE loss aims to improve the model’s ability to classify rare classes, but it is not designed for ranking the predicted relations. In order to study the model’s classification ability without being affected by the ranks of predictions, we study the mR@all for three tasks, as shown in Table 2. As we can see, our CCE loss increases the mR@all for three tasks with two different base scene graph models [2, 3]. The CCE loss works well when K is large, but not when K is small like 20 and 50, since a smaller K will cause the metric mR@K to be more limited by the ranking performance of the model. As CCE is only designed for classification, it is not surprising that it does not work well when the metric is more dependent on ranking performance.

We also plot the confusion matrices calculated on the training set to see how well trained models can fit the training data. As shown in Figure 5(a), LinkNet [2] trained with cross-entropy loss is highly biased towards the largest class “on”, and many points on the diagonal line has shallow color which shows

that the model has poor performance in classifying these classes correctly. On the other hand, Figure 5(b) shows that LinkNet trained with CCE loss achieves better performance on many classes, as the color on diagonal line is much darker than the left one, which shows that the CCE loss can significantly alleviate the class imbalance problem in scene graph parsing.

The **Scorer**, on the other hand, aims to improve the model’s performance in ranking its output relations to achieve higher $mR@K$ when K is small. As shown in Table 2, our Scorer can help improve $mR@20$ and $mR@50$ in all cases. For $K=all$, MotifNet+Scorer has small improvements over MotifNet itself, while LinkNet+Scorer has a small decrease of average 0.13% compared with LinkNet [2] without Scorer, which is reasonable since $mR@all$ is more dependent on the model’s classification ability than ranking. Although the base models trained with Scorer alone does not have significant improvement over the ones without, it is worth noting that, together with our CCE loss, the whole proposed framework can significantly improve the mean recall@ K of both LinkNet [28] and MotifNet [3] in all three tasks, as shown in Table 2.

Self-attention. In the last row of Table 2, we show the results of LinkNet [2] with our CCE+Scorer framework but without self-attention, where we simply replace the self-attention module by a fully-connected network applied on each relation triplet individually (pure point-wise ranking). From the results we can see that although the proposed framework without self-attention still outperforms the LinkNet [2] baseline, its performance is worse than the one with self-attention (point-wise and list-wise combined ranking). This observation proves that it is important to compare each predicted relation with others by passing information among them, and justifies our design of the Scorer module that unifies both list-wise and point-wise ranking methods in a novel way.

5 Conclusion

In this paper, we show that although current scene graph models achieve high *micro-averaged recall@K* ($R@K$), they suffer from not being able to detect rare classes, and that *macro-averaged recall@K* ($mR@K$) is a more suitable metric as it treats both frequent and rare classes the same. To tackle the class imbalance problem, we first formulate the task as a combination of classification and ranking, and then we design a framework which consists of a Contrasting Cross-Entropy (CCE) loss and a Scorer module respectively for these two sub-tasks. Our proposed framework is general and can be applied to any existing scene graph models trained with cross-entropy. Extensive experiments demonstrate the effectiveness of our framework with KERN [10], LinkNet [2] and MotifNet [3].

Acknowledgement

This work is supported in part by NSF under grants III-1763325, III-1909323, and SaTC-1930941.

References

1. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision* **123** (2017) 32–73
2. Woo, S., Kim, D., Cho, D., Kweon, I.S.: Linknet: Relational embedding for scene graph. In: *Advances in Neural Information Processing Systems*. (2018) 560–570
3. Zellers, R., Yatskar, M., Thomson, S., Choi, Y.: Neural motifs: Scene graph parsing with global context. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2018) 5831–5840
4. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: *Advances in neural information processing systems*. (2015) 91–99
5. Xu, D., Zhu, Y., Choy, C.B., Fei-Fei, L.: Scene graph generation by iterative message passing. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2017) 5410–5419
6. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. (2016) 852–869
7. Yang, X., Tang, K., Zhang, H., Cai, J.: Auto-encoding scene graphs for image captioning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 10685–10694
8. Johnson, J., Krishna, R., Stark, M., Li, L.J., Shamma, D., Bernstein, M., Fei-Fei, L.: Image retrieval using scene graphs. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2015) 3668–3678
9. Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2018) 1219–1228
10. Chen, T., Yu, W., Chen, R., Lin, L.: Knowledge-embedded routing network for scene graph generation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 6163–6171
11. Yang, J., Lu, J., Lee, S., Batra, D., Parikh, D.: Graph r-cnn for scene graph generation. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. (2018) 670–685
12. Dai, B., Zhang, Y., Lin, D.: Detecting visual relationships with deep relational networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2017) 3076–3086
13. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014)
14. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. *ICLR* (2017)
15. Herzig, R., Raboh, M., Chechik, G., Berant, J., Globerson, A.: Mapping images to scene graphs with permutation-invariant structured prediction. In: *Advances in Neural Information Processing Systems*. (2018) 7211–7221
16. Li, Y., Ouyang, W., Zhou, B., Shi, J., Zhang, C., Wang, X.: Factorizable net: an efficient subgraph-based framework for scene graph generation. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. (2018) 335–351
17. Yu, R., Li, A., Morariu, V.I., Davis, L.S.: Visual relationship detection with internal and external linguistic knowledge distillation. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2017) 1974–1982

18. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* **9** (1997) 1735–1780
19. Gu, J., Zhao, H., Lin, Z., Li, S., Cai, J., Ling, M.: Scene graph generation with external knowledge and image reconstruction. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 1969–1978
20. Qi, M., Li, W., Yang, Z., Wang, Y., Luo, J.: Attentive relational networks for mapping images to scene graphs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 3957–3966
21. Wang, W., Wang, R., Shan, S., Chen, X.: Exploring context and visual pattern of relationship for scene graph generation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 8188–8197
22. Zhan, Y., Yu, J., Yu, T., Tao, D.: On exploring undetermined relationships for visual relationship detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019)
23. Buda, M., Maki, A., Mazurowski, M.A.: A systematic study of the class imbalance problem in convolutional neural networks. *CoRR* **abs/1710.05381** (2017)
24. Japkowicz, N., Stephen, S.: The class imbalance problem: A systematic study. *Intelligent data analysis* **6** (2002) 429–449
25. Oksuz, K., Cam, B.C., Kalkan, S., Akbas, E.: Imbalance problems in object detection: A review. *arXiv preprint arXiv:1909.00169* (2019)
26. Rota Buló, S., Neuhold, G., Kotschieder, P.: Loss max-pooling for semantic image segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2017) 2126–2135
27. Shrivastava, A., Gupta, A., Girshick, R.: Training region-based object detectors with online hard example mining. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2016) 761–769
28. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE international conference on computer vision*. (2017) 2980–2988
29. Qian, Q., Chen, L., Li, H., Jin, R.: Dr loss: Improving object detection by distributional ranking. *arXiv preprint arXiv:1907.10156* (2019)
30. Chen, K., Li, J., Lin, W., See, J., Wang, J., Duan, L., Chen, Z., He, C., Zou, J.: Towards accurate one-stage object detection with ap-loss. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 5119–5127
31. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*. Volume 1., IEEE (2005) 539–546
32. Kim, M., Sahu, P., Gholami, B., Pavlovic, V.: Unsupervised visual domain adaptation: A deep max-margin gaussian process approach. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 4380–4390
33. Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2015) 2794–2802
34. Zhang, C., Jia, B., Gao, F., Zhu, Y., Lu, H., Zhu, S.C.: Learning perceptual inference by contrasting. In: *Advances in Neural Information Processing Systems*. (2019) 1073–1085
35. Elsayed, G., Krishnan, D., Mobahi, H., Regan, K., Bengio, S.: Large margin deep networks for classification. In: *Advances in neural information processing systems*. (2018) 842–852

36. Hayat, M., Khan, S., Zamir, S.W., Shen, J., Shao, L.: Gaussian affinity for max-margin class imbalanced learning. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 6469–6479
37. Zhang, J., Shih, K.J., Elgammal, A., Tao, A., Catanzaro, B.: Graphical contrastive losses for scene graph generation. arXiv preprint arXiv:1903.02728 (2019)
38. Liu, T.Y., et al.: Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval* **3** (2009) 225–331
39. Sculley, D.: Combined regression and ranking. In: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM (2010) 979–988
40. Ibrahim, O.A.S., Landa-Silva, D.: Es-rank: evolution strategy learning to rank approach. In: Proceedings of the Symposium on Applied Computing, ACM (2017) 944–950
41. Moon, T., Smola, A., Chang, Y., Zheng, Z.: Intervalrank: isotonic regression with listwise and pairwise constraints. In: Proceedings of the third ACM international conference on Web search and data mining, ACM (2010) 151–160
42. Wang, X., Hua, Y., Kodirov, E., Hu, G., Garnier, R., Robertson, N.M.: Ranked list loss for deep metric learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019)
43. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition, Ieee (2009) 248–255
44. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. (2017) 5998–6008
45. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)