




Graph R-CNN for Scene Graph Generation

Jianwei Yang¹(✉) , Jiasen Lu¹, Stefan Lee¹, Dhruv Batra^{1,2},
and Devi Parikh^{1,2}

¹ Georgia Institute of Technology, Atlanta, USA

{jw2yang, jiasenlu, steflee, dbatra, parikh}@gatech.edu

² Facebook AI Research, Menlo Park, USA

Abstract. We propose a novel scene graph generation model called Graph R-CNN, that is both effective and efficient at detecting objects and their relations in images. Our model contains a Relation Proposal Network (RePN) that efficiently deals with the quadratic number of potential relations between objects in an image. We also propose an attentional Graph Convolutional Network (aGCN) that effectively captures contextual information between objects and relations. Finally, we introduce a new evaluation metric that is more holistic and realistic than existing metrics. We report state-of-the-art performance on scene graph generation as evaluated using both existing and our proposed metrics.

Keywords: Graph R-CNN · Scene graph generation
Relation proposal network · Attentional graph convolutional network

1 Introduction

Visual scene understanding has traditionally focused on identifying *objects in images* – learning to predict their presence (*i.e.* image classification [9, 15, 34]) and spatial extent (*i.e.* object detection [7, 22, 31] or segmentation [21]). These object-centric techniques have matured significantly in recent years, however, representing scenes as collections of objects fails to capture relationships which may be essential for scene understanding.

A recent work [12] has instead proposed representing visual scenes as graphs containing objects, their attributes, and the relationships between them. These *scene graphs* form an interpretable structured representation of the image that can support higher-level visual intelligence tasks such as captioning [24, 39], visual question answering [1, 11, 35, 37–39], and image-grounded dialog [3]. While scene graph representations hold tremendous promise, extracting scene graphs from images – efficiently and accurately – is challenging. The natural approach of considering every pair of nodes (objects) as a potential edge (relationship) – essentially reasoning over fully-connected graphs – is often effective in modeling

J. Yang and J. Lu – Equal contribution

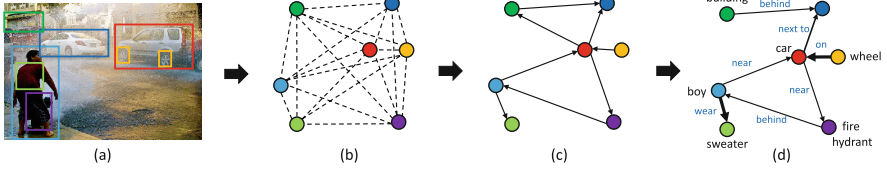


Fig. 1. Given an image (a), our proposed approach first extracts a set of objects visible in the scene and considers possible relationships between all nodes (b). Then it prunes unlikely relationships using a learned measure of ‘relatedness’, producing a sparser candidate graph structure (c). Finally, an attentional graph convolution network is applied to integrate global context and update object node and relationship edge labels.

contextual relationships but scales poorly (quadratically) with the number of objects, quickly becoming impractical. The naive fix of randomly sub-sampling edges to be considered is more efficient but not as effective since the distribution of interactions between objects is far from random – take Fig. 1(a) as an example, it is much more likely for a ‘car’ and ‘wheel’ to have a relationship than a ‘wheel’ and ‘building’. Furthermore, the types of relationships that typically occur between objects are also highly dependent on those objects.

Graph R-CNN. In this work, we propose a new framework, Graph R-CNN, for scene graph generation which effectively leverages object-relationship regularities through two mechanisms to intelligently sparsify and reason over candidate scene graphs. Our model can be factorized into three logical stages: (1) object node extraction, (2) relationship edge pruning, and (3) graph context integration, which are depicted in Fig. 1. In the object node extraction stage, we utilize a standard object detection pipeline [32]. This results in a set of localized object regions as shown in Fig. 1b. We introduce two important novelties in the rest of the pipeline to incorporate the real-world regularities in object relationships discussed above. First, we introduce a relation proposal network (RePN) that learns to efficiently compute *relatedness scores* between object pairs which are used to intelligently prune unlikely scene graph connections (as opposed to random pruning in prior work). A sparse post-pruning graph is shown in Fig. 1c. Second, given the resulting sparsely connected scene graph candidate, we apply an attentional graph convolution network (aGCN) to propagate higher-order context throughout the graph – updating each object and relationship representation based on its neighbors. In contrast to existing work, we predict per-node edge attentions, enabling our approach to learn to modulate information flow across unreliable or unlikely edges. We show refined graph labels and edge attentions (proportional to edge width) in Fig. 1d.

To validate our approach, we compare our performance with existing methods on the Visual Genome [14] dataset and find that our approach achieves an absolute gain of 5.0 on Recall@50 for scene graph generation [40]. We also perform extensive model ablations and quantify the impact of our modeling choices.

Evaluating Scene Graph Generation. Existing metrics for scene graph generation are based on recall of ⟨subject, predicate, object⟩ triplets (e.g. *SGGen* from [14]) or of objects and predicates given ground truth object localizations (e.g. *PredCls* and *PhrCls* from [14]). In order to expose a problem with these metrics, consider a method that mistakes the boy in Fig. 1a as a man but otherwise identifies that he is (1) standing behind a fire hydrant, (2) near a car, and (3) wearing a sweater. Under the triplet-based metrics, this minor error (boy vs man) would be heavily penalized despite most of the boy’s relationships being correctly identified. Metrics that provide ground-truth regions side-step this problem by focusing strictly on relationship prediction but cannot accurately reflect the test-time performance of the entire scene graph generation system.

