

Graph Density-Aware Losses for Novel Compositions in Scene Graph Generation

Boris Knyazev*^{1,2}

bknyazev@uoguelph.ca

Harm de Vries³

mail@harmdevries.com

Cătălina Cangea⁴

Catalina.Cangea@cst.cam.ac.uk

Graham W. Taylor^{1,6,2}

gwtaylor@uoguelph.ca

Aaron Courville^{6,7,5}

aaron.courville@umontreal.ca

Eugene Belilovsky⁵

belilove@mila.quebec

¹ University of Guelph

² Vector Institute for Artificial Intelligence

³ Element AI

⁴ University of Cambridge

⁵ Mila, Université de Montréal

⁶ Canada CIFAR AI Chair

⁷ CIFAR LMB Fellow

Abstract

Scene graph generation (SGG) aims to predict graph-structured descriptions of input images, in the form of objects and relationships between them. This task is becoming increasingly useful for progress at the interface of vision and language. Here, it is important—yet challenging—to perform well on novel (zero-shot) or rare (few-shot) compositions of objects and relationships. In this paper, we identify two key issues that limit such generalization. Firstly, we show that the standard loss used in this task is unintentionally a function of scene graph density. This leads to the neglect of individual edges in large sparse graphs during training, even though these contain diverse few-shot examples that are important for generalization. Secondly, the frequency of relationships can create a strong bias in this task, such that a “blind” model predicting the most frequent relationship achieves good performance. Consequently, some state-of-the-art models exploit this bias to improve results. We show that such models can suffer the most in their ability to generalize to rare compositions, evaluating two different models on the Visual Genome dataset and its more recent, improved version, GQA. To address these issues, we introduce a density-normalized edge loss, which provides more than a two-fold improvement in certain generalization metrics. Compared to other works in this direction, our enhancements require only a few lines of code and no added computational cost. We also highlight the difficulty of accurately evaluating models using existing metrics, especially on zero/few shots, and introduce a novel weighted metric.¹

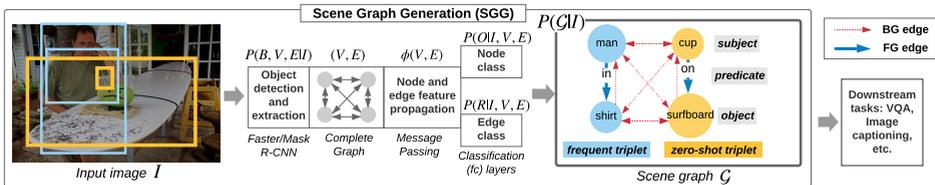


Figure 1: In this work, we improve scene graph generation $P(\mathcal{G}|I)$. In many downstream tasks, such as VQA, the result directly depends on the accuracy of predicted scene graphs.



Figure 2: **Motivation of our work.** We split the training set of Visual Genome [19] into two subsets: those with relatively small (≤ 10 nodes) and large (> 10 nodes) graphs. (a) While each subset contains a similar number of images (the three left bars), larger graphs contain more few shot labels (the three right bars). (b) Baseline methods ([13]) in this case) fail to learn from larger graphs due to their loss function. However, training on large graphs and corresponding few shot labels is important for stronger generalization. We address this limitation and significantly improve results on zero and few shots. (c, d) Small and large scene graphs typically describe simple and complex scenes respectively.

1 Introduction

In recent years, there has been growing interest to connect successes in visual perception with language and reasoning [22, 23]. This requires us to design systems that can not only recognize objects, but understand and reason about the relationships between them. This is essential for such tasks as visual question answering (VQA) [2, 5, 16] or caption generation [22, 56]. However, predicting a high-level semantic output (e.g. answer) from a low-level visual signal (e.g. image) is challenging due to a vast gap between the modalities. To bridge this gap, it would be useful to have some intermediate representation that can be relatively easily generated by the low-level module and, at the same time, can be effectively used by the high-level reasoning module. We want this representation to semantically describe the visual scene in terms of objects and relationships between them, which leads us to a structured image representation, the **scene graph** (SG) [18, 19]. A scene graph is a collection of visual relationship *triplets*: $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ (e.g. $\langle \text{cup}, \text{on}, \text{table} \rangle$). Each node in the graph corresponds to a subject or object (with a specific image location) and edges to predicates (Figure 1). Besides bridging the gap, SGs can be used to verify how well the model has understood the visual world, as opposed to just exploiting one of the biases in a dataset [9, 5, 17]. Alternative directions to SGs include, for example, attention [24].

