

Learning Visual Commonsense for Robust Scene Graph Generation

Alireza Zareian*, Zhecan Wang*, Haoxuan You*, and Shih-Fu Chang

Columbia University, New York NY 10027, USA
{az2407, zw2627, hy2612, sc250}@columbia.edu

Abstract. Scene graph generation models understand the scene through object and predicate recognition, but are prone to mistakes due to the challenges of perception in the wild. Perception errors often lead to non-sensical compositions in the output scene graph, which do not follow real-world rules and patterns, and can be corrected using commonsense knowledge. We propose the first method to acquire visual commonsense such as affordance and intuitive physics automatically from data, and use that to improve the robustness of scene understanding. To this end, we extend Transformer models to incorporate the structure of scene graphs, and train our Global-Local Attention Transformer on a scene graph corpus. Once trained, our model can be applied on any scene graph generation model and correct its obvious mistakes, resulting in more semantically plausible scene graphs. Through extensive experiments, we show our model learns commonsense better than any alternative, and improves the accuracy of state-of-the-art scene graph generation methods.

1 Introduction

In recent computer vision literature, there is a growing interest in incorporating commonsense reasoning and background knowledge into the process of visual recognition and scene understanding [9, 31, 13, 8, 33]. In Scene Graph Generation (SGG), for instance, external knowledge bases [7] and dataset statistics [2, 34] have been utilized to improve the accuracy of entity (object) and predicate (relation) recognition. The effect of these techniques is usually to correct obvious perception errors, and replace with more plausible alternatives. For instance, Figure 1 (top) shows an SGG model mistakenly classifies a **bird** as a **bear**, possibly due to the dim lighting and small object size. However, a commonsense model can correctly predict **bird**, because **bear on branch** is a less common situation, less aligned with intuitive physics, or contrary to animal behavior.

Nevertheless, existing methods to incorporate commonsense into the process of visual recognition have two major limitations. Firstly, they rely on an external source of commonsense, such as crowd-sourced or automatically mined commonsense rules, which tend to be incomplete and inaccurate [7], or statistics directly gathered from training data, which are limited to simple heuristics such

* Equal contribution.




Fig. 1. Overview of the proposed method: We propose a commonsense model that takes a scene graph generated by a perception model and refines that to make it more plausible. Then a fusion module compares the perception and commonsense outputs and generates a final graph, incorporating both signals.

as co-occurrence frequency [2]. In this paper, we propose the first method to learn graphical commonsense automatically from a scene graph corpus, which does not require external knowledge, and **acquires** commonsense by learning complex, structured patterns beyond simple heuristics.

Secondly, most existing methods are strongly vulnerable to data bias as they integrate data-driven commonsense knowledge into data-driven neural networks. For instance, the commonsense model in Figure 1 mistakes the **elephant** for a **person**, in order to avoid the bizarre triplet **elephant drawing picture**, while the **elephant** is quite clear visually, and the perception model already recognizes it correctly. None of the existing efforts to equip scene understanding with commonsense have studied the fundamental question of whether to trust perception or commonsense, *i.e.*, what you see versus what you expect. In this paper, we propose a way to disentangle perception and commonsense into two separately trained models, and introduce a method to exploit the disagreement between those two models to achieve the best of both worlds.

To this end, we first propose a mathematical formalization of visual commonsense, as a problem of auto-encoding perturbed scene graphs. Based on the new formalism, we propose a novel method to learn visual commonsense from annotated scene graphs. We extend recently successful transformers [23] by adding local attention heads to enable them to encode the structure of a scene graph, and we train them on a corpus of annotated scene graphs to predict missing elements of a scene via a masking framework similar to BERT [5]. As illustrated in Figure 2, our commonsense model learns to use its experience to imagine which entity or predicate could replace the mask, considering the structure and context of a given scene graph. Once trained, it can be stacked on top of any perception (*i.e.*, SGG) model to correct nonsensical mistakes in the generated scene graphs.

The output of the perception and commonsense models can be seen as two generated scene graphs with potential disagreements. We devise a fusion module that takes those two graphs, along with their classification confidence values, and predicts a final scene graph that reflects both perception and commonsense

knowledge. The degree to which our fusion module trusts each input varies for each image, and is determined based on the estimated confidence of each model. This way, if the perception model is uncertain about the `bird` due to darkness, the fusion module relies on the commonsense more, and if perception is confident about the `elephant` due to its clarity, the fusion module *trusts its eyes*.