To address this mismatch, we introduce a novel evaluation metric (*SGGen+*) that more holistically evaluates the performance of scene graph generation with respect to objects, attributes (if any), and relationships. Our proposed metric *SGGen+* computes the total recall for singleton entities (objects and predicates), pair entries ⟨object, attribute⟩ (if any), and triplet entities ⟨subject, predicate, object⟩. We report results on existing methods under this new metric and find our approach also outperforms the state-of-the-art significantly. More importantly, this new metric provides a more robust and holistic measure of similarity between generated and ground-truth scene graphs.

Summary of Contributions. Concretely, this work addresses the scene graph generation problem by introducing a novel model (Graph R-CNN), which can leverage object-relationship regularities, and proposes a more holistic evaluation metric (*SGGen+*) for scene graph generation. We benchmark our model against existing approaches on standard metrics and this new measure – outperforming existing approaches.

2 Related Work

Contextual Reasoning and Scene Graphs. The idea of using context to improve scene understanding has a long history in computer vision [16, 27, 28, 30]. More recently, inspired by representations studied by the graphics community, Johnson *et al.*[12] introduced the problem of extracting scene graphs from images, which generalizes the task of object detection [6, 7, 22, 31, 32] to also detecting relationships and attributes of objects.

Scene Graph Generation. A number of approaches have been proposed for the detection of both objects and their relationships [2, 17–19, 23, 26, 29, 40, 42–44, 46]. Though most of these works point out that reasoning over a quadratic number of relationships in the scene graph is intractable, each resorted to heuristic methods like random sampling to address this problem. Our work is the first to introduce a trainable relationship proposal network (RePN) that learns to prune unlikely relationship edges from the graph without sacrificing efficacy. RePN provides high-quality relationship candidates, which we find improves overall scene graph generation performance.

Most scene graph generation methods also include some mechanisms for context propagation and reasoning over a candidate scene graph in order to refine the final labeling. In [40], Xu *et al.* decomposed the problem into two sub-graphs – one for objects and one for relationships – and performed message passing. Similarly, in [17], the authors propose two message-passing strategies (parallel and sequential) for propagating information between objects and relationships. Dai *et al.* [2] address model the scene graph generation process as inference on a conditional random field (CRF). Newell *et al.* [26] proposed to directly generate scene graphs from image pixels without the use of object detector based on associative graph embeddings. In our work, we develop a novel attentional graph convolutional network (aGCN) to update node and relationship representations by propagating context between nodes in candidate scene graphs – operating both on visual and semantic features. While similar in function to the message-passing based approach above, aGCN is highly efficient and can learn to place attention on reliable edges and dampen the influence of unlikely ones.

A number of previous approaches have noted the strong regularities in scene graph generation which motivate our approach. In [23], Lu *et al.* integrates semantic priors from language to improve the detection of meaningful relationships between objects. Likewise, Li *et al.* [18] demonstrated that region captions can also provide useful context for scene graph generation. Most related to our motivation, Zeller *et al.* [42] formalize the notion of motifs (*i.e.*, regularly occurring graph structures) and examine their prevalence in the Visual Genome dataset [14]. The authors also propose a surprisingly strong baseline which directly uses frequency priors to predict relationships – explicitly integrating regularities in the graph structure.

Relationship Proposals. Our Relationship Proposal Network (RePN) is inspired and relates strongly to the region proposal network (RPN) of faster R-CNN [32] used in object detection. Our RePN is also similar in spirit to the recently-proposed relationship proposal network (Rel-PN) [45]. There are a number of subtle differences between these approaches. The Rel-PN model independently predicts proposals for subject, objects and predicates, and then re-scores all valid triples, while our RePN generates relations conditioned on objects, allowing it to learn object-pair relationship biases. Moreover, their approach is class agnostic and has not been used for scene graph generation.

Graph Convolutional Networks (GCNs). GCNs were first proposed in [13] in the context of semi-supervised learning. GCNs decompose complicated computation over graph data into a series of localized operations (typically only involving neighboring nodes) for each node at each time step. The structure and edge strengths are typically fixed prior to the computation. For completeness, we note that an upcoming publication [36] has concurrently and independently developed a similar GCN attention mechanism (as aGCN) and shown its effectiveness in other (non-computer vision) contexts.

3 Approach

In this work, we model scene graphs as graphs consisting of image regions, relationships, and their labellings. More formally, let I denote an image, V be a set of nodes corresponding to localized object regions in I , $E \in \binom{V}{2}$ denote the relationships (or edges) between objects, and O and R denote object and relationship labels respectively. Thus, the goal is to build a model for $P(S = (V, E, O, R)|I)$. In this work, we factorize the scene graph generation process into three parts:

$$P(S|I) = \underbrace{P(V|I)}_{\text{Object Region Proposal}} \underbrace{P(E|V, I)}_{\text{Relationship Proposal}} \underbrace{P(R, O|V, E, I)}_{\text{Graph Labeling}} \tag{1}$$

which separates graph construction (nodes and edges) from graph labeling. The intuition behind this factorization is straightforward. First, the object region proposal $P(V|I)$ is typically modeled using an off-the-shelf object detection system such as [32] to produce candidate regions. Notably, existing methods typically model the second relationship proposal term $P(E|V, I)$ as a uniform random sampling of potential edges between vertices V . In contrast, we propose a relationship proposal network (RePN) to directly model $P(E|V, I)$ – making our approach the first that allows for learning the entire generation process end-to-end. Finally, the graph labeling process $P(R, O|V, E, I)$ is typically treated as an iterative refinement process [2, 17, 40]. A brief pipeline is shown in Fig. 2.

