



Bridging Knowledge Graphs to Generate Scene Graphs

Alireza Zareian^(✉), Svebor Karaman, and Shih-Fu Chang

Columbia University, New York, NY 10027, USA
{az2407,sk4089,sc250}@columbia.edu

Abstract. Scene graphs are powerful representations that parse images into their abstract semantic elements, *i.e.*, objects and their interactions, which facilitates visual comprehension and explainable reasoning. On the other hand, commonsense knowledge graphs are rich repositories that encode how the world is structured, and how general concepts interact. In this paper, we present a unified formulation of these two constructs, where a scene graph is seen as an image-conditioned instantiation of a commonsense knowledge graph. Based on this new perspective, we reformulate scene graph generation as the inference of a bridge between the scene and commonsense graphs, where each entity or predicate instance in the scene graph has to be linked to its corresponding entity or predicate class in the commonsense graph. To this end, we propose a novel graph-based neural network that iteratively propagates information between the two graphs, as well as within each of them, while gradually refining their bridge in each iteration. Our Graph Bridging Network, GB-NET, successively infers edges and nodes, allowing to simultaneously exploit and refine the rich, heterogeneous structure of the interconnected scene and commonsense graphs. Through extensive experimentation, we showcase the superior accuracy of GB-NET compared to the most recent methods, resulting in a new state of the art. We publicly release the source code of our method (<https://github.com/alirezazareian/gbnet>).

1 Introduction

Extracting structured, symbolic, semantic representations from data has a long history in Natural Language Processing (NLP), under the umbrella terms *semantic parsing* at the sentence level [8,9] and *information extraction* at the document level [22,41]. The resulting *semantic graphs* or *knowledge graphs* have many applications such as question answering [7,17] and information retrieval [6,50]. In computer vision, Xu *et al.* have recently called attention to the task of Scene Graph Generation (SGG) [44], which aims at extracting a symbolic, graphical representation from a given image, where every node corresponds to a

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-58592-1_36) contains supplementary material, which is available to authorized users.

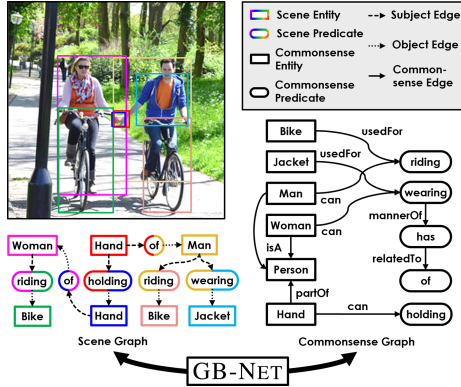


Fig. 1. Left: An example of a Visual Genome image and its ground truth scene graph. Right: A relevant portion of the commonsense graph. In this paper we formulate the task of Scene Graph Generation as the problem of creating a bridge between these two graphs. Such bridge not only classifies each scene entity and predicate, but also creates an inter-connected heterogeneous graph whose rich structure is exploited by our method (GB-NET).

localized and categorized object (entity), and every edge encodes a pairwise interaction (predicate). This has inspired two lines of follow-up work, some improving the performance on SGG [2, 10, 11, 23, 24, 31, 43, 47, 52], and others exploiting such rich structures for down-stream tasks such as Visual Question Answering (VQA) [12, 38, 39, 54], image captioning [48, 49], image retrieval [15, 37], and image synthesis [14]. In VQA for instance, SGG not only improves performance, but also promotes interpretability and enables explainable reasoning [38].

Although several methods have been proposed, the state-of-the-art performance for SGG is still far from acceptable. Most recently, [2] achieves only 16% mean recall, for matching the top 100 predicted subject-predicate-object triples against ground truth triples. This suggests the current SGG methods are insufficient to address the complexity of this task. Recently, a few papers have attempted to use external *commonsense* knowledge to advance SGG [2, 10, 52], as well as other domains [3, 16]. This commonsense can range from curated knowledge bases such as ConceptNet [27], ontologies such as WordNet [30], or automatically extracted facts such as co-occurrence frequencies [52]. The key message of those works is that a prior knowledge about the world can be very helpful when perceiving a complex scene. If we know the relationship of a **Person** and a **Bike** is most likely **riding**, we can more easily disambiguate between **riding**, **on**, and **attachedTo**, and classify their relationship more accurately. Similarly, if we know a **Man** and a **Woman** are both sub-types of **Person**, even if we only see **Man-riding-Bike** in training data, we can generalize and recognize a **Woman-riding-Bike** triplet at test time. Although this idea is intuitively promising, existing methods that implement it have major limitations, as detailed in Sect. 2, and we address those in the proposed method.

More specifically, recent methods either use ad-hoc heuristics to integrate limited types of commonsense into the scene graph generation process [2, 52], or fail to exploit the rich, graphical structure of commonsense knowledge [10]. To devise a general framework for incorporating any type of graphical knowledge into the process of scene understanding, we take inspiration from early works on knowledge representation and applying structured grammars to computer vision problems [32, 40, 55], and redefine those concepts in the light of the recent advances in graph-based deep learning. Simply put, we formulate both scene and commonsense graphs as knowledge graphs with entity and predicate nodes, and various types of edges. A scene graph node represents an entity or predicate *instance* in a specific image, while a commonsense graph node represents an entity or predicate *class*, which is a general concept independent of the image. Similarly, a scene graph edge indicates the participation of an entity instance (*e.g.* as a subject or object) in a predicate instance in a scene, while a commonsense edge states a general fact about the interaction of two concepts in the world. Figure 1 shows an example scene graph and commonsense graph side by side.

Based on this unified perspective, we reformulate the problem of scene graph generation from entity and predicate classification into the novel problem of bridging those two graphs. More specifically, we propose a method that given an image, initializes potential entity and predicate nodes, and then classifies each node by connecting it to its corresponding class node in the commonsense graph, through an edge we call a *bridge*. This establishes a connectivity between instance-level, visual knowledge and generic, commonsense knowledge. To incorporate the rich combination of visual and commonsense information in the SGG process, we propose a novel graphical neural network, that iteratively propagates messages between the scene and commonsense graphs, as well as within each of them, while gradually refining the bridge in each iteration. Our Graph Bridging Network, GB-NET, successively infers edges and nodes, allowing to simultaneously exploit and refine the rich, heterogeneous structure of the interconnected scene and commonsense graphs.