We conduct extensive experiments on the Visual Genome datasets [12], showing (1) The proposed GLAT model outperforms existing transformers and graph-based models in the task of commonsense acquisition; (2) Our model learns various types of commonsense that are absent in SGG models, such as object affordance and intuitive physics; (3) The proposed model is robust to dataset bias, and shows commonsensical behavior even in rare and zero-shot scenarios; (4) The proposed GLAT and Fusion mechanism can be applied on any SGG method to correct their mistakes and improve their accuracy. The main contributions of this paper are the following:

- We propose the first method for learning structured visual commonsense, Global-Local Attention Transformer (GLAT), which does not require any external knowledge, and outperforms conventional transformers and graph-based networks.
- We propose a cascaded fusion architecture for Scene Graph Generation, which disentangles commonsense reasoning from visual perception, and integrates them in a way that is robust to the failure of each component.
- We report experiments that showcase our model’s unique ability of learning commonsense without picking up dataset bias, and its utility in downstream scene understanding.

2 Related Work

2.1 Commonsense in computer vision

Incorporating commonsense knowledge has been explored in various computer vision tasks such as object recognition [3, 28, 14], object detection [13], semantic segmentation [19], action recognition [9], visual relation detection [31], scene graph generation [2, 34, 7], and visual question answering [18, 22]. There are two aspects to study about these methods: where their commonsense comes from, and how they use it.

Most methods either adopt an external curated knowledge base such as ConceptNet [21, 7, 28, 14, 19, 18], or acquire commonsense automatically by collecting statistics over an often annotated corpus [3, 2, 34, 13, 22, 31]. Nevertheless, the former group are limited to incomplete external knowledge, and the latter are based on ad-hoc, hard-coded heuristics such as the co-occurrence frequency of categories. Our method is the first to formulate visual commonsense as a machine learning task, and train a graph-based neural network to solve it. There are a third group of works that focus on a particular type of commonsense by designing a specialized model, such as intuitive physics [6], or object affordance [4]. We put forth a more general framework that includes but is not limited to physics

and affordance, by exploiting scene graphs as a versatile semantic representation. The most similar to our work is [26], which only models object co-occurrence patterns, while we also incorporate object relationships and scene graph structure.

When it comes to utilizing commonsense, existing methods integrate it within the inference pipeline, either by retrieving a set of relevant facts from a knowledge base and feeding as additional features to the model [18, 7, 22], or by employing a graph-based message propagation process to embed the structure of the knowledge graph within the intermediate representations of the model [3, 9, 2, 28, 14]. Some other methods distill the knowledge during training through auxiliary objectives, making the inference simple and free of external knowledge [19, 31]. Nevertheless, in all those approaches, commonsense is seamlessly infused into the model and cannot be disentangled. This makes it hard to study and evaluate commonsense and perception separately, or control their influence. Few methods have modeled commonsense as a standalone module which is late-fused into the prediction of the perception model [13, 34]. Yet, we are the first to devise separate perception and commonsense models, and adaptively weigh their importance based on their confidence, before fusing their predictions.

2.2 Commonsense in scene graph generation

Zellers *et al.* [34] were the first to explicitly incorporate commonsense into the process of scene graph generation. They biased predicate classification logits using a pre-computed frequency prior that is a static distribution, given each entity class pair. Although this significantly improved their overall accuracy, the improvement is mainly due to the fact that they favor frequent triplets over others, which is statistically rewarding. Even if their model classifies the relation between a **person** and a **hat** as **holding**, their frequency bias would most likely change that to **wearing**, which is more frequent.

More recently, Chen *et al.* [2] employed a less explicit way to incorporate the frequency prior within the process of entity and predicate classification. They embed the frequencies into the edge weights of their inference graph, and utilize those weights within their message propagation process. This improves the results especially on less frequent predicates, since it less strictly enforces the statistics on the final decision. However, this way commonsense is integrated implicitly into the SGG model and cannot be probed or studied in isolation. We remove the adverse effect of statistical bias while keeping the commonsense model disentangled from perception.

Gu *et al.* [7] exploits ConceptNet [21] rather than dataset statistics, which is a large-scale knowledge graph comprising relational facts about concepts, *e.g.* **dog is-a animal** or **fork is-used-for eating**. Given each detected object, they retrieve ConceptNet facts involving that object class, and employ a recurrent neural net and an attention mechanism to encode those facts into the object features, before classifying objects and predicates. Nevertheless, ConceptNet is not exhaustive, since it is extremely hard to compile all commonsense facts. Our

method does not depend on a limited source of external knowledge, and acquires commonsense automatically, via a generalizable neural network.