Scene graph generation (SGG) is the task of predicting a SG given an input image. The inferred SG can be used directly for downstream tasks such as VQA [18, 58], image captioning [22, 56] or retrieval [9, 18, 30]. A model which performs well on SGG should demonstrate the ability to ground visual concepts to images and generalize to compositions of objects and predicates in new contexts. In real world images, some compositions (e.g. $\langle \text{cup}, \text{on}, \text{table} \rangle$) appear more frequently than others (e.g. $\langle \text{cup}, \text{on}, \text{surfboard} \rangle$ or $\langle \text{cup}, \text{under}, \text{table} \rangle$), which creates a strong frequency bias. This makes it particularly challenging for models to generalize to novel (zero shot) and rare (few shot) compositions, even though

each of the subjects, objects and predicates have been observed at training time. The problem is exacerbated by the test set and evaluation metrics, which do not penalize models that blindly rely on such bias. Indeed, Zellers et al. [57] has pointed out that SGG models largely exploit simple co-occurrence information. In fact, the performance of models predicting solely based on frequency (i.e. a cup is most likely to be *on* a *table*) is not far from the state-of-the-art using common metrics (see `FREQ` in Table 1).

In this work, we reveal that (a) the frequency bias exploited by certain models leads to poor generalization on few-shot and zero-shot compositions; (b) existing models disproportionately penalize large graphs, even if these often contain many of the infrequent visual relationships, which leads to performance degradation on few and zero-shot cases (Figure 2). We address these challenges and show that our suggested improvements can provide benefits for two strong baseline models [53, 57]. Overall, we make the following four **contributions**:

1. **Improved loss**: we introduce a density-normalized edge loss, which improves results on all metrics, especially for few and zero shots (Section 3.2);
2. **Novel weighted metric**: we illustrate several issues in the evaluation of few and zero shot compositions of objects and predicates, proposing a novel weighted metric which can better track the performance of this critical desiderata (Section 3.3);
3. **Frequency bias**: we demonstrate a negative effect of the frequency bias, proposed in Neural Motifs [57], on few and zero shot performance (Section 4);
4. **Scaling to GQA**: in addition to evaluating on Visual Genome (VG) [19], we confirm the usefulness of our loss and metrics on GQA [16] – an improved version of VG. Differing in its graph properties (Section 4), GQA has not been used to evaluate SGG models before.

2 Related Work

Zero shot learning. In vision tasks, such as image classification, zero shot learning has been extensively studied, and the main approaches are based on attributes [20] and semantic embeddings [11, 52]. The first approach is related to the zero-shot problem we address in this work: it assumes that all individual attributes of objects (color, shape, etc.) are observed during training, such that novel classes can be detected at test time based on *compositions* of their attributes. Zero shot learning in scene graphs is similar: all individual subjects, objects and predicates are observed during training, but most of their compositions are not. This task was first evaluated in [22] on the VRD dataset using a joint vision-language model. Several follow-up works attempted to improve upon it: by learning a translation operator in the embedding space [59], clustering in a weakly-supervised fashion [25], using conditional random fields [8] or optimizing a cycle-consistency loss to learn object-agnostic features [55]. Augmentation using generative networks to generate more examples of rare cases is another promising approach [51]; but, it was only evaluated in the predicate classification task. In our work, we also consider subject/object classification to enable the classification of the whole triplets, making the “image to scene graph” pipeline complete. Most recently, Tang et al. [30] proposed learning causal graphs and showed strong performance in zero shot cases.

While these works improve generalization, none of them has identified the challenges and importance of learning from large graphs for generalization. By concentrating the model’s capacity on smaller graphs and neglecting larger graphs, baseline models limit the variability of training data, useful for stronger generalization [12]. Our loss enables this learning, in-

creasing the effective data variability. Moreover, previous gains typically incur a large computational cost, while our loss has negligible cost and can be easily added to other models.

Few shot predicates. Several recent works have addressed the problem of imbalanced and few shot predicate classes [6, 7, 10, 29, 50, 41]. However, compared to our work, these works have not considered the imbalance between foreground and background edges, which is more severe than other predicate classes (Figure 3) and is important to be fixed as we show in this work. Moreover, we argue that the compositional generalization, not addressed in those works, can be more difficult than generalization to rare predicates. For example, the triplet <cup, on, surfboard> is challenging to be predicted correctly as a whole; even though ‘on’ can be the most frequent predicate, it has never been observed together with ‘cup’ and ‘surfboard’. Experimental results in previous work [22, 30, 31, 55, 59] highlight this difficulty. Throughout this work, by “few shot” we assume triplets, not predicates.

“Unbiasing” methods. Our idea is similar to the Focal loss [21], which addresses the imbalance between foreground and background objects in the object detection task. However, directly applying the focal loss to Visual Genome is challenging, due to the large amount of missing and mislabeled examples in the dataset. In this case, concentrating the model’s capacity on “hard” examples can be equivalent to putting more weight on noise, which can hurt performance. Tang et al. [50] compared the focal loss and other unbiasing methods, such as upsampling and upweighting, and did not report significantly better results.

3 Methods

In this section, we will review a standard loss used to train scene graph generation models (Section 3.1) and describe our improved loss, (6) (Section 3.2). We will then discuss issues with evaluating rarer combinations and propose a new weighted metric, (8) (Section 3.3).