To evaluate the effectiveness of our method, we conduct extensive experiments on the Visual Genome [20] dataset. The proposed GB-NET outperforms the state of the art consistently in various performance metrics. Through ablation studies, we show how each of the proposed ideas contribute to the results. We also publicly release a comprehensive software package based on [52] and [2], to reproduce all the numbers reported in this paper. We provide further quantitative, qualitative, and speed analysis in our Supplementary Material, as well as additional implementation details.

2 Related Work

2.1 Scene Graph Generation

Most SGG methods are based on an object detection backbone that extracts region proposals from the input image. They utilize some kind of information

propagation module to incorporate context, and then classify each region to an object class, as well as each pair of regions to a relation class [2, 44, 47, 52]. Our method has two key differences with this conventional process: firstly, our information propagation network operates on a larger graph which consists of not only object nodes, but also predicate nodes and commonsense graph nodes, and has a more complex structure. Secondly, we do not classify each object and relation using classifiers, but instead use a pairwise matching mechanism to connect them to corresponding class nodes in the commonsense graph.

More recently, a few methods [2, 10, 52] have used external knowledge to enhance scene graph generation. This external knowledge is sometimes referred to as “commonsense”, because it encodes ontological knowledge about classes, rather than specific instances. Despite encouraging results, these methods have major limitations. Specifically, [52] used triplet frequency to bias the logits of their predicate classifier, and [2] used such frequencies to initialize edge weights on their graphs. Such external priors have been also shown beneficial for recognizing objects [45, 46] and relationships [26, 53], that are building blocks for SGG. Nevertheless, neither of those methods can incorporate other types or knowledge, such as the semantic hierarchy concepts, or object affordances. Gu *et al.* [10] propose a more general way to incorporate knowledge in SGG, by retrieving a set of relevant facts for each object from a pool of commonsense facts. However, their method does not utilize the structure of the commonsense graph, and treats knowledge as a set of triplets. Our method considers commonsense as a general graph with several types of edges, explicitly integrates that graph with the scene graph by connecting corresponding nodes, and incorporates the rich structure of commonsense by graphical message passing.

2.2 Graph-Based Neural Networks

By Graph-based Neural Networks (GNN), we refer to the family of models that take a graph as input, and iteratively update the representation of each node by applying a learnable function (a.k.a., message) on the node’s neighbors. Graph Convolutional Networks (GCN) [19], Gated Graph Neural Networks (GGNN) [25], and others are all specific implementations of this general model. Most SGG methods use some variant of GNNs to propagate information between region proposals [2, 24, 44, 47]. Our message passing method, detailed in Sect. 4, resembles GGNN but instead of propagating messages through a static graph, we update (some) edges as well. Few methods exist that dynamically update edges during message passing [35, 51], but we are the first to refine edges between a scene graph and an external knowledge graph.

Apart from SGG, GNNs have been used in several other computer vision tasks, often in order to propagate context information across different objects in a scene. For instance, [28] injects a GNN into a Faster R-CNN [36] framework to contextualize the features of region proposals before classifying them. This improves the results since the presence of a **table** can affect the detection of a **chair**. On the other hand, some methods utilize GNNs on graphs that represent the ontology of concepts, rather than objects in a scene [16, 21, 29, 42].

This often enables generalization to unseen or infrequent concepts by incorporating their relationship with frequently seen concepts. More similarly to our work, Chen *et al.* [3] were the first to bring those two ideas together, and form a graph by objects in an image as well as object classes in the ontology. Nevertheless, the class nodes in that work were merely an auxiliary means to improve object features before classification. In contrast, we classify the nodes by explicitly inferring their connection to their corresponding class nodes. Moreover, we iteratively refine the bridge between scene and commonsense graphs to enhance our prediction. Furthermore, their task only involves objects and object classes, while we explore a more complex structure where predicates play an important role as well.

3 Problem Formulation

In this section, we first formalize the concepts of knowledge graph in general, and commonsense graph and scene graph in particular. Leveraging their similarities, we then reformulate the problem of scene graph generation as bridging these two graphs.

3.1 Knowledge Graphs

We define a knowledge graph as a set of entity and predicate nodes ($\mathcal{N}_E, \mathcal{N}_P$), each with a semantic label, and a set of directed, weighted edges \mathcal{E} from a predefined set of types. Denoting by Δ a node type (here, either entity E or predicate P), the set of edges encoding the relation r between nodes of type Δ and Δ' is defined as

$$\mathcal{E}_r^{\Delta \rightarrow \Delta'} \subseteq \mathcal{N}_\Delta \times \mathcal{N}_{\Delta'} \rightarrow \mathbb{R}. \quad (1)$$

A *commonsense graph* is a type of knowledge graph in which each node represents the general concept of its semantic label, and hence each semantic label (entity or predicate class) appears in exactly one node. In such a graph, each edge encodes a relational fact involving a pair of concepts, such as `Hand-partOf-Person` and `Cup-usedFor-Drinking`. Formally, we define the set of commonsense entity (CE) nodes \mathcal{N}_{CE} and commonsense predicate (CP) nodes \mathcal{N}_{CP} as all entity and predicate classes in our task. Commonsense edges \mathcal{E}_C consist of 4 distinct subsets, depending on the source and destination node type:

$$\begin{aligned} \mathcal{E}_C = & \{ \mathcal{E}_r^{CE \rightarrow CP} \} \cup \{ \mathcal{E}_r^{CP \rightarrow CE} \} \cup \\ & \{ \mathcal{E}_r^{CE \rightarrow CE} \} \cup \{ \mathcal{E}_r^{CP \rightarrow CP} \}. \end{aligned} \quad (2)$$

A *scene graph* is a different type of knowledge graph where: (a) each scene entity (SE) node is associated with a bounding box, referring to an image region, (b) each scene predicate (SP) node is associated with an ordered pair of SE nodes, namely a subject and an object, and (c) there are two types of undirected edges which connect each SP to its corresponding subject and object respectively.