2.3 Transformers and graph-based neural networks

Transformers were originally proposed to replace recurrent neural networks for machine translation, by stacking several layers of multi-head attention [23]. Ever since, transformers have been successful in various vision and language tasks [5, 27, 16]. Particularly, BERT [5] randomly replaces some words from a given sentence with a special MASK token and tries to reconstruct those words. Through this self-supervised game, BERT acquires natural language, and can transfer its language knowledge to perform well in other NLP tasks. We use a similar self-supervised strategy to learn to complete missing pieces of a scene graph. Rather than language, our model acquires the ability to imagine a scene in a structured, semantic way, which is a hallmark of human commonsense.

Transformers treat their input as a set of tokens, and discard any form of structure among them. To preserve the order of tokens in a sentence, BERT augments the initial embedding of each token with a position embedding before feeding into transformers. Scene graphs, on the other hand, have a more complex structure that cannot be embedded in such a trivial way. Recently, Graph-based Neural Networks (GNN) have been successful to encode graph structures into node representations, by applying several layers of neighborhood aggregation. More specifically, each layer of a GNN represents each node by a trainable function that takes the node as well as its neighbors as input. Graph convolutional nets [11], gated graph neural nets [15], and graph attention nets [24] all implement this idea with different computational models for neighborhood aggregation. GNNs have been widely utilized for scene graph generation by incorporating context [29, 30, 32], but we are the first to exploit GNNs to learn visual commonsense.

We adopt graph attention nets due to their similarity to transformers in using attention. The main difference of graph attention nets to transformers is that instead of representing each node by an attention over all other nodes, they only compute an attention over immediate neighbors. Inspired by that, we use a BERT-like transformer network, but replace half of its attention heads by local attention, simply by enforcing the attention between non-neighbor nodes to zero. Through ablation experiments in Section 4, we show the proposed Global-Local Attention Transformers (GLAT) outperforms conventional transformers, as well as widely used graph-based models such as graph convolution nets and graph attention nets.

3 Method

In this section, we first formalize the task, and propose a novel formulation of visual commonsense in connection with visual perception. We then provide

an overview of the proposed architecture (Figure 1), followed by an in-depth description of each proposed module.

We define a scene graph as $G = (\mathcal{N}_e, \mathcal{N}_p, \mathcal{E}_s, \mathcal{E}_o)$, where \mathcal{N}_e is a set of entity nodes, \mathcal{N}_p is a set of predicate nodes, \mathcal{E}_s is a set of edges from each predicate to its subject (which is an entity node), and \mathcal{E}_o is a set of edges from each predicate to its object (that also is an entity node). Each entity node is represented with an entity class $c_e \in \mathcal{C}_e$ and a bounding box $b \in [0, 1]^4$, while each predicate node is represented with a predicate class $c_p \in \mathcal{C}_p$ and is connected to exactly one subject and one object. Note that this formulation of scene graph is slightly different from the conventional one [29], as we formulate predicates as nodes rather than edges. This tweak does not cause any limitation since every scene graph can be converted from the conventional representation to our representation. However, this formulation allows multiple predicates between the same pair of entities, and it also enables us to define a unified attention over all nodes no matter entity or predicate.

Given a training dataset with many images $I \in [0, 1]^{h \times w \times c}$ paired with ground truth scene graphs G_T , our goal is to train a model that takes a new image and predicts a scene graph that maximizes $p(G|I)$. This is equivalent of maximizing $p(I|G)p(G)$, which breaks the problem into what we call *perception* and *commonsense*. In our proposed intuition, commonsense is the mankind’s ability to predict which situations are possible and which are not, or in other words, what makes *sense* and what does not. This can be seen as a prior distribution $p(G)$ over all possible situations in the world, represented as scene graphs. Perception, on the other hand, is the ability to form symbolic belief from raw sensory data, which are respectively G and I in our case. Although the goal of computer vision is to solve the Maximum a Posteriori (MAP) problem (maximizing $p(G|I)$), neural nets often fail to estimate the posterior, unless the prior is explicitly enforced in the model definition [17]. This is while in computer vision, the prior is often overlooked, or inaccurately considered to be a uniform distribution, making MAP equivalent to Maximum Likelihood (ML), *i.e.*, finding G that maximizes $p(I|G)$ [20].