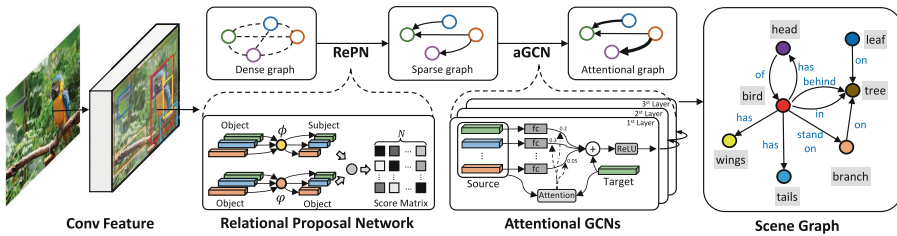


Fig. 2. The pipeline of our proposed Graph R-CNN framework. Given an image, our model first uses RPN to propose object regions, and then prunes the connections between object regions through our relation proposal network (RePN). Attentional GCN is then applied to integrate contextual information from neighboring nodes in the graph. Finally, the scene graph is obtained on the right side.

In the following, we discuss the components of our proposed Graph R-CNN model corresponding to each of the terms in Eq. 1. First, we discuss our use of Faster R-CNN [32] for node generation in Sect. 3.1. Then in Sect. 3.2 we introduce our novel relation proposal network architecture to intelligently generate edges. Finally, in Sect. 3.3 we present our graph convolutional network [13] with learned attention to adaptively integrate global context for graph labeling.

3.1 Object Proposals

In our approach, we use the Faster R-CNN [32] framework to extract a set of n object proposals from an input image. Each object proposal i is associated with a spatial region $r_i^o = [x_i, y_i, w_i, h_i]$, a pooled feature vector x_i^o , and an initial estimated label distribution p_i^o over classes $C = \{1, \dots, k\}$. We denote the collection of these vectors for all n proposals as the matrices $R^o \in \mathbb{R}^{n \times 4}$, $X^o \in \mathbb{R}^{n \times d}$, and $P^o \in \mathbb{R}^{n \times |C|}$ respectively.

3.2 Relation Proposal Network

Given the n proposed object nodes from the previous step, there are $O(n^2)$ possible connections between them; however, as previously discussed, most object pairs are unlikely to have relationships due to regularities in real-world object interactions. To model these regularities, we introduce a relation proposal network (RePN) which learns to efficiently estimate the *relatedness* of an object pair. By pruning edges corresponding to unlikely relations, the RePN can efficiently sparsify the candidate scene graph – retaining likely edges and suppressing noise introduced from unlikely ones.

In this paper, we exploit the estimated class distributions (P^o) to infer relatedness – essentially learning soft class-relationships priors. This choice aligns well with our intuition that certain classes are relatively unlikely to interact compared with some other classes. Concretely, given initial object classification distributions P^o , we score all $n * (n - 1)$ directional pairs $\{\mathbf{p}_i^o, \mathbf{p}_j^o | i \neq j\}$, computing the relatedness as $s_{ij} = f(\mathbf{p}_i^o, \mathbf{p}_j^o)$ where $f(\cdot, \cdot)$ is a learned relatedness function. One straightforward implementation of $f(\cdot, \cdot)$ could be passing the concatenation $[\mathbf{p}_i^o, \mathbf{p}_j^o]$ as input to a multi-layer perceptron which outputs the score. However, this approach would consume a great deal of memory and computation given the quadratic number of object pairs. To avoid this, we instead consider an asymmetric kernel function:

$$f(\mathbf{p}_i^o, \mathbf{p}_j^o) = \langle \Phi(\mathbf{p}_i^o), \Psi(\mathbf{p}_j^o) \rangle, i \neq j \quad (2)$$

where $\Phi(\cdot)$ and $\Psi(\cdot)$ are projection functions for subjects and objects in the relationships respectively¹. This decomposition allows the score matrix $S = \{s_{ij}\}^{n \times n}$ to be computed *with only two projection processes for X^o followed by a matrix multiplication*. We use two multi-layer perceptrons (MLPs) with identical architecture (but different parameters) for $\Phi(\cdot)$ and $\Psi(\cdot)$. We also apply a sigmoid function element-wise to S such that all relatedness scores range from 0 to 1.

After obtaining the score matrix for all object pairs, we sort the scores in descending order and choose top K pairs. We then apply non-maximal suppression (NMS) to filter out object pairs that have significant overlap with others.

¹ We distinguish between the first and last object in a relationship as subject and object respectively, that is, $\langle \text{subject}, \text{relationship}, \text{object} \rangle$.

Each relationship has a pair of bounding boxes, and the combination order matters. We compute the overlap between two object pairs $\{u, v\}$ and $\{p, q\}$ as:

$$IoU(\{u, v\}, \{p, q\}) = \frac{I(r_u^o, r_p^o) + I(r_v^o, r_q^o)}{U(r_u^o, r_p^o) + U(r_v^o, r_q^o)} \tag{3}$$

where operator I computes the intersection area between two boxes and U the union area. The remaining m object pairs are considered as candidates having meaningful relationships \mathbf{E} . With \mathbf{E} , we obtain a graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, which is much sparser than the original fully connected graph. Along with the edges proposed for the graph, we get the visual representations $X^r = \{\mathbf{x}_1^r, \dots, \mathbf{x}_m^r\}$ for all m relationships by extracting features from the union box of each object pair.