Here because we define knowledge edges to be directed, we model each undirected subject or object edge as two directed edges in the opposite directions, each with a distinct type. More specifically,

$$\begin{aligned}
 \mathcal{N}_{SE} &\subseteq [0, 1]^4 \times \mathcal{N}_{CE}, \\
 \mathcal{N}_{SP} &\subseteq \mathcal{N}_{SE} \times \mathcal{N}_{SE} \times \mathcal{N}_{CP}, \\
 \mathcal{E}_S &= \{ \mathcal{E}_{\text{subjectOf}}^{\text{SE} \rightarrow \text{SP}}, \mathcal{E}_{\text{objectOf}}^{\text{SE} \rightarrow \text{SP}}, \\
 &\quad \mathcal{E}_{\text{hasSubject}}^{\text{SP} \rightarrow \text{SE}}, \mathcal{E}_{\text{hasObject}}^{\text{SP} \rightarrow \text{SE}} \},
 \end{aligned} \tag{3}$$

where $[0, 1]^4$ is the set of possible bounding boxes, and $\mathcal{N}_{SE} \times \mathcal{N}_{SE} \times \mathcal{N}_{CP}$ is the set of all possible triples that consist of two scene entity nodes and a scene predicate node. Figure 1 shows an example of scene graph and commonsense graph side by side, to make their similarities clearer. Here we assume every scene graph node has a label that exists in the commonsense graph, since in reality some objects and predicates might belong to background classes, we consider a special commonsense node as background entity and another for background predicate.

3.2 Bridging Knowledge Graphs

Considering the similarity between the commonsense and scene graph formulations, we make a subtle refinement in the formulation to bridge these two graphs. Specifically, we remove the class from SE and SP nodes and instead encode it into a set of *bridge* edges \mathcal{E}_B that connect each SE or SP node to its corresponding class, *i.e.*, a CE or CP node respectively:

$$\begin{aligned}
 \mathcal{N}_{SE}^? &\subseteq [0, 1]^4, \\
 \mathcal{N}_{SP}^? &\subseteq \mathcal{N}_{SE} \times \mathcal{N}_{SE}, \\
 \mathcal{E}_B &= \{ \mathcal{E}_{\text{classifiedTo}}^{\text{SE} \rightarrow \text{CE}}, \mathcal{E}_{\text{classifiedTo}}^{\text{SP} \rightarrow \text{CP}}, \\
 &\quad \mathcal{E}_{\text{hasInstance}}^{\text{CE} \rightarrow \text{SE}}, \mathcal{E}_{\text{hasInstance}}^{\text{CP} \rightarrow \text{SP}} \},
 \end{aligned} \tag{4}$$

where $.^?$ means the nodes are implicit, *i.e.*, their classes are unknown. Each edge of type `classifiedTo`, connects an entity or predicate to its corresponding label in the commonsense graph, and has a reverse edge of type `hasInstance` which connects the commonsense node back to the instance. Based on this reformulation, we can define the problem of SGG as the extraction of implicit entity and predicate nodes from the image (hereafter called *scene graph proposal*), and then classifying them by connecting each entity or predicate to the corresponding node in the commonsense graph. Accordingly, Given an input image I and a provided and fixed commonsense graph, the goal of SGG with commonsense knowledge is to maximize

$$\begin{aligned}
 &p(\mathcal{N}_{SE}, \mathcal{N}_{SP}, \mathcal{E}_S | I, \mathcal{N}_{CE}, \mathcal{N}_{CP}, \mathcal{E}_C) \\
 &= p(\mathcal{N}_{SE}^?, \mathcal{N}_{SP}^?, \mathcal{E}_S | I) \times \\
 &p(\mathcal{E}_B | I, \mathcal{N}_{CE}, \mathcal{N}_{CP}, \mathcal{E}_C, \mathcal{N}_{SE}^?, \mathcal{N}_{SP}^?, \mathcal{E}_S).
 \end{aligned} \tag{5}$$

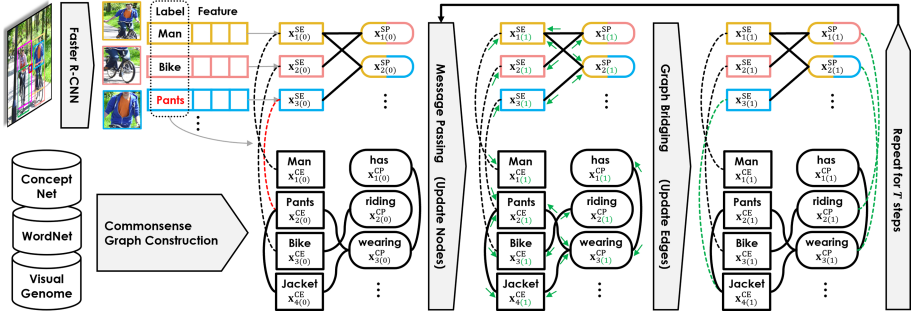


Fig. 2. An illustrative example of the GB-NET process. First, we initialize the scene graph and entity bridges using a Faster R-CNN. Then we propagate messages to update node representations, and use them to update the entity and predicate bridges. This is repeated T times and the final bridge determines the output label of each node.

In this paper, the first term is implemented as a region proposal network that infers $\mathcal{N}_{SE}^?$ given the image, followed by a simple predicate proposal algorithm that considers all possible entity pairs as $\mathcal{N}_{SP}^?$. The second term is fulfilled by the proposed GB-NET which infers bridge edges by incorporating the rich structure of the scene and commonsense graphs. Note that unlike most existing methods [2, 52], we do not factorize this into predicting entity classes given the image, and then predicate classes given entities. Therefore, our formulation is more general and allows the proposed method to classify entities and predicates jointly.