3.1 Overview of Scene Graph Generation

In scene graph generation, given an image I , we aim to output a scene graph $\mathcal{G} = (O, R)$ consisting of a set of subjects and objects (O) as nodes and a set of relationships or predicates (R) between them as edges (Figure 1). So the task is to maximize the probability $P(\mathcal{G}|I)$, which can be expressed as $P(O, R|I) = P(O|I) P(R|I, O)$. Except for works that directly learn from pixels [23], the task is commonly [33, 34, 37] reformulated by first detecting bounding boxes B and extracting corresponding object and edge features, $V = f(I, B)$ and $E = g(I, B)$ respectively, using some functions f, g (e.g. a ConvNet followed by ROI Align [13]):

$$P(\mathcal{G}|I) = P(V, E|I) P(O, R|V, E, I). \quad (1)$$

The advantage of this approach is that solving $P(O, R|V, E, I)$ is easier than solving $P(O, R|I)$. At the same time, to compute $P(V, E|I)$ we can use pretrained object detectors [13, 26]. Therefore, we follow [33, 34, 37] and use this approach to scene graph generation.

In practice, we can assume that the pretrained object detector is fixed or that ground truth bounding boxes B are available, so we can assume $P(V, E|I)$ is constant. In addition, following [22, 33, 34], we can assume conditional independence of variables O and R : $P(O, R|V, E, I) = P(O|V, E, I)P(R|V, E, I)$. We thus obtain the scene graph generation loss:

$$-\log P(\mathcal{G}|I) = -\log P(O|V, E, I) - \log P(R|V, E, I). \quad (2)$$

Some models [57] reasonably do not assume the conditional independence of O and R , solving $P(O|V, E, I)P(R|O, V, E, I)$. However, such a model must be carefully regularized, since it can start to ignore (V, E, I) and mainly rely on the frequency distribution $P(R|O)$ as a stronger signal. For example, the model can learn that between ‘cup’ and ‘table’ the relationship is most likely to be ‘on’, regardless the visual signal. As we show, this can hurt generalization.

Eq. (2) is commonly handled as a multitask classification problem, where each task is optimized by the cross-entropy loss \mathcal{L} . In particular, given a batch of scene graphs with N nodes and M edges in total, the loss is the following:

$$\mathcal{L} = \mathcal{L}_{node} + \mathcal{L}_{edge} = \frac{1}{N} \sum_i \mathcal{L}_{obj,i} + \frac{1}{M} \sum_{ij} \mathcal{L}_{rel,ij}. \quad (3)$$

Node and edge features (V, E) output by the detector form a complete graph without self-loops (Figure 1). So, conventionally [53, 54, 57], the loss is applied to all edges: $M \approx N^2$. These edges can be divided into foreground (FG), corresponding to annotated edges, and background (BG), corresponding to not annotated edges: $M = M_{FG} + M_{BG}$. The BG edge type is similar to a “negative” class in the object detection task and has a similar purpose. Without training on BG edges, at test time the model would label all pairs of nodes as “positive”, i.e. having some relationship, when often it is not the case (at least, given the vocabulary in the datasets). Therefore, not using the BG type can hurt the quality of predicted scene graphs and can lower recall.

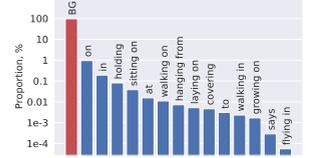


Figure 3: Predicate distribution in Visual Genome (split [53]). BG edges (note the log scale) dominate, with >96% of all edges, creating an extreme imbalance. For clarity, only some most and least frequent predicate classes are shown.

3.2 Hyperparameter-free Normalization of the Edge Loss

Baseline loss as a function of graph density. In scene graph datasets such as Visual Genome, the number of BG edges is greater than FG ones (Figure 3), yet the baseline loss (3) does not explicitly differentiate between BG and other edges. If we assume a fixed probability for two objects to have a relationship, then as the number of nodes grows we can expect fewer of them to have a relationship. Thus the graph density can vary based on the number of nodes (Figure 4), a fact not taken into account in Eq. (3). To avoid this, we start by decoupling the edge term of (3) into the foreground (FG) and background (BG) terms:

$$\mathcal{L}_{edge} = \frac{1}{M} \sum_{ij} \mathcal{L}_{rel,ij} = \frac{1}{M_{FG} + M_{BG}} \left[\underbrace{\sum_{ij \in \mathcal{E}}^{M_{FG}} \mathcal{L}_{rel,ij}}_{\text{FG edges}} + \underbrace{\sum_{ij \notin \mathcal{E}}^{M_{BG}} \mathcal{L}_{rel,ij}}_{\text{BG edges}} \right], \quad (4)$$

where \mathcal{E} is a set of FG edges, M_{FG} is the number of FG edges ($|\mathcal{E}|$) and M_{BG} is the number of BG edges. Next, we denote FG and BG edge losses averaged per batch as $\mathcal{L}_{FG} = 1/M_{FG} \sum_{ij \in \mathcal{E}}^{M_{FG}} \mathcal{L}_{rel,ij}$ and $\mathcal{L}_{BG} = 1/M_{BG} \sum_{ij \notin \mathcal{E}}^{M_{BG}} \mathcal{L}_{rel,ij}$, respectively. Then, using the definition of graph density as a proportion of FG edges to all edges, $d = M_{FG}/(M_{BG} + M_{FG})$, we can express the total baseline loss equivalent to Eq. (3) as a function of graph density:

$$\mathcal{L} = \mathcal{L}_{node} + d\mathcal{L}_{FG} + (1-d)\mathcal{L}_{BG}. \quad (5)$$