3.3 Attentional GCN

To integrate contextual information informed by the graph structure, we propose an attentional graph convolutional network (aGCN). Before we describe our proposed aGCN, let us briefly recap a ‘vanilla’ GCN in which each node i has a representation $z_i \in \mathbb{R}^d$, as proposed in [13]. Briefly, for a target node i in the graph, the representations of its neighboring nodes $\{z_j \mid j \in \mathcal{N}(i)\}$ are first transformed via a learned linear transformation W . Then, these transformed representations are gathered with predetermined weights α , followed by a non-linear function σ (ReLU [25]). This layer-wise propagation can be written as:

$$z_i^{(l+1)} = \sigma \left(z_i^{(l)} + \sum_{j \in \mathcal{N}(i)} \alpha_{ij} W z_j^{(l)} \right) \tag{4}$$

or equivalently we can collect node representations into a matrix $Z \in \mathbb{R}^{d \times Tn}$

$$z_i^{(l+1)} = \sigma \left(W Z^{(l)} \alpha_i \right) \tag{5}$$

for $\alpha_i \in [0, 1]^n$ with 0 entries for nodes not neighboring i and $\alpha_{ii} = 1$. In conventional GCN, the connections in the graph are known and coefficient vector α_i are preset based on the symmetrically normalized adjacency matrix of features.

In this paper, we extend the conventional GCN to an attentional version, which we refer to as aGCN, by learning to adjust α . To predict attention from node features, we learn a 2-layer MLP over concatenated node features and compute a softmax over the resulting scores. The attention for node i is

$$u_{ij} = w_h^T \sigma(W_a [z_i^{(l)}, z_j^{(l)}]) \tag{6}$$

$$\alpha_i = \text{softmax}(u_i), \tag{7}$$

where w_h and W_a are learned parameters and $[\cdot, \cdot]$ is the concatenation operation. By definition, we set $\alpha_{ii} = 1$ and $\alpha_{ij} = 0 \forall j \notin \mathcal{N}(i)$. As attention is a function of node features, each iteration results in altered attentions which affects successive iterations.

aGCN for Scene Graph Generation. Recall that from the previous sections we have a set of N object regions and m relationships. From these, we construct a graph G with nodes corresponding to object and relationship proposals. We insert edges between relation nodes and their associated objects. We also add skip-connect edges directly between all object nodes. These connections allow information to flow directly between object nodes. Recent work has shown that reasoning about object correlation can improve detection performance [10]. We apply aGCN to this graph to update object and relationship representations based on global context.

Note that our graph captures a number of different types of connections (*i.e.* `object` \leftrightarrow `relationship`, `relationship` \leftrightarrow `subject` and `object` \leftrightarrow `object`). In addition, the information flow across each connection may be asymmetric (the informativeness of `subject` on `relationship` might be quite different from `relationship` to `subject`). We learn different transformations for each type and ordering – denoting the linear transform from node type a to node type b as W^{ab} with s =subjects, o =objects, and r =relationships. Using the same notation as in Eq. 5 and writing object and relationship features as Z^o and Z^r , we write the representation update for object nodes as

$$z_i^o = \sigma \left(\underbrace{W^{\text{skip}} Z^o \alpha^{\text{skip}}}_{\text{Message from Other Objects}} + \underbrace{W^{sr} Z^r \alpha^{sr} + W^{or} Z^r \alpha^{or}}_{\text{Messages from Neighboring Relationships}} \right) \quad (8)$$

with $\alpha_{ii}^{\text{skip}}=1$ and similarly for relationship nodes as

$$z_i^r = \sigma \left(z_i^r + \underbrace{W^{rs} Z^o \alpha^{rs} + W^{ro} Z^o \alpha^{ro}}_{\text{Messages from Neighboring Objects}} \right). \quad (9)$$

where α are computed at each iteration as in Eq. 7.

One open choice is how to initialize the object and relationship node representations z which could potentially be set to any intermediate feature representation or even the pre-softmax output corresponding to class labels. In practice, we run both a visual and semantic aGCN computation – one with visual features and the other using pre-softmax outputs. In this way, we can reason about both lower-level visual details (*i.e.* two people are likely talking if they are facing one another) as well as higher-level semantic co-occurrences (*i.e.* cars have wheels). Further, we set the attention in the semantic aGCN to be that of the visual aGCN – effectively modulating the flow of semantic information based on visual cues. This also enforces that real-world objects and relationships represented in both graphs interact with others in the same manner.

3.4 Loss Function

In Graph R-CNN, we factorize the scene graph generation process into three sub-processes: $P(\mathbf{R}, \mathbf{O} | \mathbf{V}, \mathbf{E}, \mathbf{I})$, $P(\mathbf{E} | \mathbf{V}, \mathbf{I})$, $P(\mathbf{V} | \mathbf{I})$, which were described above. During training, each of these sub-processes are trained with supervision. For $P(\mathbf{V} | \mathbf{I})$, we use the same loss as used in RPN, which consists of a binary cross

entropy loss on proposals and a regression loss for anchors. For $P(\mathbf{E}|\mathbf{V}, \mathbf{I})$, we use another binary cross entropy loss on the relation proposals. For the final scene graph generation $P(\mathbf{R}, \mathbf{O}|\mathbf{V}, \mathbf{E}, \mathbf{I})$, two multi-class cross entropy losses are used for object classification and predicate classification.