We propose the first method to explicitly approximate the MAP inference by devising an explicit prior model (commonsense). Since posterior inference is intractable, we propose a two-stage framework as an approximation: We first adopt any off-the-shelve SGG model as the *perception model*, which takes an input image and produces a perception-driven scene graph, G_P , that approximately maximizes the likelihood. Then we propose a *commonsense model*, which takes G_P as input, and produces a commonsense-driven scene graph, G_C , to approximately maximize the posterior, *i.e.*,

$$G_P = f_P(I) \approx \underset{G}{\operatorname{argmax}} p(I|G), \quad (1)$$

$$G_C = f_C(G_P) \approx \underset{G}{\operatorname{argmax}} p(I|G)p(G), \quad (2)$$

where f_P and f_C are the perception and commonsense models. The commonsense model can be seen as a graph-based extension of denoising autoencoders [25],




Fig. 2. The proposed Global-Local Attention Transformer (GLAT), and its training framework: We augment transformers with local attention heads to help them encode the structure of scene graphs within node embeddings. The decoder takes the embeddings of a perturbed scene graph and reconstructs the correct scene graph without having access to the image. Note this figure only shows the commonsense block of our overall pipeline shown in Figure 1.

which evidently can learn the generative distribution of data $[10, 1]$, that is $p(G)$ in our case. Accordingly, f_C can take any scene graph as input and produce a more plausible graph by only slightly changing the input. A key design choice here is the fact that f_C does not take the image as input. Otherwise, it would be hard to ensure it is purely learning commonsense and not perception.

Ideally, G_C is the best decision to make, since it maximizes the posterior distribution. However, in practice autoencoders tend to under-represent long-tailed distributions and only capture the modes. This means the commonsense model may fail to predict less common structures, in favor of more statistically rewarding alternatives. To alleviate this problem, we propose a *fusion module* that takes G_P and G_C as input, and outputs a fused scene graph, G_F , which is the final output of our system. This can be seen as a decision-making agent that has to decide how much to trust each model, based on how confident they are.

Figure 1 illustrates an overview of the proposed architecture. In the rest of this section, we elaborate each module in detail.

3.1 Global-Local Attention Transformers

We propose the first graph-based visual commonsense model, which learns a generative distribution over the semantic structure of real-world scenes, through a denoising autoencoder framework. Inspired by BERT [5], which reconstructs

masked tokens in a sentence through stacked layers of multi-head attention, we propose Global-Local Attention Transformers (GLAT) that take a graph with masked nodes as input, and reconstructs the missing nodes. Figure 2 illustrates how GLAT works. Given an input scene graph G_P , we represent node i as a one-hot vector $x_i^{(0)}$, that includes entity and predicate categories, as well as a special MASK class. We stack node representations as rows of a matrix $X^{(0)}$ for notation purposes.

GLAT takes $X^{(0)}$ as input and represents each node by encoding the structure and context. To this end, it applies L layers of multi-head attention on the input nodes. Each layer l creates new node representations $X^{(l)}$, by applying a linear layer on the concatenated output of that layer’s attention heads. More specifically,

$$X^{(l)} = \underset{h \in \mathcal{H}_l}{\text{concatenate}} \left[h(X^{(l-1)}) \right] \times W_l + b_l, \quad (3)$$

where \mathcal{H}_l is the set of attention heads for layer l , W_l and b_l are trainable fusion weights and bias for that layer, and the concatenation operates along columns. We use two types of attention head, namely global and local. Each node can attend to all other nodes through global attention, while only its neighbors through local attention. We further divide local heads based on the type of edge they use, in order to differentiate the way subjects and objects interact with predicates, and vice versa. Therefore, we can write:

$$\mathcal{H}_l = \mathcal{H}_l^G \cup \mathcal{H}_l^{LS} \cup \mathcal{H}_l^{LO}. \quad (4)$$

All heads within each subset are identical, except they have distinct parameters that are initialized and trained independently. Each global head h^G operates as a typical self-attention would:

$$h^G(X) = [q(X)^T k(X)] v(X), \quad (5)$$

where q, k, v are query, key, and value heads, each a fully connected network, typically (but not necessarily) with a single linear layer. A local attention is the same, except queries can only interact with keys of their immediate neighbor nodes. For instance in subject heads,

$$h^{LS}(X) = [q(X)^T k(X) \odot A_s] v(X), \quad (6)$$

where A_s is the adjacency matrix of subject edges, which is 1 between from each predicate to its subject and vice versa, and 0 elsewhere. We similarly define A_o and h^{LO} for object edges.