4 Method

The proposed method is illustrated in Fig. 2. Given an image, our model first applies a Faster R-CNN [36] to detect objects, and represents them as scene entity (SE) nodes. It also creates a scene predicate (SP) node for each pair of entities, which forms a scene graph proposal, yet to be classified. Given this graph and a background commonsense graph, each with fixed internal connectivity, our goal is to create *bridge* edges between the two graphs that connect each instance (SE and SP node) to its corresponding class (CE and CP node). To this end, our model initializes entity bridges by connecting each SE to the CE that matches the label predicted by Faster R-CNN, and propagates messages among all nodes, through every edge type with dedicated message passing parameters. Given the updated node representations, it computes a pairwise similarity between every SP node and every CP node, and finds maximal similarity pairs to connect scene predicates to their corresponding classes, via predicate bridges. It also does the same for entity nodes to potentially refine their bridges too. Given the new bridges, it propagates messages again, and repeats this process for a predefined number of steps. The final state of the bridge determines which class each node belongs to, resulting in the output scene graph.

4.1 Graph Initialization

The object detector outputs a set of n detected objects, each with a bounding box b_j , a label distribution p_j and an RoI-aligned [36] feature vector \mathbf{v}_j . Then we allocate a *scene entity node* (SE) for each object, and a *scene predicate node* (SP) for each pair of objects, representing the potential predicate with the two entities as its subject and object. Each entity is initialized using its RoI features \mathbf{v}_j , and each predicate is initialized using the RoI features \mathbf{u}_j of a bounding box enclosing the union of its subject and object. Formally, we can write, *i.e.*,

$$\mathbf{x}_j^{\text{SE}} = \phi_{\text{init}}^{\text{SE}}(\mathbf{v}_j), \quad \text{and} \quad \mathbf{x}_j^{\text{SP}} = \phi_{\text{init}}^{\text{SP}}(\mathbf{u}_j), \quad (6)$$

where $\phi_{\text{init}}^{\text{SE}}$ and $\phi_{\text{init}}^{\text{SP}}$ are two fully connected networks that are branched from the backbone after ROI-align. To form a scene graph proposal, we connect each predicate node to its subject and object via labeled edges. Specifically, we define the following 4 edge types: for a triplet $s - p - o$, we connect p to s using a `hasSubject` edge, p to o using a `hasObject` edge, s to p using a `subjectOf` edge, and o to p using an `objectOf` edge. The reason we have two directions as separate types is that in the message passing phase, the way we use predicate information to update entities should be different from the way we use entities to update predicates.

On the other hand, we initialize the commonsense graph with *commonsense entity nodes* (CE) and *commonsense predicate nodes* (CP) using a linear projection of their word embeddings:

$$\mathbf{x}_i^{\text{CE}} = \phi_{\text{init}}^{\text{CE}}(\mathbf{e}_i^n), \quad \text{and} \quad \mathbf{x}_i^{\text{CP}} = \phi_{\text{init}}^{\text{CP}}(\mathbf{e}_i^p). \quad (7)$$

The commonsense graph also has various types of edges, such as `UsedFor` and `PartOf`, as detailed in Sect. 5.2. Our method is independent of the types of commonsense edges, and can utilize any provided graph from any source.

So far, we have two isolated graphs, scene and commonsense. An SE node representing a detected `Person` intuitively refers to the `Person` concept in the ontology, and hence the `Person` node in the commonsense graph. Therefore, we connect each SE node to the CE node that corresponds the semantic label predicted by Faster R-CNN, via a `classifiedTo` edge type. Instead of a hard classification, we connect each entity to top K_{bridge} classes using p_j (class distribution predicted by Faster R-CNN) as weights. We also create a reverse connection from each CE node to corresponding SE nodes, using an `hasInstance` edge, but with the same weights p_j . As mentioned earlier, this is to make sure information flows from commonsense to scene as well as scene to commonsense, but not in the same way. We similarly define two other edge types, `classifiedTo` and `hasInstance` for predicates, which are initially an empty set, and will be updated to bridge SP nodes to CP nodes as we explain in the following. These 4 edge types can be seen as flexible *bridges* that connect the two fixed graphs, which are considered latent variables to be determined by the model.

This forms a heterogeneous graph with four types of nodes (SE, SP, CE, and CP) and various types of edges: scene graph edges \mathcal{E}_S such as `subjectOf`, commonsense edges \mathcal{E}_C such as `usedFor`, and bridge edges \mathcal{E}_B such as `classifiedTo`.

Next, we explain how our proposed method updates node representations and bridge edges, while keeps commonsense and scene edges constant.

4.2 Successive Message Passing and Bridging

Given a heterogeneous graph as described above, we employ a variant of GGNN [25] to propagate information among nodes. First, each node representation is fed into a fully connected network to compute *outgoing* messages, that is

$$\mathbf{m}_i^{\Delta \rightarrow} = \phi_{\text{send}}^{\Delta}(\mathbf{x}_i^{\Delta}), \quad (8)$$

for each i and node type Δ , where ϕ_{send} is a trainable *send head* which has shared weights across nodes of each type. After computing outgoing messages, we send them through all outgoing edges, multiplying by the edge weight. Then for each node, we aggregate incoming messages, by first adding across edges of the same type, and then concatenating across edge types. We compute the *incoming* message for each node by applying another fully connected network on the aggregated messages:

$$\mathbf{m}_j^{\Delta \leftarrow} = \phi_{\text{receive}}^{\Delta} \left(\bigcup_{\Delta'} \bigcup_{\mathcal{E}_k \in \mathcal{E}^{\Delta' \rightarrow \Delta}} \sum_{(i,j,a_{ij}^k) \in \mathcal{E}_k} a_{ij}^k \mathbf{m}_i^{\Delta' \rightarrow} \right), \quad (9)$$

where ϕ_{receive} is a trainable *receive head* and \cup denotes concatenation. Note that the first concatenation is over all 4 node types, the second concatenation is over all edge types from Δ' to Δ , and the sum is over all edges of that type, where i and j are the head and tail nodes, and a_{ij}^k is the edge weight. Given the incoming message for each node, we update the representation of the node using a Gated Recurrent Unit (GRU) update rule, following [4]:

$$\begin{aligned} \mathbf{z}_j^{\Delta} &= \sigma(W_z^{\Delta} \mathbf{m}_j^{\Delta \leftarrow} + U_z^{\Delta} \mathbf{x}_j^{\Delta}), \\ \mathbf{r}_j^{\Delta} &= \sigma(W_r^{\Delta} \mathbf{m}_j^{\Delta \leftarrow} + U_r^{\Delta} \mathbf{x}_j^{\Delta}), \\ \mathbf{h}_j^{\Delta} &= \tanh(W_h^{\Delta} \mathbf{m}_j^{\Delta \leftarrow} + U_h^{\Delta} (\mathbf{r}_j^{\Delta} \odot \mathbf{x}_j^{\Delta})), \\ \mathbf{x}_j^{\Delta} &\Leftarrow (1 - \mathbf{z}_j^{\Delta}) \odot \mathbf{x}_j^{\Delta} + \mathbf{z}_j^{\Delta} \odot \mathbf{h}_j^{\Delta}, \end{aligned} \quad (10)$$

where σ is the sigmoid function, and W^{Δ} and U^{Δ} are trainable matrices that are shared across nodes of the same type, but distinct for each node type Δ . This update rule can be seen as an extension of GGNN [25] to heterogeneous graphs, with a more complex message aggregation strategy. Note that \Leftarrow means we update the node representation. Mathematically, this means $\mathbf{x}_{j(t+1)}^{\Delta} = U(\mathbf{x}_{j(t)}^{\Delta})$, where U is the aforementioned update rule and (t) denotes iteration number. For simplicity, we drop this subscript throughout this paper.

So far, we have explained how to update node representations using graph edges. Now using the new node representations, we should update the bridge

edges \mathcal{E}_B that connect scene nodes to commonsense nodes. To this end, we compute a pairwise similarity from each SE to all CE nodes, and from each SP to all CP nodes.

$$\mathbf{a}_{ij}^{\text{EB}} = \frac{\exp\langle \mathbf{x}_i^{\text{SE}}, \mathbf{x}_j^{\text{CE}} \rangle_{\text{EB}}}{\sum_{j'} \exp\langle \mathbf{x}_i^{\text{SE}}, \mathbf{x}_{j'}^{\text{CE}} \rangle_{\text{EB}}}, \quad \text{where} \quad \langle \mathbf{x}, \mathbf{y} \rangle_{\text{EB}} = \phi_{\text{att}}^{\text{SE}}(\mathbf{x})^T \phi_{\text{att}}^{\text{CE}}(\mathbf{y}), \quad (11)$$

and similarly for predicates,

$$\mathbf{a}_{ij}^{\text{PB}} = \frac{\exp\langle \mathbf{x}_i^{\text{SP}}, \mathbf{x}_j^{\text{CP}} \rangle_{\text{PB}}}{\sum_{j'} \exp\langle \mathbf{x}_i^{\text{SP}}, \mathbf{x}_{j'}^{\text{CP}} \rangle_{\text{PB}}}, \quad \text{where} \quad \langle \mathbf{x}, \mathbf{y} \rangle_{\text{PB}} = \phi_{\text{att}}^{\text{SP}}(\mathbf{x})^T \phi_{\text{att}}^{\text{CP}}(\mathbf{y}). \quad (12)$$

Here $\phi_{\text{att}}^{\Delta}$ is a fully connected network that resembles *attention head* in transformers. Note that since $\phi_{\text{att}}^{\Delta}$ is not shared across node types, our similarity metric is asymmetric. We use each $\mathbf{a}_{ij}^{\text{EB}}$ to set the edge weight of the `classifiedTo` edge from \mathbf{x}_i^{SE} to \mathbf{x}_j^{CE} , as well as the `hasInstance` edge from \mathbf{x}_j^{CE} to \mathbf{x}_i^{SE} . Similarly we use each $\mathbf{a}_{ij}^{\text{PB}}$ to set the weight of edges between \mathbf{x}_i^{SP} and \mathbf{x}_j^{CP} . In preliminary experiments we realised that such fully connected bridges hurt performance in large graphs. Hence, we only keep the top K_{bridge} values of $\mathbf{a}_{ij}^{\text{EB}}$ for each i , and set the rest to zero. We do the same thing for predicates, keeping the top K_{bridge} values of $\mathbf{a}_{ij}^{\text{PB}}$ for each i . Given the updated bridges, we propagate messages again to update node representations, and iterate for a fixed number of steps, T . The final values of $\mathbf{a}_{ij}^{\text{EB}}$ and $\mathbf{a}_{ij}^{\text{PB}}$ are the outputs of our model, which can be used to classify each entity and predicate in the scene graph.

4.3 Training

We closely follow [2] which itself follows [52] for training procedure. Specifically, given the output and ground truth graphs, we align output entities and predicates to ground truth counterparts. To align entities we use IoU and predicates will be aligned naturally since they correspond to aligned pairs of entities. Then we use the output probability scores of each node to define a cross-entropy loss. The sum of all node-level loss values will be the objective function to be minimized using Adam [18].

Due to the highly imbalanced predicate statistics in Visual Genome, we observed that best-performing models usually concentrate their performance merely on the most frequent classes such as `on` and `wearing`. To alleviate this, we modify the basic cross-entropy objective that is commonly used by assigning an importance weight to each class. We follow the recently proposed class-balanced loss [5] where the weight of each class is inversely proportional to its frequency. More specifically, we use the following loss function for each predicate node:

$$\mathcal{L}_i^P = -\frac{1 - \beta}{1 - \beta^{n_j}} \log \mathbf{a}_{ij}^{\text{PB}}, \quad (13)$$

where j is the class index of the ground truth predicate aligned with i , n_j is the frequency of class j in training data, and β is a hyperparameter. Note that

$\beta = 0$ leads to a regular cross-entropy loss, and the more it approaches 1, the more strictly it suppresses frequent classes. To be fair in comparison with other methods, we include a variant of our method without reweighting, which still outperforms all other methods.

5 Experiments

Following the literature, we use the large-scale Visual Genome benchmark [20] to evaluate our method. We first show our GB-NET outperforms the state of the art, by extensively evaluating it on 24 performance metrics. Then we present an ablation study to illustrate how each innovation contributes to the performance. In the Supplementary Material, we also provide a per-class performance breakdown to show the consistency and robustness of our performance across frequent and rare classes. That is accompanied by a computational speed analysis, and several qualitative examples of our generated graphs compared to the state of the art, side by side.