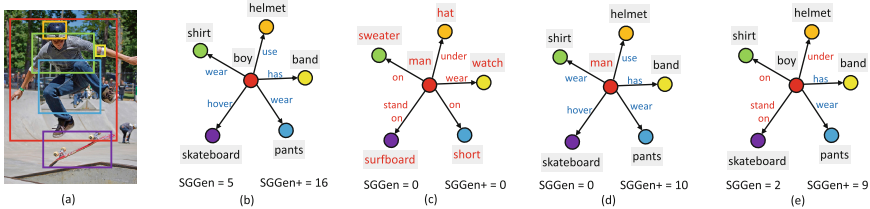


Fig. 3. A example to demonstrate the difference between *SGGGen* and *SGGGen+*. Given the input image (a), its ground truth scene graph is depicted in (b). (c)–(e) are three generated scene graphs. For clarity, we merely show the connections with *boy*. At the bottom of each graph, we compare the number of correct predictions for two metrics.

4 Evaluating Scene Graph Generation

Scene graph generation is naturally a structured prediction problem over attributed graphs, and how to correctly and efficiently evaluate predictions is an under-examined problem in prior work on scene graph generation. We note that graph similarity based on minimum graph edit distance has been well-studied in graph theory [5]; however, computing exact solution is NP-complete and approximation APX-hard [20].

Prior work has circumvented these issues by evaluating scene graph generation under a simple triplet-recall based metric introduced in [40]. Under this metric which we will refer to as *SGGGen*, the ground truth scene graph is represented as a set of $\langle \text{object}, \text{relationship}, \text{subject} \rangle$ triplets and recall is computed via exact match. That is to say, a triplet is considered ‘matched’ in a generated scene graph if all three elements have been correctly labeled, and both *object* and *subject* nodes have been properly localized (*i.e.*, bounding box IoU > 0.5). While simple to compute, this metric results in some unintuitive notions of similarity that we demonstrate in Fig. 3.

Figure 3a shows an input image overlaid with bounding box localizations of correspondingly colored nodes in the ground truth scene graph shown in (b). (c), (d), and (e) present erroneously labeled scene graphs corresponding to these same localizations. Even a casual examination of (c) and (d) yields the stark difference in their accuracy – while (d) has merely mislabeled the boy as a man, (c) has failed to accurately predict even a single node or relationship! Despite these differences, neither recalls a single complete triplet and are both scored identically under *SGGGen* (*i.e.*, 0).

To address this issue, we propose a new metric called *SGGGen+* as the augmentation of *SGGGen*. *SGGGen+* not only considers the triplets in the graph, but

also the singletons (object and predicate). The computation of **SGGen+** can be formulated as:

$$Recall = \frac{C(O) + C(P) + C(T)}{N} \quad (10)$$

where $C(\cdot)$ is a counting operation, and hence $C(O)$ is the number of object nodes correctly localized and recognized; $C(P)$ is for predicate. Since the location of predicate depends on the location of subject and object, only if both subject and object are correctly localized and the predicate is correctly recognized, we will count it as one. $C(T)$ is for triplet, which is the same as **SGGen**. Here, N is the number of entries (the sum of number of objects, predicates and relationships) in the ground truth graph. In Fig. 3, using our **SGGen+**, the recall for graph (c) is still 0, since all predictions are wrong. However, the recall for graph (d) is not 0 anymore since most of the object and all predicate predictions are correct, except for one wrong prediction for the red node. Based on our new metric, we can obtain a much comprehensive measurement of scene graph similarity.

5 Experiments

Recently, there are some inconsistencies in existing work on scene graph generation in terms of data preprocessing, data split, and evaluation. This makes it difficult to systematically benchmark progress and cleanly compare numbers across papers. So we first clarify the details of our experimental settings.

Datasets. There are a number of splits of the Visual Genome dataset that have been used in the scene graph generation literature [18, 40, 45]. The most commonly used is the one proposed in [40]. Hence, in our experiments, we follow their preprocessing strategy and dataset split. After preprocessing, the dataset is split into training and test sets, which contains 75,651 images and 32,422 images, respectively. In this dataset, the top-frequent 150 object classes and 50 relation classes are selected. Each image has around 11.5 objects and 6.2 relationships in the scene graph.

Training. For training, multiple strategies have been used in literature. In [18, 26, 40], the authors used two-stage training, where the object detector is pre-trained, followed by the joint training of the whole scene graph generation model. To be consistent with previous work [18, 40], we also adopt the two-stage training – we first train the object detector and then train the whole model jointly until convergence.

Metrics. We use four metrics for evaluating scene graph generation, including three previously used metrics and our proposed **SGGen+** metric:

- **Predicate Classification (PredCls):** The performance for recognizing the relation between two objects given the ground truth locations.
- **Phrase Classification (PhrCls):** The performance for recognizing two object categories and their relation given the ground truth locations.

- **Scene Graph Generation (SGGen)**: The performance for detecting objects (IoU > 0.5) *and* recognizing the relations between object pairs.
- **Comprehensive Scene Graph Generation (SGGen+)**: Besides the triplets counted by SGGen, it considers the singletons and pairs (if any), as described earlier.

Evaluation. In our experiments, we multiply the classification scores for subjects, objects and their relationships, then sort them in descending order. Based on this order, we compute the recall at top 50 and top 100, respectively. Another difference in existing literature in the evaluation protocol is w.r.t. the **PhrCls** and **PredCls** metrics. Some previous works [18, 26] used different models to evaluate along different metrics. However, such a comparison is unfair since the models could be trained to overfit the respective metrics. For meaningful evaluation, we evaluate a single model – the one obtained after joint training – across all metrics.