Once we get contextualized, structure-aware representations $x_i^{(L)}$ for each node i , we devise a simple decoder to generate the output scene graph G_C , using a fully connected network that classifies each node to an entity or predicate class, and another fully connected network that classifies each pair of nodes into an edge type (subject, object or no edge). We train the encoder and decoder end-to-end, by randomly adding noise to annotated scene graphs from Visual Genome, feeding the noisy graph to GLAT, reconstructing nodes and edges, and

comparing each with the original scene graph before perturbation. We train the network using two cross-entropy loss terms on the node and edge classifiers. The details of training including the perturbation process are explained in Section 4.1.

3.2 Fusing Perception and Commonsense

The perception and commonsense models each predict the output node categories using a classifier that computes a probability distribution over all classes by applying a softmax on its logits. The class with highest probability is chosen and assigned a confidence score equal to its softmax probability. More specifically, node i from G_P has a logit vector L_i^P that has $|\mathcal{C}_e|$ or $|\mathcal{C}_p|$ dimensions depending of whether it is an entity node or predicate node. Similarly node i from G_C has a logit vector L_i^C . Note that these two nodes correspond to the same entity or predicate in the image, since GLAT does not change the order of nodes. Then the confidence of each node can be written as

$$q_i^P = \max_j \frac{\exp(L_i^P[j])}{\sum_k \exp(L_i^P[k])}, \quad (7)$$

and similarly q_i^C is defined given L_i^C .

The fusion module takes each node of G_P and the corresponding node of G_C , and computes a new logit vector for that node, as a weighted average of L_i^P and L_i^C . The weights determine the contribution of each model in the final prediction, and thus have to be proportional to the confidence of each model. Therefore, we compute the fused logits as:

$$L_i^F = \frac{q_i^P L_i^P + q_i^C L_i^C}{q_i^P + q_i^C}. \quad (8)$$

Finally, a softmax is applied on L_i^F to compute the final classification distribution for node i .

4 Experiments

In this section, we describe our experiments on the Visual Genome (VG) dataset in detail. We first evaluate how well our GLAT model learns visual commonsense, by comparing it to other models on the task of masked scene graph reconstruction. Then we provide a statistical analysis of our model prediction to show the kinds of commonsense knowledge it acquires, and distinguish it from bias. Next, we evaluate how effective GLAT and our fusion mechanism are for the downstream task of SGG, when applied on various perception models. We also provide several examples of how the commonsense model corrects the perceived output, and how the fusion model combines the two.

4.1 Implementation details

We train the perception and commonsense models separately using the ground truth scene graphs G_T from VG [12], particularly the version most widely used for SGG [29], which has 150 entity and 50 predicate classes. We then stack commonsense on top of perception and fine-tune it on VG, this time with actual scene graphs generated by perception, to adapt to the downstream task. The fusion module does not have trainable parameters and is thus only used during inference. We use the 75k VG scene graphs for training all models, and use the other 25k for test. We hold a small portion of the train set for validation. Our GLAT model (and other baselines when applicable) have 6 layers, each with 8 attention heads, and has a 300-dimensional representation for each node. While training GLAT, we randomly mask 30% of the nodes, which is the average number of nodes mistaken by a typical SGG model. We average the classification loss over all nodes and edges classified by the decoder, no matter masked or not. For fine-tuning and inference, we prune the output of the perception model before feeding to GLAT, by keeping the top 100 most confident predicates and all entities connected to those.

4.2 Evaluating commonsense

Once GLAT is trained, we evaluate it on the same task of reconstructing ground truth VG graphs that are perturbed by randomly masking 30% of their nodes. We evaluate the accuracy of our model in classifying the masked nodes, and report the accuracy (Table 1) separately for entity nodes and predicate nodes, as well as overall. This is a good measure of how well the model has learned commonsense, because it mimics mankind’s ability to imagine what would a real-world scene look like, given some context. In Figure 2, for instance, given the fact that there is a person riding something that is masked, we can immediately tell it is probably a bike, a motorcycle, or a horse. If we also know there is a mountain behind the masked object, and the masked object has a face and legs (not shown in the figure for brevity), then we can more certainly imagine it is a horse. By incorporating the global context of the scene, as well as the local structure of the graph, GLAT is able to effectively imagine the scene and predict the class of the entity or predicate that was masked, at a significantly higher accuracy compared to all baselines.

More specifically, we compare GLAT to: (1) A transformer [5] that is the same as our model, except it only has global heads; (2) A Graph Attention Net [24] which is also the same as our model, but only with local heads; and (3) A Graph Convolutional Network [11], which has only one local head at each layer, and the attention is fixed to be equal for all neighbors of each node. We also compare our method with the frequency prior used by Zellers *et al.* [34], which can only be applied for masked predicates, and simply predicts the most frequent predicate given its subject and object. As Table 1 shows, our method significantly outperforms all aforementioned baselines, which are a good representative of any existing method to learn semantic graph reconstruction.

