

Differentiable Scene Graphs

Moshiko Raboh^{1*}, Roei Herzig^{1*}, Jonathan Berant^{1,4}, Gal Chechik^{2,3}, Amir Globerson¹

¹Tel Aviv University, ²Bar-Ilan University, ³NVIDIA Research, ⁴AI2

Abstract

Reasoning about complex visual scenes involves perception of entities and their relations. Scene Graphs (SGs) provide a natural representation for reasoning tasks, by assigning labels to both entities (nodes) and relations (edges). Reasoning systems based on SGs are typically trained in a two-step procedure: first, a model is trained to predict SGs from images, and next a separate model is trained to reason based on the predicted SGs. However, it would seem preferable to train such systems in an end-to-end manner. The challenge, which we address here is that scene-graph representations are non-differentiable and therefore it isn't clear how to use them as intermediate components. Here we propose Differentiable Scene Graphs (DSGs), an image representation that is amenable to differentiable end-to-end optimization, and requires supervision only from the downstream tasks. DSGs provide a dense representation for all regions and pairs of regions, and do not spend modelling capacity on regions of the image that do not contain objects or relations of interest. We evaluate our model on the challenging task of identifying referring relationships (RR) in three benchmark datasets: Visual Genome, VRD and CLEVR. Using DSGs as an intermediate representation leads to new state-of-the-art performance. The full code is available at <https://github.com/shikorab/DSG>.

1. Introduction

Understanding the full semantics of rich visual scenes is a complex task that involves detecting individual entities, as well as reasoning about the combination of entities and the relations between them. To represent entities and their relations jointly, it is natural to view them as a graph, where nodes are entities and edges represent relations. Such representations are often called *Scene Graphs* (SGs) [23]. Because SGs allow to explicitly reason about images, substantial efforts have been made to infer them from raw images [22, 23, 50, 33, 57, 17, 58].

*Equal Contribution.




Figure 1. **Differentiable Scene Graphs**: An intermediate “graph-like” representation that provides a distributed representation for each entity and pair of entities in an image. Differentiable scene graphs can be learned with gradient descent in an end-to-end manner, only using supervision about a downstream visual reasoning task (referring relationships here).

While scene graphs have been shown to be useful for various tasks [22, 23, 20], using them as a component in a visual reasoning system is challenging: (a) Because scene graphs are discrete representations, it is difficult to learn them in an end-to-end fashion from a downstream task. (b) The alternative is to pre-train SG predictors separately from supervised data, but this requires arduous and prohibitive manual annotation. Moreover, pre-trained SG predictors have *low coverage*, because the set of labels they are pre-trained on rarely fits the needs of a downstream task. For example, given an image of a parade and a question “*point to the officer on the black horse*”, that horse might not be a node in the graph, and the term “officer” might not be in the vocabulary. Given these limitations, it is an open question how to make scene graphs useful for visual reasoning applications.

In this work, we describe *Differentiable Scene-Graphs* (DSG), which address the above challenges (Figure 1). DSGs are an **intermediate representation trained end-to-end from the supervision for a downstream reasoning task**. The key idea is to relax the discrete properties of scene graphs such that each entity and relation is described with a dense differentiable descriptor.

We demonstrate the benefits of DSGs in the task of resolving *referring relationships* (RR) [27] (see Figure 1). Here, given an image and a triplet query (*subject, relation, object*), a model has to find the bounding boxes of the sub-

ject and object that participate in the relation.

We train an RR model with DSGs as an intermediate component. As such, DSGs are not trained with direct supervision about entities and relations, but using several supervision signals about the downstream RR task. We evaluate our approach on three standard RR datasets: Visual Genome [28], VRD [35] and CLEVR [21], and find that DSGs substantially improve performance compared to state-of-the-art approaches [35, 27].

To conclude, our novel contributions are: (1) A new *Differentiable Scene-Graph* representation for visual reasoning, which captures information about multiple entities in an image and their relations. We describe how DSGs can be trained end-to-end with a downstream visual reasoning task without direct supervision of manually annotated scene-graphs. (2) A new architecture for the task of referring relationships, using DSGs as its central component. (3) New state-of-the-art results on the task of referring relationships on the Visual Genome, VRD and CLEVR datasets.

2. Referring Relationship: The Learning Setup

In the referring relationship task [27] we are given an image I and a subject-relation-object query $q = \langle s, r, o \rangle$. The goal is to output a bounding box \mathcal{B}_s for the subject, and another bounding box \mathcal{B}_o for the object. In practice there are sometimes several boxes for each. See Fig. 1 for a sample query and expected output.