5.1 Task Description

Visual Genome [20] consists of 108,077 images with annotated objects (entities) and pairwise relationships (predicates), which is then post-processed by [44] to create scene graphs. They use the most frequent 150 entity classes and 50 predicate classes to filter the annotations. Figure 1 shows an example of their post-processed scene graphs which we use as ground truth. We closely follow their evaluation settings such as train and test splits.

The task of scene graph generation, as described in Sect. 4, is equivalent to the SGG_{EN} scenario proposed by [44] and followed ever since. Given an image, the task of SGG_{EN} is to jointly infer entities and predicates from scratch. Since this task is limited by the quality of the object proposals, [44] also introduced two other tasks that more clearly evaluate entity and predicate recognition. In SGCLS, we take localization (here region proposal network) out of the picture, by providing the model with ground truth bounding boxes during test, simulating a *perfect* proposal model. In PREDCLS, we take object detection for granted, and provide the model with not only ground truth bounding boxes, but also their true entity class. In each task, the main evaluation metric is average per-image recall of the top K subject-predicate-object triplets. The confidence of a triplet that is used for ranking is computed by multiplying the classification confidence of all three elements. Given the ground truth scene graph, each predicate forms a triplet, which we match against the top K triplets in the output scene graph. A triplet is matched if all three elements are classified correctly, and the bounding boxes of subject and object match with an IoU of at least 0.5. Besides the choice of K , there are two other choices to be made: (1) Whether or not to enforce the so-called *Graph Constraint* (GC), which limits the top K triplets to only one predicate for each ordered entity pair, and (2) Whether to compute the recall for each predicate class separately and take the mean (mR), or compute a single

recall for all triplets (R) [2]. We comprehensively report both mean and overall recall, both with and without GC, and conventionally use both 50 and 100 for K , resulting in 8 metrics for each task, 24 in total.

5.2 Implementation Details

We use three-layer fully connected networks with ReLU activation for all trainable networks ϕ_{init} , ϕ_{send} , ϕ_{receive} and ϕ_{att} . We set the dimension of node representations to 1024, and perform 3 message passing steps, except in ablation experiments where we try 1, 2 and 3. We tried various values for β . Generally the higher it is, mean recall improves and recall falls. We found 0.999 is a good trade-off, and chose $K_{\text{bridge}} = 5$ empirically. All hyperparameters are tuned using a validation set randomly selected from training data. We borrow the Faster R-CNN trained by [52] and shared among all our baselines, which has a VGG-16 backbone and predicts 128 proposals.

In our commonsense graph, the nodes are the 151 entity classes and 51 predicate classes that are fixed by [44], including background. We use the GloVE [33] embedding of category titles to initialize their node representation (via ϕ_{init}), and fix GloVE during training. We compile our commonsense edges from three sources, WordNet [30], ConceptNet [27], and Visual Genome. To summarize, there are three groups of edge types in our commonsense graph. We have **SimilarTo** from WordNet hierarchy, we have **PartOf**, **RelatedTo**, **IsA**, **MannerOf**, and **UsedFor** from ConceptNet, and finally from VG training data we have conditional probabilities of subject given predicate, predicate given subject, subject given object, *etc.* We explain the details in the supplementary material. The process of compiling and pruning the knowledge graph is semi-automatic and takes less than a day from a single person. We make it publicly available as a part of our code. We have also tried using each individual source (e.g. only ConceptNet) independently, which requires less effort, and does not significantly impact the performance. There are also recent approaches to automate the process of commonsense knowledge graph construction [1, 13], which can be utilized to further reduce the manual labor.

5.3 Main Results

Table 1 summarizes our results in comparison to the state of the art. IMP+ refers to the re-implementation of [44] by [52] using their new Faster R-CNN backbone. That method does not use any external knowledge and only uses message passing among the entities and predicates and then classifies each. Hence, it can be seen as a strong, but knowledge-free baseline. FREQ is a simple baseline proposed by [52], which predicts the most frequent predicate for any given pair of entity classes, solely based on statistics from the training data. FREQ surprisingly outperforms IMP+, confirming the efficacy of commonsense in SGG.

SMN [52] applies bi-directional LSTMs on top of the entity features, then classifies each entity and each pair. They bias their classifier logits using statistics from FREQ, which improves their total recall significantly, at the expense

Table 1. Evaluation in terms of mean and overall triplet recall, at top 50 and top 100, with and without Graph Constraint (GC), for the three tasks of SGG_{EN}, SGCLS and PREDCLS. Numbers are in percentage. All baseline numbers were borrowed from [2]. Top two methods for each metric is shown in **bold** and *italic* respectively.

Task	Metric	GC	Method					
			IMP+	FREQ	SMN	KERN	GB-NET	GB-NET- β
SGG _{EN}	mR@50	Y	3.8	4.3	5.3	<i>6.4</i>	6.1	7.1
		N	5.4	5.9	9.3	<i>11.7</i>	9.8	11.7
	mR@100	Y	4.8	5.6	6.1	7.3	<i>7.3</i>	8.5
		N	8.0	8.9	12.9	<i>16.0</i>	14.0	16.6
	R@50	Y	20.7	23.5	27.2	<i>27.1</i>	26.4	26.3
		N	22.0	25.3	<i>30.5</i>	30.9	29.4	29.3
	R@100	Y	24.5	27.6	30.3	29.8	<i>30.0</i>	29.9
		N	27.4	30.9	<i>35.8</i>	35.8	35.1	35.0
SGCLS	mR@50	Y	5.8	6.8	7.1	9.4	<i>9.6</i>	12.7
		N	12.1	13.5	15.4	19.8	<i>21.4</i>	25.6
	mR@100	Y	6.0	7.8	7.6	10.0	<i>10.2</i>	13.4
		N	16.9	19.6	20.6	26.2	<i>29.1</i>	32.1
	R@50	Y	34.6	32.4	35.8	36.7	38.0	<i>37.3</i>
		N	43.4	40.5	44.5	45.9	47.7	<i>46.9</i>
	R@100	Y	35.4	34.0	36.5	37.4	38.8	<i>38.0</i>
		N	47.2	43.7	47.7	49.0	51.1	<i>50.3</i>
PREDCLS	mR@50	Y	9.8	13.3	13.3	17.7	<i>19.3</i>	22.1
		N	20.3	24.8	27.5	36.3	<i>41.1</i>	44.5
	mR@100	Y	10.5	15.8	14.4	19.2	<i>20.9</i>	24.0
		N	28.9	37.3	37.9	49.0	<i>55.4</i>	58.7
	R@50	Y	59.3	59.9	65.2	65.8	66.6	<i>66.6</i>
		N	75.2	71.3	81.1	81.9	83.6	<i>83.5</i>
	R@100	Y	61.3	64.1	67.1	67.6	68.2	<i>68.2</i>
		N	83.6	81.2	88.3	88.9	90.5	<i>90.3</i>