Table 1. Ablation study on Visual Genome. All numbers are in percentage, and graph constraint is enforced

Method	Entity	Predicate	Both
Triplet Frequency [34]	-	44.4	-
Graph Convolutional Nets [11] (local-only, fixed attention)	8.7	43.4	19.7
Graph Attention Nets [24] (local-only)	12.0	45.0	22.3
Transformers [5] (global-only)	14.0	42.3	22.9
Global-Local Attention Transformers (ours)	22.3	60.7	34.4

To provide a better sense of the commonsense knowledge our model learns, we apply GLAT on the entire VG test set, using the procedure detailed below (Section 4.3), and collect its prediction statistics in a diverse set of situations. We elaborate using an example, shown in the top left cell of Table 2. Out of all triplets from all scene graphs produced by our model, we collect those triplets that match the certain template of **person [X] horse**, and show our sorted top 5 predictions in terms of frequency. The 5 predicates most often predicted by our method between a **person** and a **horse** are **on**, **riding**, **near**, **watching**, and **behind**. These are all possible interactions between a person and a horse, and all follow the affordance properties of both **person** and **horse**. Nevertheless, when we get the same statistics from the output of a state-of-the-art scene graph generation model (IMP [29]), we observe that it frequently predicts **person wearing horse**, which does not follow the affordance of **horse**. This can be attributed to the high frequency of **wearing** in VG annotation, which biases the IMP model, while our commonsense model is prone to such bias, and has learned affordances through the self-supervised training framework.

Table 2 provides several more scenarios like this, demonstrating our proficiency in three types of commonsense: object affordance, intuitive physics, and object composition. As an example of physics, we choose the triplet template **[X] under bed**, and show that our model predicts plausible objects such as **pot**, **shoe**, **drawer**, **book**, and **sneaker**. This is while IMP predicts **bed under bed**, **counter under bed**, and **sink under bed**, which are all physically counter-intuitive. More interestingly, one of our frequent predictions, **book under bed**, is a composition that does not exist in training data, suggesting the knowledge acquired by GLAT is not merely a biased memory of frequent compositions in training data.

The last type of commonsense in our illustration is object composition, *i.e.*, the fact that certain objects are physical parts of other objects. For **[X] has ear**, we predict **head**, **cat**, **elephant**, **zebra**, and **person**, out of which **head has ear** and **person has ear** are not within the 10 most frequent triplets in training data that match the template. Yet our model frequently predicts them, demonstrating its unbiased knowledge. Not to mention, 4 out of 5 top predictions made by IMP are nonsensical.

Table 2. Prediction statistics of our method compared to IMP [29] in various situations, showcasing our model’s commonsense knowledge, and its robustness to dataset bias. Each row is designated for a certain type of commonsense, and has three examples in three pairs of columns. Each pair of columns show the top 5 most frequent triplets matching a certain template from our model’s prediction, compared to IMP. **Black** triplets are commonsensically correct, **red** triplets are wrong, **blue** are commonsensically correct but statistically rare in training data, and **green** are correct but never seen in training data.

	Template 1		Template 2		Template 3	
	IMP + GLAT	IMP	IMP + GLAT	IMP	IMP + GLAT	IMP
Object Affordance	person on horse person riding horse person near horse person watching horse person behind horse	person on horse person riding horse person near horse person wearing horse person behind horse	person has flower person with flower person near flower person watching flower person holding flower	person on flower person has flower person near flower person wearing flower person holding flower	person sitting on chair person sitting on bench person sitting on seat person sitting on stand person sitting on rock	person sitting on chair person sitting on bench person sitting on table person sitting on seat person sitting on person
Intuitive Physics	orange in basket fruit in basket banana in basket food in basket paper in basket	basket in basket man in basket woman in basket orange in basket fruit in basket	airplane above plane airplane above car airplane above bird airplane above beach airplane above vehicle	airplane above table	pot under bed shoe under bed drawer under bed book under bed sneaker under bed	bed under bed drawer under bed counter under bed sink under bed shoe under bed
Object Composition	head has ear cat has ear elephant has ear zebra has ear person has ear	leg has ear head has ear ear has ear tree has ear nose has ear	skier has head skier has hand skier has leg skier has arm skier has face	skier has leg skier has arm skier has pole skier has person skier has tree	leg of person leg of man leg of giraffe leg of zebra leg of elephant	leg of leg leg of man leg of person leg of tree leg of head