Following [27], we focus on training a referring relationship predictor from labeled data. Namely, we use a training set consisting of images, queries and the correct boxes for these queries. We denote these by $\{(I_j, q_j, (\mathcal{B}_j^s, \mathcal{B}_j^o))\}_{j=1}^N$. As in [27], we assume that the vocabulary of query components (subject, object and relation) is fixed.

In our model, we break this task into two components that we optimize in parallel. We fine-tune the position of bounding boxes such that they cover entities tightly, and we also label each box as one of the following four possible labels. The labels “Subject” and “Object” correspond to the ‘s’ and ‘o’ entities in the query. The label “Other” corresponds to boxes that contain entities (e.g., person or any other category that can appear as a subject or object in queries) that are not the subject or the object of the query. Finally, the label “Background” corresponds to boxes that do not contain an entity. We refer to the above two modules as **Box Refiner** and **Referring Relationships Classifier**.

3. Differentiable Scene Graphs

We begin by discussing the motivation and potential advantages of using intermediate scene-graph-like representations, as compared to standard scene graphs. We then explain how DSGs fit into the full architecture of our model.

3.1. Why use intermediate DSG layers?

A “perfect” scene graph (representing all entities and relations) captures most of the information needed for visual reasoning, and thus should be useful as an intermediate representation. Such a SG can then be used by downstream reasoning algorithms, using the predicted SG as an input. Unfortunately, learning to predict “perfect” scene graphs for any downstream task is unlikely. First, there is rarely enough data to train good SG predictors, and second, learning to predict SGs in a way that is independent of the downstream task, tends to yield less relevant SGs.

Instead, we propose an intermediate representation, which we refer to as a “Differentiable Scene Graph” layer (DSG). A DSG captures the relational information as in a scene graph but can be trained end-to-end in a task-specific manner (Fig. 2). Like SGs, a DSG keeps descriptors for visual entities and their relations. Unlike SGs, whose nodes and edges are annotated by discrete values (labels), a DSG contains a dense distributed representation vector for each detected entity (referred to as a *node descriptor*) and each pair of entities (referred to as an *edge descriptor*). These representations are themselves learned functions of the input image (see supplementary for more details). Like SGs, a DSG only describes candidate boxes which cover entities of interest and their relations. Unlike SGs, each DSG descriptor captures not only the local information about a node, but also information about its context. Most importantly, because DSGs are differentiable, they are used as input to downstream visual-reasoning modules (in our case, a referring relationships module).

DGSs provide several computational and modelling advantages:

Differentiability. Because node and edge descriptors are differentiable functions of detected boxes, and are fed into a differentiable reasoning module, the entire pipeline can be trained with gradient descent.

Dense descriptors. By keeping dense descriptors for nodes and edges, the DSG keeps more information about possible semantics of nodes and edges, instead of committing too early to hard sparse representations. This allows it to better fit downstream tasks.

Supervision using downstream tasks. Collecting supervised labels for training scene graphs is hard and costly. DGSs can be trained using training data that is available for downstream tasks, saving costly labeling efforts. On the other hand, when labeled scene graphs are available for given images, that data can be used when training the DSG, using an additional loss component.

Holistic representation. DSG descriptors are computed by integrating global information from the entire image using graph neural networks (see supplemental materials). Combining information across the image increases the accuracy of object and relation descriptors.