5.1 Implementation Details

We use Faster R-CNN [32] associated with VGG16 [33] as the backbone based on the PyTorch re-implementation [41]. During training, the number of proposals from RPN is 256. For each proposal, we perform ROI Align [8] pooling, to get a 7×7 response map, which is then fed to a two-layer MLP to obtain each proposal’s representation. In RePN, the projection functions $\Phi(\cdot)$ and $\Psi(\cdot)$ are simply two-layer MLPs. During training, we sample 128 object pairs from the quadratic number of candidates. We then obtain the union of boxes of the two objects and extract a representation for the union. The threshold for box-pair NMS is 0.7. In aGCN, to obtain the attention for one node pair, we first project the object/predicate features into 256-d and then concatenate them into 512-d, which is then fed to a two-layer MLP with a 1-d output. For aGCN, we use two aGCN layers at the feature level and semantic level, respectively. The attention on the graph is updated in each aGCN layer at the feature level, which is then fixed and sent to the aGCN at the semantic level.

Training. As mentioned, we perform stage-wise training – we first pretrain Faster R-CNN for object detection, and then fix the parameters in the backbone to train the scene graph generation model. SGD is used as the optimizer, with initial learning rate $1e-2$ for both training stages.

5.2 Analysis on New Metric

We first quantitatively demonstrate the difference between our proposed metric **SGGen+** and **SGGen**. We compare them by perturbing ground truth scene graphs. We consider assigning random incorrect labels to objects; perturbing objects 1) without relationships, 2) with relationships, and 3) both. We vary the fraction of nodes which are perturbed among {20%, 50%, 100%}. Recall is reported for both metrics. As shown in Table 1, **SGGen** is completely insensitive to the perturbation

of objects without relationships (staying at 100 consistently) since it only considers relationship triplets. Note that there are on average 50.1% objects without relationships in the dataset, which **SGGen** omits. On the other hand, **SGGen** is overly sensitive to label errors on objects with relationships (reporting 54.1 at only 20% perturbation where the overall scene graph is still quite accurate). Note that even at 100% perturbation the object localizations and relationships are still correct such that **SGGen+** provides a non-zero score, unlike **SGGen** which considers the graph entirely wrong. Overall, we hope this analysis demonstrates that **SGGen+** is more comprehensive compared to **SGGen**.

Table 1. Comparisons between **SGGen** and **SGGen+** under different perturbations.

Perturb Type	none	w/o relationship				w/ relationship			both		
Perturb Ratio	0%	20%	50%	100%	20%	50%	100%	20%	50%	100%	
SGGen	100.0	100.0	100.0	100.0	54.1	22.1	0.0	62.2	24.2	0.0	
SGGen+	100.0	94.5	89.1	76.8	84.3	69.6	47.9	80.1	56.6	22.8	

5.3 Quantitative Comparison

We compare our Graph R-CNN with recent proposed methods, including Iterative Message Passing (IMP) [40], Multi-level scene Description Network (MSDN) [18]. Furthermore, we evaluate the neural motif frequency baseline proposed in [42]. Note that previous methods often use slightly different pre-training procedures or data split or extra supervisions. For a fair comparison and to control for such orthogonal variations, we reimplemented IMP, MSDN and frequency baseline in our codebase. Then, we re-train IMP and MSDN based on our backbone – specifically, we used the same pre-trained object detector, and then jointly train the scene graph generator until convergence. We denote these as IMP[†] and MSDN[†]. Using the same pre-trained object detector, we report the neural motif frequency baseline in [42] as NM-Freq[†].

Table 2. Comparison on Visual Genome test set [14]. We reimplemented IMP [40] and MSDN [18] using the same object detection backbone for fair comparison.

Method	SGGen+		SGGen		PhrCls		PredCls	
	R@50	R@100	R@50	R@100	R@50	R@100	R@50	R@100
IMP [40]	-	-	3.4	4.2	21.7	24.4	44.8	53.0
MSDN [18]	-	-	7.7	10.5	19.3	21.8	63.1	66.4
Pixel2Graph [26]	-	-	9.7	11.3	26.5	30.0	68.0	75.2
IMP [†] [40]	25.6	27.7	6.4	8.0	20.6	22.4	40.8	45.2
MSDN [†] [18]	25.8	28.2	7.0	9.1	27.6	29.9	53.2	57.9
NM-Freq [†] [42]	26.4	27.8	6.9	9.1	23.8	27.2	41.8	48.8
Graph R-CNN (Us)	28.5	35.9	11.4	13.7	29.6	31.6	54.2	59.1

We report the scene graph generation performance in Table 2. The top three rows are numbers reported in the original paper, and the bottom four rows are the numbers from our re-implementations. First, we note that our re-implementations of IMP and MSDN (IMP[†] and MSDN[†]) result in performance that is close to or better than the originally reported numbers under some metrics (but not all), which establishes that the takeaway messages next are indeed due to our proposed architectural choices – relation proposal network and attentional GCNs. Next, we notice that Graph R-CNN outperforms IMP[†] and MSDN[†]. This indicates that our proposed Graph R-CNN model is more effective to extract the scene graph from images. Our approach also outperforms the frequency baseline on all metrics, demonstrating that our model has not just learned simple co-occurrence statistics from training data, but rather also captures context in individual images. More comprehensively, we compare with IMP and MSDN on the efficiency over training and inference. IMP uses 2.15× while MSDN uses 1.86× our method. During inference, IMP is 3.27× while MSDN is 3.80× slower than our Graph R-CNN. This is mainly due to the simplified architecture design (especially the aGCN for context propagation) in our model.