4.3 Evaluating scene graph generation

Now that we showed the efficacy of GLAT in learning visual commonsense and correcting perturbed scene graphs, we apply and evaluate it on the downstream task of scene graph generation. We adopt existing SGG models as our perception model, and compare their output G_P , to the ones corrected by our commonsense model G_C , as well as the final output of our system after fusion G_F . We compare those 3 outputs for 3 different choices of perception model, all of which have competitive state-of-the-art performance. More specifically, we use Iterative Message Passing (IMP [29]) as a strong baseline that is not augmented by commonsense. We also use Stacked Neural Motifs (SNM [34]) that late-fuse a frequency prior with their output, and Knowledge-Embedded Routing Networks (KERN [2]) that encode frequency prior within their internal message passing.

To evaluate, we conventionally compute the mean recall of the top 50 (mR@50) and top 100 (mR@100) triplets predicted by each model. Each triplet is considered correct if the subject, predicate, and object are all classified correctly, and the bounding box of the subject and object have more than 50% overlap (intersection over union) with the ground truth. We compute the recall for the triplets of each predicate class separately, and average over classes. The aforementioned metrics are measured in 2 sub-tasks: (1) SGCLS is the main scenario where we classify entities and predicates given annotated bounding boxes. This way the performance is not limited by proposal quality. (2) PREDCLS provides the model with ground truth object labels, which helps evaluation focus on predicate recognition accuracy. Table 3 shows the full comparison of all methods on all metrics. We observe that GLAT improves the performance of IMP which does not have

Table 3. The mean recall of our method compared to the state of the art on the task of scene graph generation, evaluated on the Visual Genome dataset [29], following the experiment settings of [34]. All baseline numbers were borrowed from [2], and all numbers are in percentage

Method	PREDCLS		SGCLS	
	mR@50	mR@100	mR@50	mR@100
IMP [29]	9.8	10.5	5.8	6.0
IMP + GLAT	11.1	11.9	6.2	6.5
IMP + GLAT + Fusion	12.1	12.9	6.6	7.0
SNM [34]	13.3	14.4	7.1	7.5
SNM + GLAT	13.6	14.6	7.3	7.8
SNM + GLAT + Fusion	14.1	15.3	7.5	7.9
KERN [2]	17.7	19.2	9.4	10.0
KERN + GLAT	17.6	19.1	9.3	10.0
KERN + GLAT + Fusion	17.8	19.3	9.9	10.4

commonsense, but does not significantly change the performance of SNM and KERN which already use dataset statistics. However, our full model which uses both the output of the perception model as well as commonsense model consistently improves SGG performance. In the supplementary material, we provide a more detailed analysis by breaking the results down into subgroups based on triplet frequency, and showing our performance boost is consistent in frequent and rare situations.

Finally, we provide several examples in Figure 3 to illustrate how our commonsense model fixes perception errors in difficult scenarios, and improves the robustness of our model. To save space, we merge the three scene graphs predicted by the perception, commonsense, and fusion models into a single graph, and emphasize any node or edge where these three models disagree. In example (a), the **chair** is not fully visible, and the visible part does not visually show the action of **sitting**, thus the perception model incorrectly predicts **wearing**, which is likely to be also affected by the bias due to the prevalence of **wearing** annotations in Visual Genome. However, it is trivial for the commonsense model that the affordance of **chair** is **sitting**. The fusion module correctly prefers the output of the commonsense model, due to its higher confidence. In (b), the perception model mistakes the **head** of the **bird** for a **hat**, due to the complexity of the lighting and the similarity of foreground and background colors. This might be also affected by the bias of **head** instances in VG, which are usually human heads, and the fact that **hat** instances typically co-occur with a **head**. Nevertheless, our commonsense model has the knowledge of object composition and knows **birds** typically have **heads** but not **hats**. Example (c) is an unusual case of **holding**, in terms of visual attributes such as arm pose. Hence, the perception model fails to predict **holding** correctly, while our commonsense model corrects that mistake by incorporating the affordance of **bottle**. Finally, in (d), the **person** is perceived **under** the **tower** due to the camera angle, but for the




Fig. 3. Example scene graphs generated by the perception, commonsense, and fusion modules, merged into one graph. Entities are shown as rectangular nodes and predicates are shown as directed edges from subject to object. For entities and predicates that are identically classified by the perception and commonsense model, we simply show the predicted label. But in cases where the perception and commonsense models disagree, we show both of their predictions as well as the final output chosen by the fusion module. We show mistakes in red, with the ground truth in parentheses.

commonsense model that is unlikely due to intuitive physics. Hence, it corrects the mistake and the fusion module accepts that fix. More examples are provided in the supplementary material.