Density-normalized edge loss. Eq. (5) and Figure 4 allow us to notice two issues:

1. **Discrepancy of the loss between graphs of different sizes.** Since d exponentially decreases with graph size (Fig. 4, left), FG edges of larger graphs are weighted less than edges

of smaller graphs in the loss (Fig.4, middle), making the model neglect larger graphs.

2. Discrepancy between object and edge losses. Due to d tending to be small on average, L_{edge} is much smaller than L_{node} , so the model might focus mainly on L_{node} (Fig.4, right).

We propose to address both issues by normalizing FG and BG terms by graph density d :

$$\mathcal{L} = \mathcal{L}_{node} + \gamma[\mathcal{L}_{FG} + M_{BG}/M_{FG}\mathcal{L}_{BG}]. \quad (6)$$

where $\gamma = 1$ in our default hyperparameter-free variant and $\gamma \neq 1$ only to empirically analyze the loss (Table 4). Even though the BG term still depends on graph density, we found it to be less sensitive to variations in d , since the BG loss quickly converges to some stable value, performing a role of regularization (Figure 4, right). We examine this in detail in Section 4.

3.3 Weighted Triplet Recall

The common evaluation metric for scene graph prediction is image-level Recall@K or R@K [63, 64, 67]. To compute it, we first need to extract the top- K triplets, Top_K , from the *entire image* based on ranked predictions of a model. Given a set of ground truth triplets, GT, the image-level R@K is computed as (see Figure 5 for a visualization):

$$\text{R@K} = |\text{Top}_K \cap \text{GT}|/|\text{GT}|. \quad (7)$$

There are four issues with this metric (we discuss them in detail in *Supp. Material*):

- (a) The frequency bias of triplets means more frequent triplets will dominate the metric.
- (b) The denominator in (7) creates discrepancies between images with different $|\text{GT}|$ (the number of ground truth triplets in an image), especially pronounced in few/zero shots.
- (c) Evaluation of zero ($n = 0$) and different few shot cases $n = 1, 5, 10, \dots$ [63] leads to many R@K results [61]. This complicates the analysis. Instead, we want a single metric for all n .
- (d) Two ways of computing the image-level recall [23, 67], graph *constrained* and *unconstrained*, lead to very different results and complicate the comparison (Figure 5).

To address issue (a), the predicate-normalized metric, mean recall (mR@K) [6, 69] and weighted mR@K were introduced [61]. These metrics, however, only address the imbalance of predicate classes, not whole triplets. Early work [9, 72] used triplet-level Recall@K (or $R_{tr}@K$) for some tasks, which is based on ranking predicted triplets for each ground truth

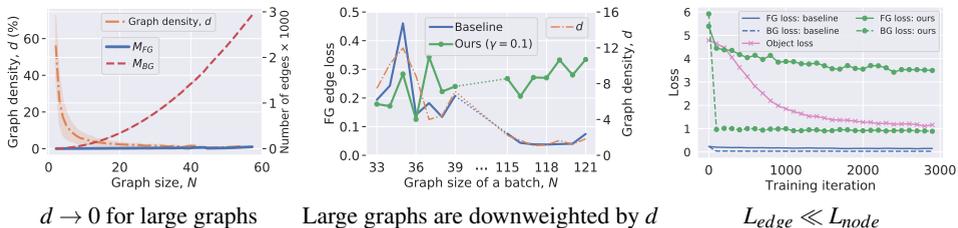


Figure 4: **(left)** The number of FG edges grows much more slowly with graph size than the number of BG edges (on VG: $M_{FG} \approx 0.5N$). This leads to: **(middle)** Downweighting of the FG loss on larger graphs and effectively limiting the amount and variability of training data, since large graphs contain a lot of labeled data (Figure 2); here, losses of converged models for batches sorted by graph size are shown. **(right)** Downweighting of the edge loss L_{edge} overall compared to L_{node} , even though both tasks are equally important to correctly predicting a scene graph. Our normalization fixes both issues.