of higher bias against less frequent classes, as revealed by [2]. More recently, KERN [2] encodes VG statistics into the edge weights of the graph, which is then incorporated by propagating messages. Since it encodes statistics more implicitly, KERN is less biased compared to SMN, which improves mR. Our method improves both R and mR significantly, and our class-balanced model, GB-NET- β , further enhances mR (+2.7% in average) without hurting R by much (-0.2%).

We observed that the state of the art performance has been saturated in the SGG_{EN} setting, especially for overall recall. This is partly because object detection performance is a bottleneck that limits the performance. It is worth noting

Table 2. Ablation study on Visual Genome. All numbers are in percentage, and graph constraint is enforced.

Method	SGGEN				PREDCLS			
	mR@50	mR@100	R@50	R@100	mR@50	mR@100	R@50	R@100
No Knowledge	5.5	6.6	25.3	28.8	15.4	16.8	62.5	64.5
$T = 1$	5.6	6.7	24.9	28.5	15.6	17.1	62.1	64.2
$T = 2$	5.7	6.9	26.1	29.7	18.2	19.7	66.7	68.4
GB-NET	6.1	7.3	26.4	30.0	18.2	19.7	67.0	68.6

that mean recall is a more important metric than overall recall, since most SGG methods tend to score a high overall recall by investing on few most frequent classes, and ignoring the rest [2]. As shown in Table 1, our method achieves significant improvements in mean recall. We provide in-depth performance analysis by comparing our recall per predicate class with that of the state of the art, as well as qualitative analysis in the Supplementary Material.

There are other recent SGG methods that are not used for comparison here, because their evaluation settings are not identical to ours, and their code is not publicly available to the best of our knowledge [10, 34]. For instance, [34] reports only 8 out of our 24 evaluation metrics, and although our method is superior in 6 metrics out of those 8, that is not sufficient to fairly compare the two methods.

5.4 Ablation Study

To further explain our performance improvement, Table 2 compares our full method with its weaker variants. Specifically, to investigate the effectiveness of commonsense knowledge, we remove the commonsense graph and instead classify each node in our graph using a 2-layer fully connected classifier after message passing. This negatively impacts performance in all metrics, proving our method is able to exploit commonsense knowledge through the proposed bridging technique. Moreover, to highlight the importance of our proposed message passing and bridge refinement process, we repeated the experiments with fewer steps. We observe the performance drops significantly with fewer steps, proving the effectiveness of our model, but it saturates as we go beyond 3 steps.

6 Conclusion

We proposed a new method for Scene Graph Generation that incorporates external commonsense knowledge in a novel, graphical neural framework. We unified the formulation of scene graph and commonsense graph as two types of knowledge graph, which are fused into a single graph through a dynamic message passing and bridging algorithm. Our method iteratively propagates messages to update nodes, then compares nodes to update bridge edges, and repeats until the two graphs are carefully connected. Through extensive experiments, we showed our method outperforms the state of the art in various metrics.

Acknowledgement. This work was supported in part by Contract N6600119C4032 (NIWC and DARPA). The views expressed are those of the authors and do not reflect the official policy of the Department of Defense or the U.S. Government.

References

1. Bosselut, A., Rashkin, H., Sap, M., Malaviya, C., Celikyilmaz, A., Choi, Y.: COMET: commonsense transformers for automatic knowledge graph construction. arXiv preprint [arXiv:1906.05317](https://arxiv.org/abs/1906.05317) (2019)
2. 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, pp. 6163–6171 (2019)
3. Chen, X., Li, L.J., Fei-Fei, L., Gupta, A.: Iterative visual reasoning beyond convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7239–7248 (2018)
4. Cho, K., et al.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. arXiv preprint [arXiv:1406.1078](https://arxiv.org/abs/1406.1078) (2014)
5. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9268–9277 (2019)
6. Dietz, L., Kotov, A., Meij, E.: Utilizing knowledge graphs for text-centric information retrieval. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, pp. 1387–1390. ACM (2018)
7. Fader, A., Zettlemoyer, L., Etzioni, O.: Open question answering over curated and extracted knowledge bases. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1156–1165. ACM (2014)
8. Flanigan, J., Thomson, S., Carbonell, J., Dyer, C., Smith, N.A.: A discriminative graph-based parser for the abstract meaning representation. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1426–1436 (2014)
9. Gardner, M., Dasigi, P., Iyer, S., Suhr, A., Zettlemoyer, L.: Neural semantic parsing. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts, pp. 17–18 (2018)
10. 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, pp. 1969–1978 (2019)
11. 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, pp. 7211–7221 (2018)
12. Hudson, D.A., Manning, C.D.: Learning by abstraction: the neural state machine. arXiv preprint [arXiv:1907.03950](https://arxiv.org/abs/1907.03950) (2019)
13. Ilievski, F., Szekely, P., Cheng, J., Zhang, F., Qasemi, E.: Consolidating commonsense knowledge. arXiv preprint [arXiv:2006.06114](https://arxiv.org/abs/2006.06114) (2020)
14. Johnson, J., Gupta, A., Fei-Fei, L.: Image generation from scene graphs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1219–1228 (2018)
15. Johnson, J., et al.: Image retrieval using scene graphs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3668–3678 (2015)