5.4 Ablation Study

In Graph R-CNN, we proposed two novel modules – relation proposal network (RePN) and attentional GCNs (aGCN). In this sub-section, we perform ablation studies to get a clear sense of how these different components affect the final performance. The left-most columns in Table 3 indicate whether or not we used

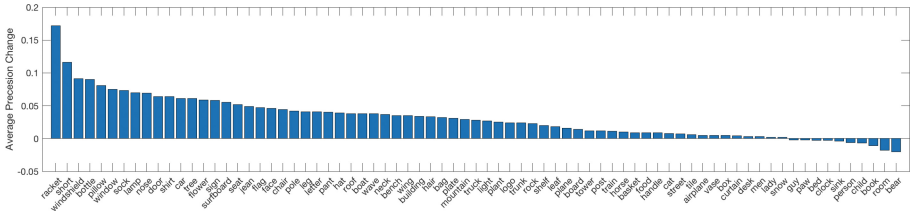


Fig. 4. Per category object detection performance change after adding RePN.

Table 3. Ablation studies on Graph R-CNN. We report the performance based on four scene graph generation metrics and the object detection performance in mAP@0.5.

RePN	GCN	aGCN	Detection	SGGen+		SGGen		PhrCls		PredCls	
			mAP@0.5	R@50	R@100	R@50	R@100	R@50	R@100	R@50	R@100
-	-	-	20.4	25.9	27.9	6.1	7.9	17.8	19.9	33.5	38.4
✓	-	-	23.6	27.6	34.8	8.7	11.1	18.3	20.4	34.5	39.5
✓	✓	-	23.4	28.1	35.3	10.8	13.4	27.2	29.5	52.3	57.2
✓	-	✓	23.0	28.5	35.9	11.4	13.7	29.4	31.6	54.2	59.1

RePN, GCN, and attentional GCN (aGCN) in our approach. The results are reported in the remaining columns of Table 3. We also report object detection performance $mAP@0.5$ following Pascal VOC’s metric [4].

In Table 3, we find RePN boosts *SGGen* and *SGGen+* significantly. This indicates that our RePN can effectively prune the spurious connections between objects to achieve high recall for the correct relationships. We also notice it improves object detection significantly. In Fig. 4 we show the per category object detection performance change when RePN is added. For visual clarity, we dropped every other column when producing the plot. We can see that almost all object categories improve after adding RePN. Interestingly, we find the detection performance on categories like *racket*, *short*, *windshield*, *bottle* are most significantly improved. Note that many of these classes are smaller objects that have strong relationships with other objects, e.g. rackets are often carried by people. Evaluating *PhrCls* and *PredCls* involves using the ground truth object locations. Since the number of objects in images (typically <25) is much less than the number of object proposals (64), the number of relation pairs is already very small. As a result, RePN has less effect on these two metrics.

By adding the aGCNs into our model, the performance is further improved. These improvements demonstrate that the aGCN in our Graph R-CNN can capture meaningful context across the graph. We also compare the performance of our model with and without attention. We see that by adding attention on top of GCNs, the performance is higher. This indicates that controlling the extent to which contextual information flows through the edges is important. These results align with our intuitions mentioned in the introduction. Figure 5 shows generated scene graphs for test images. With RePN and aGCN, our model is able to generate higher recall scene graphs. The green ellipsoids shows the correct relationship predictions in the generated scene graph.

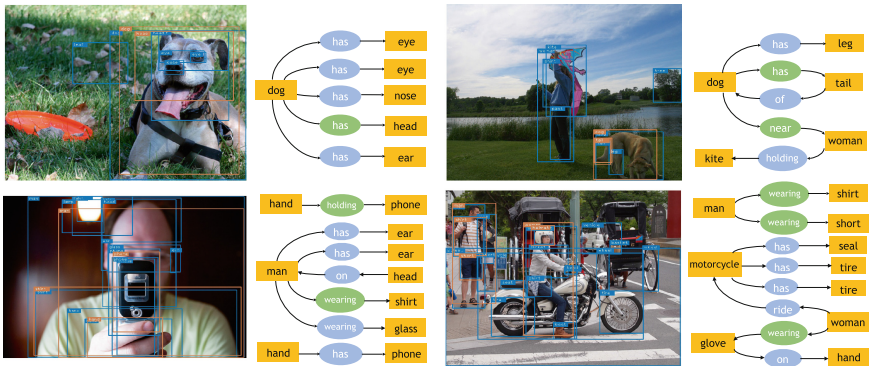


Fig. 5. Qualitative results from Graph R-CNN. In images, blue and orange bounding boxes are ground truths and correct predictions, respectively. In scene graphs, blue ellipsoids are ground truth relationships while green ones denote correct predictions.

6 Conclusion

In this work, we introduce a new model for scene graph generation – Graph R-CNN. Our model includes a relation proposal network (RePN) that efficiently and intelligently prunes out pairs of objects that are unlikely to be related, and an attentional graph convolutional network (aGCN) that effectively propagates contextual information across the graph. We also introduce a novel scene graph generation evaluation metric (SGGen+) that is more fine-grained and realistic than existing metrics. Our approach outperforms existing methods for scene graph generation, as evaluated using existing metrics and our proposed metric.