subject-object pair independently; the pairs without relationships are not evaluated. Hence, $R_{tr}@K$ is similar to top- K accuracy. This metric avoids issues (b) and (d), but the issues of the frequency bias (a) and unseen/rare cases (c) still remain. To alleviate these, we adapt this metric to better track unseen and rare cases. We call our novel metric Weighted Triplet Recall $wR_{tr}@K$, which computes a recall at each triplet and reweights the average result based on the frequency of the GT triplet in the training set:

$$wR_{tr}@K = \sum_t^T w_t [\text{rank}_t \leq K], \quad (8)$$

where T is the number of all test triplets, $[\cdot]$ is the Iverson bracket, $w_t = \frac{1}{(n_t+1)\sum_l 1/(n_l+1)} \in [0, 1]$ and n_t is the number of occurrences of triplet t in the training set; $n_t + 1$ is used to handle zero-shot triplets; $\sum_l w_l = 1$. Since $wR_{tr}@K$ is still a triplet-level metric, we avoid issues (b) and (d). Our metric is also robust to the frequency-bias (a), since frequent triplets (with high n_t) are downweighted proportionally, which we confirm by evaluating the FREQ model from [57]. Finally, a single $wR_{tr}@K$ value shows zero and few shot performance linearly aggregated for all $n \geq 0$, solving issue (c).

4 Experiments

Datasets. We evaluate our loss and metric on Visual Genome [19]. Since it is a noisy dataset, several “clean” variants were introduced. We mainly experiment with the most common variant (VG) [33], which consists of the 150 most frequent object and 50 predicates classes. An alternative variant (VTE) [69] has been often used for zero-shot evaluation. Surprisingly, we found that the VG split [33] is better suited for this task, given a larger variability of zero-shot triplets in the test set (see *Supp. Material*). Recently, GQA [16] was introduced, where scene graphs were cleaned to automatically construct question answer pairs. GQA has more object and predicate classes, so that zero and few shot triplets are more likely to occur at test time. To the best of our knowledge, scene graph generation (SGG) results have not been reported on GQA before, even though some VQA models have relied on SGG [15].

Training and evaluation details. We experiment with two models: Message Passing (MP) [34] and Neural Motifs (NM) [34]. We use publicly available implementations of MP and NM², with all architecture details and hyperparameters kept the same (except for the small changes outlined in *Supp. Material*). We perform more experiments with Message Passing, since our experiments revealed that it better generalizes to zero and few shot cases, while performing only slightly worse on other metrics. In addition, it is a relatively simple model, which makes

²<https://github.com/rowanz/neural-motifs>

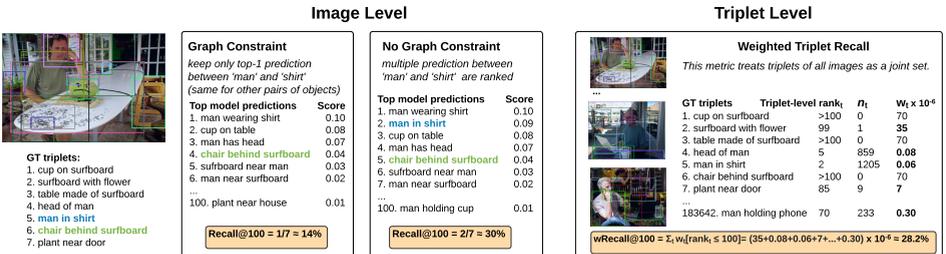


Figure 5: Existing image-level recall metrics *versus* our proposed weighted triplet recall. We first make unweighted predictions $\text{rank}_t \leq K$ for all GT triplets in all test images, then reweight them according to the frequency distribution (8). Computing our metric per image would be noisy.

Data-set	Model	Loss	Scene Graph Classification					Predicate Classification				
			R@100	R _{ZS} @100	R _{tr} @20	wR _{tr} @20	mR@100	R@50	R _{ZS} @50	R _{tr} @5	wR _{tr} @5	mR@50
Visual Genome	FREQ [63]	–	45.4	0.5	51.7	18.3	19.1	69.8	0.3	89.8	31.0	22.1
	MP [14, 15]	BASILINE (3)	47.2	8.2	51.9	26.2	17.3	74.8	23.3	86.6	51.3	20.6
		OURS (6)	48.6	9.1	52.6	28.2	26.5	78.2	28.4	89.4	58.4	32.1
	NM [14]	BASILINE (3)	48.1	5.7	51.9	26.5	20.4	80.5	11.1	91.0	51.8	26.9
		OURS (6)	48.4	7.1	52.0	27.7	25.5	82.0	16.7	92.0	56.4	34.8
		OURS (6), NO FREQ	48.4	8.9	51.8	28.0	26.1	82.5	26.6	92.4	60.3	35.8
KERN* [8]	BASILINE (3)	49.0	3.7	52.6	27.7	26.2	81.9	5.8	91.9	49.1	36.3	
	ReIDN* [41]	BASILINE (3)	50.8 [†]	–	–	–	–	93.7 [†]	–	–	–	
GQA	MP [14, 15]	BASILINE (3)	27.1	2.8	31.9	8.9	1.6	59.7	34.9	96.4	88.4	1.8
		OURS (6)	27.6	3.0	32.2	8.9	2.8	61.0	37.2	96.9	89.5	2.9
GQA-nLR	MP [14, 15]	BASILINE (3)	24.9	3.0	30.2	12.4	2.8	58.1	21.7	71.6	47.0	4.6
		OURS (6)	25.0	3.2	29.4	12.6	7.0	62.4	26.2	77.9	55.0	12.1