16. Kato, K., Li, Y., Gupta, A.: Compositional learning for human object interaction. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018*. LNCS, vol. 11218, pp. 247–264. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01264-9_15
17. Khashabi, D., Khot, T., Sabharwal, A., Roth, D.: Question answering as global reasoning over semantic abstractions. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
18. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
19. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
20. Krishna, R., et al.: Visual genome: connecting language and vision using crowd-sourced dense image annotations. *Int. J. Comput. Vis.* **123**(1), 32–73 (2017)
21. Lee, C.W., Fang, W., Yeh, C.K., Frank Wang, Y.C.: Multi-label zero-shot learning with structured knowledge graphs. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1576–1585 (2018)
22. Li, M., et al.: Multilingual entity, relation, event and human value extraction. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pp. 110–115 (2019)
23. Li, Y., Ouyang, W., Zhou, B., Shi, J., Zhang, C., Wang, X.: Factorizable net: an efficient subgraph-based framework for scene graph generation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *ECCV 2018*. LNCS, vol. 11205, pp. 346–363. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01246-5_21
24. Li, Y., Ouyang, W., Zhou, B., Wang, K., Wang, X.: Scene graph generation from objects, phrases and region captions. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1261–1270 (2017)
25. Li, Y., Tarlow, D., Brockschmidt, M., Zemel, R.: Gated graph sequence neural networks. arXiv preprint [arXiv:1511.05493](https://arxiv.org/abs/1511.05493) (2015)
26. Liang, K., Guo, Y., Chang, H., Chen, X.: Visual relationship detection with deep structural ranking. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
27. Liu, H., Singh, P.: ConceptNet—a practical commonsense reasoning tool-kit. *BT Technol. J.* **22**(4), 211–226 (2004)
28. Liu, Y., Wang, R., Shan, S., Chen, X.: Structure inference net: object detection using scene-level context and instance-level relationships. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6985–6994 (2018)
29. Marino, K., Salakhutdinov, R., Gupta, A.: The more you know: using knowledge graphs for image classification. arXiv preprint [arXiv:1612.04844](https://arxiv.org/abs/1612.04844) (2016)
30. Miller, G.A.: WordNet: a lexical database for English. *Commun. ACM* **38**(11), 39–41 (1995)
31. Newell, A., Deng, J.: Pixels to graphs by associative embedding. In: *Advances in Neural Information Processing Systems*, pp. 2171–2180 (2017)
32. Pei, M., Jia, Y., Zhu, S.C.: Parsing video events with goal inference and intent prediction. In: *2011 International Conference on Computer Vision*, pp. 487–494. IEEE (2011)
33. Pennington, J., Socher, R., Manning, C.: Glove: global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543 (2014)

34. 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, pp. 3957–3966 (2019)
35. Qi, S., Wang, W., Jia, B., Shen, J., Zhu, S.-C.: Learning human-object interactions by graph parsing neural networks. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11213, pp. 407–423. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01240-3_25
36. 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, pp. 91–99 (2015)
37. Schuster, S., Krishna, R., Chang, A., Fei-Fei, L., Manning, C.D.: Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In: Proceedings of the Fourth Workshop on Vision and Language, pp. 70–80 (2015)
38. Shi, J., Zhang, H., Li, J.: Explainable and explicit visual reasoning over scene graphs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8376–8384 (2019)
39. Teney, D., Liu, L., van den Hengel, A.: Graph-structured representations for visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–9 (2017)
40. Tu, K., Meng, M., Lee, M.W., Choe, T.E., Zhu, S.C.: Joint video and text parsing for understanding events and answering queries. *IEEE MultiMedia* **21**(2), 42–70 (2014)
41. Wadden, D., Wennberg, U., Luan, Y., Hajishirzi, H.: Entity, relation, and event extraction with contextualized span representations. arXiv preprint [arXiv:1909.03546](https://arxiv.org/abs/1909.03546) (2019)
42. Wang, X., Ye, Y., Gupta, A.: Zero-shot recognition via semantic embeddings and knowledge graphs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6857–6866 (2018)
43. Woo, S., Kim, D., Cho, D., Kweon, I.S.: Linknet: relational embedding for scene graph. In: Advances in Neural Information Processing Systems, pp. 560–570 (2018)
44. 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, pp. 5410–5419 (2017)
45. Xu, H., Jiang, C., Liang, X., Li, Z.: Spatial-aware graph relation network for large-scale object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9298–9307 (2019)
46. Xu, H., Jiang, C., Liang, X., Lin, L., Li, Z.: Reasoning-RCNN: unifying adaptive global reasoning into large-scale object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6419–6428 (2019)
47. Yang, J., Lu, J., Lee, S., Batra, D., Parikh, D.: Graph R-CNN for scene graph generation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11205, pp. 690–706. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01246-5_41
48. 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, pp. 10685–10694 (2019)
49. Yao, T., Pan, Y., Li, Y., Mei, T.: Exploring visual relationship for image captioning. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) Computer Vision – ECCV 2018. LNCS, vol. 11218, pp. 711–727. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01264-9_42

50. Yu, J., Lu, Y., Qin, Z., Zhang, W., Liu, Y., Tan, J., Guo, L.: Modeling text with graph convolutional network for cross-modal information retrieval. In: Hong, R., Cheng, W.-H., Yamasaki, T., Wang, M., Ngo, C.-W. (eds.) PCM 2018. LNCS, vol. 11164, pp. 223–234. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00776-8_21
51. Zareian, A., Karaman, S., Chang, S.F.: Weakly supervised visual semantic parsing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3736–3745 (2020)
52. 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, pp. 5831–5840 (2018)
53. 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, pp. 5128–5137 (2019)
54. Zhang, C., Chao, W.L., Xuan, D.: An empirical study on leveraging scene graphs for visual question answering. arXiv preprint [arXiv:1907.12133](https://arxiv.org/abs/1907.12133) (2019)
55. Zhao, Y., Zhu, S.C.: Image parsing with stochastic scene grammar. In: Advances in Neural Information Processing Systems, pp. 73–81 (2011)