Acknowledgements. This work was supported in part by NSF, AFRL, DARPA, Siemens, Google, Amazon, ONR YIPs and ONR Grants N00014-16-1-{2713,2793}.

References

1. Antol, S., et al.: VQA: visual question answering. In: ICCV, pp. 2425–2433 (2015)
2. Dai, B., Zhang, Y., Lin, D.: Detecting visual relationships with deep relational networks. In: CVPR (2017)
3. Das, A., et al.: Visual dialog. In: CVPR (2017)
4. Everingham, M., Van Gool, L., Williams, C., Winn, J., Zisserman, A.: The Pascal visual object classes challenge 2012 results, vol. 5 (2012). <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>
5. Gao, X., Xiao, B., Tao, D., Li, X.: A survey of graph edit distance. *Pattern Anal. Appl.* **13**(1), 113–129 (2010)
6. Girshick, R.: Fast R-CNN. In: CVPR (2015)
7. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR (2014)
8. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: ICCV (2017)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
10. Hu, H., Gu, J., Zhang, Z., Dai, J., Wei, Y.: Relation networks for object detection. In: CVPR (2018)
11. Johnson, J., et al.: CLEVR: a diagnostic dataset for compositional language and elementary visual reasoning. In: CVPR (2017)
12. Johnson, J., et al.: Image retrieval using scene graphs. In: CVPR (2015)
13. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017)
14. Krishna, R., et al.: Visual genome: Connecting language and vision using crowd-sourced dense image annotations. *IJCV* **123**(1), 32–73 (2017)
15. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: NIPS (2012)
16. Ladicky, L., Russell, C., Kohli, P., Torr, P.H.S.: Graph cut based inference with co-occurrence statistics. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010. LNCS, vol. 6315, pp. 239–253. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15555-0_18
17. Li, Y., Ouyang, W., Wang, X.: ViP-CNN: a visual phrase reasoning convolutional neural network for visual relationship detection. In: CVPR (2017)

18. Li, Y., Ouyang, W., Zhou, B., Wang, K., Wang, X.: Scene graph generation from objects, phrases and region captions. In: ICCV (2017)
19. Liang, X., Lee, L., Xing, E.P.: Deep variation-structured reinforcement learning for visual relationship and attribute detection. In: CVPR (2017)
20. Lin, C.-L.: Hardness of approximating graph transformation problem. In: Du, D.-Z., Zhang, X.-S. (eds.) ISAAC 1994. LNCS, vol. 834, pp. 74–82. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-58325-4_168
21. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017)
22. Liu, W., et al.: SSD: single shot multibox detector. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 21–37. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_2
23. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L.: Visual relationship detection with language priors. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 852–869. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_51
24. Lu, J., Yang, J., Batra, D., Parikh, D.: Neural baby talk. In: CVPR (2018)
25. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: ICML (2010)
26. Newell, A., Deng, J.: Pixels to graphs by associative embedding. In: NIPS (2017)
27. Oliva, A., Torralba, A.: The role of context in object recognition. *Trends Cogn. Sci.* **11**(12), 520–527 (2007)
28. Parikh, D., Zitnick, C.L., Chen, T.: From appearance to context-based recognition: Dense labeling in small images. In: CVPR (2008)
29. Peyre, J., Laptev, I., Schmid, C., Sivic, J.: Weakly-supervised learning of visual relations. In: ICCV (2017)
30. Rabinovich, A., Vedaldi, A., Galleguillos, C., Wiewiora, E., Belongie, S.: Objects in context. In: ICCV (2007)
31. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: unified, real-time object detection. In: CVPR (2016)
32. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS (2015)
33. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
34. Szegedy, C., et al.: Going deeper with convolutions. In: CVPR (2015)
35. Teney, D., Liu, L., Hengel, A.V.d.: Graph-structured representations for visual question answering. In: CVPR (2017)
36. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: ICLR (2018)
37. Wang, P., Wu, Q., Shen, C., Dick, A., van den Hengel, A.: FVQA: fact-based visual question answering. In: PAMI (2017)
38. Wang, P., Wu, Q., Shen, C., van den Hengel, A.: The VQA-machine: learning how to use existing vision algorithms to answer new questions. In: CVPR (2017)
39. Wu, Q., Shen, C., Wang, P., Dick, A., van den Hengel, A.: Image captioning and visual question answering based on attributes and external knowledge. In: PAMI (2017)
40. Xu, D., Zhu, Y., Choy, C.B., Fei-Fei, L.: Scene graph generation by iterative message passing. In: CVPR (2017)
41. Yang, J., Lu, J., Batra, D., Parikh, D.: A faster Pytorch implementation of faster R-CNN (2017). <https://github.com/jwyang/faster-rcnn.pytorch>

42. Zellers, R., Yatskar, M., Thomson, S., Choi, Y.: Neural motifs: Scene graph parsing with global context. In: CVPR (2018)
43. Zhang, H., Kyaw, Z., Chang, S.F., Chua, T.S.: Visual translation embedding network for visual relation detection. In: CVPR (2017)
44. Zhang, H., Kyaw, Z., Yu, J., Chang, S.F.: Ppr-fcn: weakly supervised visual relation detection via parallel pairwise r-fcn (2017)
45. Zhang, J., Elhoseiny, M., Cohen, S., Chang, W., Elgammal, A.: Relationship proposal networks. In: CVPR (2017)
46. Zhuang, B., Liu, L., Shen, C., Reid, I.: Towards context-aware interaction recognition for visual relationship detection. In: ICCV (2017)