Table 1: Results on Visual Genome (split [63]) and GQA [14]. We obtain particularly strong results in columns R_{ZS}, wR_{tr} and mR in each of the two tasks. denotes cases with $\geq 15\%$ relative difference between the baseline and our result; denotes a difference of $\geq 50\%$. Best results for each dataset (VG, GQA and GQA-nLR) are bolded. GQA-nLR: our version of GQA with left/right spatial relationships excluded, where scene graphs become much sparser (see *Supp. Material* for dataset details). *Results are provided for the reference and evaluating our loss with these methods is left for future work. [†]The correctness of this evaluation is discussed in [48].

the analysis of its performance easier. We evaluate a trained model on two tasks, according to [63]: scene graph classification (**SGCls**), in which the model must label objects and relationships between them given ground truth bounding boxes, i.e. $P(O, R|I, B)$; and predicate classification (**PredCls**), in which the model only needs to label a predicate, i.e. $P(R|I, B, O)$. Results on **SGGen**, $P(\mathcal{G}|I)$, which includes detecting bounding boxes first, directly depend on **SGCls**, and are reported and discussed in *Supp. Material*.

4.1 Results

Table 1 shows our main results, where for each task we report five metrics: image-level recall on all triplets (R@K) and zero-shot triplets (R_{ZS}@K), triplet-level recall R_{tr}@K and our weighted triplet recall (wR_{tr}@K), and mR@K. We compute recalls without the graph constraint since, as we discuss in *Supp. Material*, this is a more accurate metric. We denote graph-constrained results as **SGCls-GC**, **PredCls-GC** and report them only in Tables 2, 3.

VG results. We can observe that both Message Passing (MP) and Neural Motifs (NM) greatly benefit from our density-normalized loss on all reported metrics. Larger gaps are achieved on metrics evaluating zero and few shots. For example, in PredCls on Visual Genome, MP with our loss is 22% better (in relative terms) on zero shots, while NM with our loss is 50% better. The gains arising from other zero shot and weighted metrics are also significant. **GQA results.** On GQA, our loss also consistently improves results, especially in PredCls. However, the gap is lower compared to VG. There are two reasons for this: 1) scene graphs in GQA are much denser (see *Supp. Material*), i.e. the imbalance between FG and BG edges is less pronounced, which means that in the baseline loss the edge term is not diminished to the extent it is in VG; and 2) the training set of GQA is more diverse than VG (with 15 times more labeled triplets), which makes the baseline model generalize well on zero and few shots. We confirm these arguments by training and evaluating on our version of GQA: **GQA-nLR** with left and right predicate classes excluded making scene graph properties, in particular sparsity, more similar to those of VG.

Effect of the Frequency Bias (FREQ) on Zero and Few Shot Performance. The FREQ model [67] simply predicts the most frequent predicate between a subject and an object, $P(R|O)$. Its effect on few shot generalization has not been empirically studied before. We study this by adding/abating FREQ from baseline MP and NM on Visual Genome (Figure 6). Our results indicate that FREQ only marginally improves results on unweighted metrics. At the same time, perhaps unsurprisingly, it leads to severe drops in zero shot and weighted metrics, especially in NM. For example, by ablating FREQ from NM, we improve $\text{PredCls-R}_{ZS}@50$ from 11% to 25%. This also highlights that the existing recall metrics are a poor choice of model.

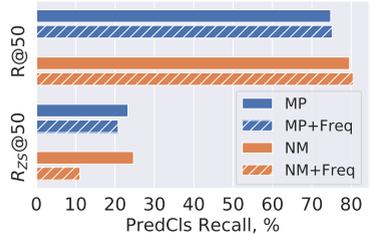


Figure 6: Ablating FREQ. FREQ only marginally improves results on R@50. At the same time, it leads to large drops in zero-shot recall $R_{ZS}@50$ (and our weighted triplet recall, see Table 1).

to understand the effectiveness of a

Why does loss normalization help more on few and zero shots?

The baseline loss effectively ignores edge labels of large graphs, because it is scaled by a small d in those cases (Figure 4). To validate that, we split the training set of Visual Genome in two subsets, with a comparable number of images in each: with relatively small and large graphs evaluating on the original test set in both cases. We observe that the baseline model does not learn well from large graphs, while our loss enables this learning (Figure 7). Moreover, when trained on small graphs only, the baseline is even better in PredCls than when trained on all graphs. This is because in the latter case, large graphs, when present in a batch, make the whole batch more sparse, down-weighting the edge loss of small graphs as well. At the same time, larger graphs predictably contain more labels, including many few shot labels (Figure 2). Together, these two factors make the baseline ignore many few shot triplets pertaining to larger graphs at training time, so the model cannot generalize to them at test time. Since the baseline essentially observes less variability during training, it leads to poor generalization on zero shots as well. This argument aligns well to the works from other domains [42], showing that generalization strongly depends on the diversity of samples during training. Our loss fixes the issue of learning from larger graphs, which, given the reasons above, directly affects the ability to generalize.