5 Conclusion

We presented the first method to learn visual commonsense automatically from a scene graph corpus. Our method learns structured commonsense patterns, rather than simple co-occurrence statistics, through a novel self-supervised training strategy. Our unique way of augmenting transformers with local attention heads significantly outperforms transformers, as well as widely used graph-based models such as graph convolutional nets. Furthermore, we proposed a novel architecture for scene graph generation, which consists of two individual models, perception and commonsense, which are trained differently, and can complement each other under uncertainty, improving the overall robustness. To this end, we proposed a fusion mechanism to combine the output of those two models based on their confidences, and showed our model correctly determines when to trust its perception and when to fall back on its commonsense. Experiments show the effectiveness of our method for scene graph generation, and encourage future work to apply the same methodology on other computer vision tasks.

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. Alain, G., Bengio, Y.: What regularized auto-encoders learn from the data-generating distribution. *The Journal of Machine Learning Research* **15**(1), 3563–3593 (2014)
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. Chuang, C.Y., Li, J., Torralba, A., Fidler, S.: Learning to act properly: Predicting and explaining affordances from images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 975–983 (2018)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
6. Groth, O., Fuchs, F.B., Posner, I., Vedaldi, A.: Shapestacks: Learning vision-based physical intuition for generalised object stacking. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 702–717 (2018)
7. 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)
8. Jiang, C., Xu, H., Liang, X., Lin, L.: Hybrid knowledge routed modules for large-scale object detection. In: *Advances in Neural Information Processing Systems*. pp. 1552–1563 (2018)
9. Kato, K., Li, Y., Gupta, A.: Compositional learning for human object interaction. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 234–251 (2018)
10. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
11. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016)
12. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision* **123**(1), 32–73 (2017)
13. Kumar Singh, K., Divvala, S., Farhadi, A., Jae Lee, Y.: Dock: Detecting objects by transferring common-sense knowledge. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 492–508 (2018)
14. 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)
15. Li, Y., Tarlow, D., Brockschmidt, M., Zemel, R.: Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493* (2015)
16. Lu, J., Batra, D., Parikh, D., Lee, S.: VIlbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: *Advances in Neural Information Processing Systems*. pp. 13–23 (2019)
17. Malinin, A., Gales, M.: Predictive uncertainty estimation via prior networks. In: *Advances in Neural Information Processing Systems*. pp. 7047–7058 (2018)

18. Narasimhan, M., Schwing, A.G.: Straight to the facts: Learning knowledge base retrieval for factual visual question answering. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 451–468 (2018)
19. Qi, M., Wang, Y., Qin, J., Li, A.: Ke-gan: Knowledge embedded generative adversarial networks for semi-supervised scene parsing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5237–5246 (2019)
20. Romaszko, L., Williams, C.K., Moreno, P., Kohli, P.: Vision-as-inverse-graphics: Obtaining a rich 3d explanation of a scene from a single image. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 851–859 (2017)
21. Speer, R., Chin, J., Havasi, C.: Conceptnet 5.5: An open multilingual graph of general knowledge. In: Thirty-First AAAI Conference on Artificial Intelligence (2017)
22. Su, Z., Zhu, C., Dong, Y., Cai, D., Chen, Y., Li, J.: Learning visual knowledge memory networks for visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7736–7745 (2018)
23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
24. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
25. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th international conference on Machine learning. pp. 1096–1103 (2008)
26. Wang, T., Huang, J., Zhang, H., Sun, Q.: Visual commonsense r-cnn. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10760–10770 (2020)
27. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7794–7803 (2018)
28. 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)
29. 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)
30. Yang, J., Lu, J., Lee, S., Batra, D., Parikh, D.: Graph r-cnn for scene graph generation. In: Proceedings of the European conference on computer vision (ECCV). pp. 670–685 (2018)
31. Yu, R., Li, A., Morariu, V.I., Davis, L.S.: Visual relationship detection with internal and external linguistic knowledge distillation. In: Proceedings of the IEEE international conference on computer vision. pp. 1974–1982 (2017)
32. 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)
33. Zellers, R., Bisk, Y., Farhadi, A., Choi, Y.: From recognition to cognition: Visual commonsense reasoning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6720–6731 (2019)
34. 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)