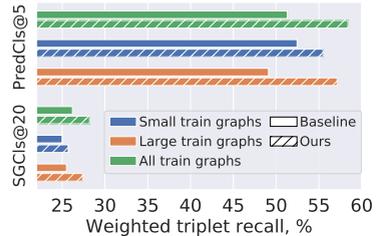


Figure 7: Learning from small ($N \leq 10$) vs. large ($N > 10$) graphs. Our loss makes models learn from larger graphs more effectively, which is important for generalization, because such graphs contain a lot of labels (see Figure 2).

Alternative approaches. We compare our loss to ones with tuned hyperparameters α, β, λ :

$$\mathcal{L} = \mathcal{L}_{node} + \alpha \mathcal{L}_{FG} + \beta \mathcal{L}_{BG}, \quad (9)$$

$$\mathcal{L} = \mathcal{L}_{node} + \lambda \mathcal{L}_{edge}. \quad (10)$$

Our main finding (Table 4) is that, while these losses can give similar or better results in some cases, the parameters α, β and λ do not generally transfer across datasets and must be tuned every time, which can be problematic at larger scale [40]. In contrast, our loss does not require tuning and achieves comparable performance.

Model	SGCls-GC	PredCls-GC
VTE [54, 69]	–	16.4
STA [54]	–	18.9
ST-GAN [54]	–	19.0
MP, baseline (3)	2.3	20.4
MP, ours (6)	3.1	21.4

Table 2: Zero-shot results ($R_{ZS}@100$) on the VTE split [54].

Model	SGCls-GC	PredCls-GC
NM+SUM+TDE [50]	4.5	18.2
NM, baseline (3)	1.7	9.5
MP, baseline (3)	3.2	20.1
NM, no Freq, ours (6)	3.9	20.4
MP, ours (6)	4.2	21.5

Table 3: Zero-shot results ($R_{ZS}@100$) on the VG split [53]. Tang et al. [50] uses ResNeXt-101 as a backbone, which helps to improve results.

Tuning dataset	Hyperparams	Loss	Testing on VG		Testing on GQA	
			SGCls@100	PredCls@50	SGCls@100	PredCls@50
No tune (baseline)	–	(3)	47.2/8.2	74.8/23.3	27.1/2.8	59.7/34.9
VG	$\lambda = 20$	(10)	48.9/9.2	78.3/27.9	26.6/2.6	60.4/36.9
VG	$\alpha = 0.5, \beta = 20$	(9)	49.1/9.4	78.2/27.8	27.1/2.9	60.5/36.3
GQA	$\lambda = 5$	(10)	48.8/9.2	78.0/26.8	27.8/2.9	60.5/36.1
GQA	$\alpha = 1, \beta = 5$	(9)	48.6/8.7	77.4/27.8	27.5/2.9	60.7/36.6
No tune (ours, independ. norm)	$\alpha = \beta = 1$	(9)	47.5/8.4	74.3/25.3	27.4/2.9	59.5/35.4
No tune (ours, no upweight)	$\gamma = 0.05/0.2$ for VG/GQA	(6)	48.7/ 9.6	78.3/28.2	27.4/2.9	61.1/36.8
No tune (ours)	$\gamma = 1$	(6)	48.6/9.1	78.2/ 28.4	27.6/3.0	61.0/ 37.2

Table 4: Comparing our loss to other approaches using MP [53, 57] and $R/R_{ZS}@K$ metrics.

Finally, to study the effect of density normalization separately from upweighting the edge loss (which is a side effect of our normalization), we also consider downweighting our edge term (6) by some $\gamma < 1$ to cancel out this upweighting effect. This ensures a similar range for the losses in our comparison. We found (Table 4) that the results are still significantly better than the baseline and, in some cases, even better than our hyperparameter-free loss. This further confirms that normalization of the graph density is important on its own. When carefully fine-tuned, the effects of normalization and upweighting are complimentary (*e.g.* when α, β or γ are fine-tuned, the results tend to be better).

Comparison to other zero shot works. We also compare to previous works studying zero shot generalization (Tables 2 and 3). For comprehensive evaluation, we test on both VTE and VG splits. We achieve superior results on VTE, even by just using the baseline MP, because, as shown in our main results, it generalizes well. On the VG split, we obtain results that compete with a more recent Total Direct Effect (TDE) method [50], even though the latter uses a more advanced detector and feature extractor. In all cases, our loss improves baseline results and, except for $R_{ZS}@100$ in SGClS, leads to state-of-the-art generalization. Our loss and TDE can be applied to a wide range of models, beyond MP and NM, to potentially have a complementary effect on generalization, which is interesting to study in future work.

5 Conclusions

Scene graphs are a useful semantic representation of images, accelerating research in many applications, such as visual question answering. It is vital for the SGG model to perform well on unseen or rare compositions of objects and predicates, which are inevitable due to an extremely long tail of the triplets distribution. We show that strong baseline models do not effectively learn from all labels, leading to poor generalization on few/zero shots. Moreover, current evaluation metrics do not reflect this problem, exacerbating it instead. We also show that learning well from larger graphs is essential to enable stronger generalization. To this end, we modify the loss commonly used in SGG and achieve significant improvements and, in certain cases, state-of-the-art results, on both the existing and our novel weighted metric.

Acknowledgments

BK is funded by the Mila internship, the Vector Institute and the University of Guelph. CC is funded by DREAM CDT. EB is funded by IVADO. This research was developed with funding from DARPA. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. The authors also acknowledge support from the Canadian Institute for Advanced Research and the Canada Foundation for Innovation. We are also thankful to Brendan Duke for the help with setting up the compute environment. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute: <http://www.vectorinstitute.ai/#partners>.

References

- [1] Ankesh Anand, Eugene Belilovsky, Kyle Kastner, Hugo Larochelle, and Aaron Courville. Blind-fold baselines for embodied QA. *arXiv preprint arXiv:1811.05013*, 2018.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [3] Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron Courville. Systematic generalization: what is required and can it be learned? *arXiv preprint arXiv:1811.12889*, 2018.
- [4] Eugene Belilovsky, Matthew Blaschko, Jamie Kiros, Raquel Urtasun, and Richard Zemel. Joint embeddings of scene graphs and images. 2017.
- [5] Cătălina Cangea, Eugene Belilovsky, Pietro Liò, and Aaron Courville. Videonavqa: Bridging the gap between visual and embodied question answering. *arXiv preprint arXiv:1908.04950*, 2019.
- [6] 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.
- [7] 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.
- [8] Weilin Cong, William Wang, and Wang-Chien Lee. Scene graph generation via conditional random fields. *arXiv preprint arXiv:1811.08075*, 2018.
- [9] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *Proceedings of the IEEE conference on computer vision and Pattern recognition*, pages 3076–3086, 2017.
- [10] Apoorva Dornadula, Austin Narcomey, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationships as functions: Enabling few-shot scene graph prediction. In *ArXiv*, 2019.
- [11] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pages 2121–2129, 2013.

- [12] 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 International Conference on Computer Vision*, pages 10323–10332, 2019.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [14] Felix Hill, Andrew Lampinen, Rosalia Schneider, Stephen Clark, Matthew Botvinick, James L. McClelland, and Adam Santoro. Environmental drivers of systematicity and generalization in a situated agent, 2019.
- [15] Drew Hudson and Christopher D Manning. Learning by abstraction: The neural state machine. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 5903–5916. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/8825-learning-by-abstraction-the-neural-state-machine.pdf>.
- [16] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6700–6709, 2019.
- [17] Allan Jabri, Armand Joulin, and Laurens Van Der Maaten. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer, 2016.
- [18] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015.
- [19] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, and et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, Feb 2017. ISSN 1573-1405. doi: 10.1007/s11263-016-0981-7. URL <http://dx.doi.org/10.1007/s11263-016-0981-7>.
- [20] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465, 2013.
- [21] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [22] 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.
- [23] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. In *Advances in neural information processing systems*, pages 2171–2180, 2017.
- [24] Will Norcliffe-Brown, Stathis Vafeias, and Sarah Parisot. Learning conditioned graph structures for interpretable visual question answering. In *Advances in Neural Information Processing Systems*, pages 8334–8343, 2018.
- [25] Julia Peyre, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Weakly-supervised learning of visual relations. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5179–5188, 2017.

- [26] 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.
- [27] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019.
- [28] Kaihua Tang. Scene graph benchmark in pytorch, 2020. URL <https://github.com/KaihuaTang/Scene-Graph-Benchmark.pytorch/blob/master/METRICS.md#topk-accuracy-ak>.
- [29] 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.
- [30] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [31] Xiaogang Wang, Qianru Sun, Marcelo ANG, and Tat-Seng CHUA. Generating expensive relationship features from cheap objects. 2019.
- [32] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 69–77, 2016.
- [33] 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.
- [34] 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.
- [35] Xu Yang, Hanwang Zhang, and Jianfei Cai. Shuffle-then-assemble: Learning object-agnostic visual relationship features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 36–52, 2018.
- [36] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10685–10694, 2019.
- [37] 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.
- [38] Cheng Zhang, Wei-Lun Chao, and Dong Xuan. An empirical study on leveraging scene graphs for visual question answering. *arXiv preprint arXiv:1907.12133*, 2019.
- [39] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5532–5540, 2017.
- [40] Ji Zhang, Yannis Kalantidis, Marcus Rohrbach, Manohar Paluri, Ahmed Elgammal, and Mohamed Elhoseiny. Large-scale visual relationship understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9185–9194, 2019.

- [41] 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.
- [42] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and VQA. *arXiv preprint arXiv:1909.11059*, 2019.