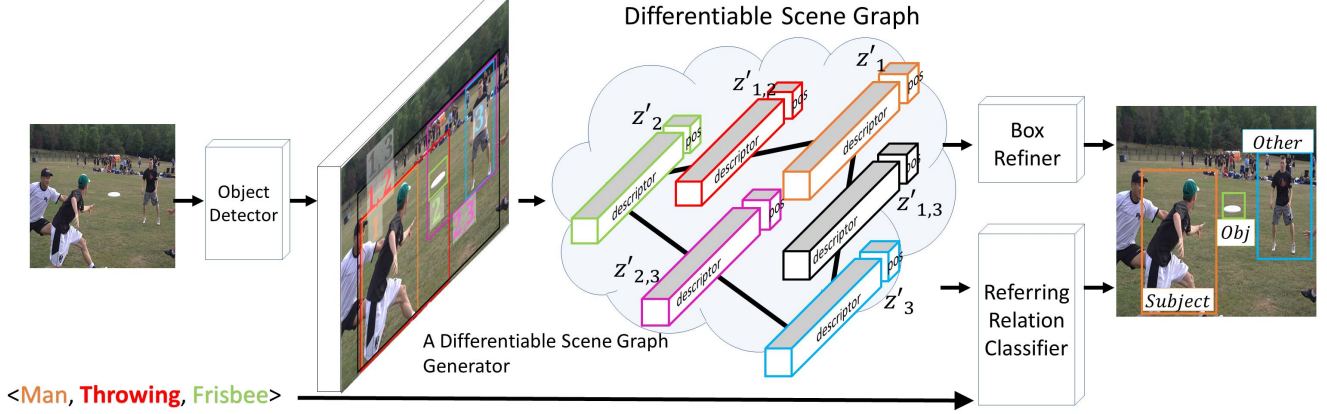


Figure 2. **The proposed architecture.** The input consists of an image and a relationship query triplet $\langle \text{subject}, \text{relation}, \text{object} \rangle$. (1) A detector produces a set of bounding box proposals. (2) An *ROI-Align* layer extracts object features from the backbone using the boxes. In parallel, every pair of box proposals is used for computing a union box, and pairwise features are extracted in the same way as object features. (3) These features are used as inputs to a Differentiable Scene-Graph Generator Module which outputs the Differential Scene Graph, a set of node and edge features that result from applying a graph convolutional network to the input features. (4) The DSG is used for both refining the original box proposals, as well as a Referring Relationships Classifier, which classifies each bounding box proposal as either Subject, Object, Other or Background. The ground-truth label of a proposal box will be Other if this proposal appears in another query relationship for this image. Otherwise the ground truth label will be Background.

3.2. The DSG Model for Referring relationships

We now describe how DSGs can be combined with other modules to solve a visual reasoning task. The architecture of the model is illustrated in Fig. 2. First, the model extracts bounding boxes for entities and relations in the image. Next, it creates a differentiable scene-graph over these bounding boxes. Then, DSG features are used by two output modules, aimed at answering a referring-relationship query: a Box Refiner module that refines the bounding box of the relevant entities, and a Referring Relationships Classifier module that classifies each box as Subject, Object, Other or Background. We now describe these components in more detail.

Object Detector. We detect candidate entities using a standard region proposal network (RPN) [43], and denote their bounding boxes by b_1, \dots, b_B (B may vary between images). We also extract a feature vector f_i for each box and concatenate it with the box coordinates, yielding $z_i = [f_i; b_i]$. See details in the supplemental material

Relation Feature Extractor. Given any two bounding boxes b_i and b_j we consider the smallest box that contains the two boxes (their “union” box). We denote this “relation box” by $b_{i,j}$ and its features by $f_{i,j}$. Finally, we denote the concatenation of the features $f_{i,j}$ and box coordinates $b_{i,j}$ by $z_{i,j} = [f_{i,j}; b_{i,j}]$.

Differentiable Scene-Graph Generator. As discussed above, the goal of the DSG Generator is to transform the above features z_i and $z_{i,j}$ into differentiable representations of the underlying scene graph. Namely, map these features into a new set of dense vectors z'_i and $z'_{i,j}$ rep-

resenting entities and relations. This mapping is intended to incorporate the relevant context of each feature vector. Namely, the representation z'_i contains information about the i^{th} entity, together with its image-wide context.

There are various possible approaches to achieve this mapping. Here we use the model proposed by [17], which uses a graph neural network for this transformation (see supplemental material).

Multi-task objective. In many domains, training with multi-task objectives can improve the accuracy of individual tasks, because auxiliary tasks operate as regularizers, pushing internal representations away from overfitting and towards capturing useful properties of the input. We follow this idea here and define a multi-task objective that has three components: (a) a Referring Relationships Classifier matches boxes to subject and object query terms. (b) A Box Refiner predicts accurate tight bounding boxes. (c) A Box Labeler recognizes visual entities in boxes if relevant ground truth is available.

Fig. 3 illustrates the effect of the first two components, and how they operate together to refine the bounding boxes and match them to the query terms. Specifically, Fig. 3c, shows how box-refinement produces boxes that are tight around objects and subjects, and Fig. 3d shows how RR classification matches boxes to query terms.

(A) Referring Relationships Classifier. Given a DSG representation, we use it for answering referring relationship queries. Recall that the output of an RR query $\langle \text{subject}, \text{relation}, \text{object} \rangle$ should be bounding boxes $\mathcal{B}_s, \mathcal{B}_o$ containing subjects and objects that participate in the query relation. Our model has already computed B bounding boxes

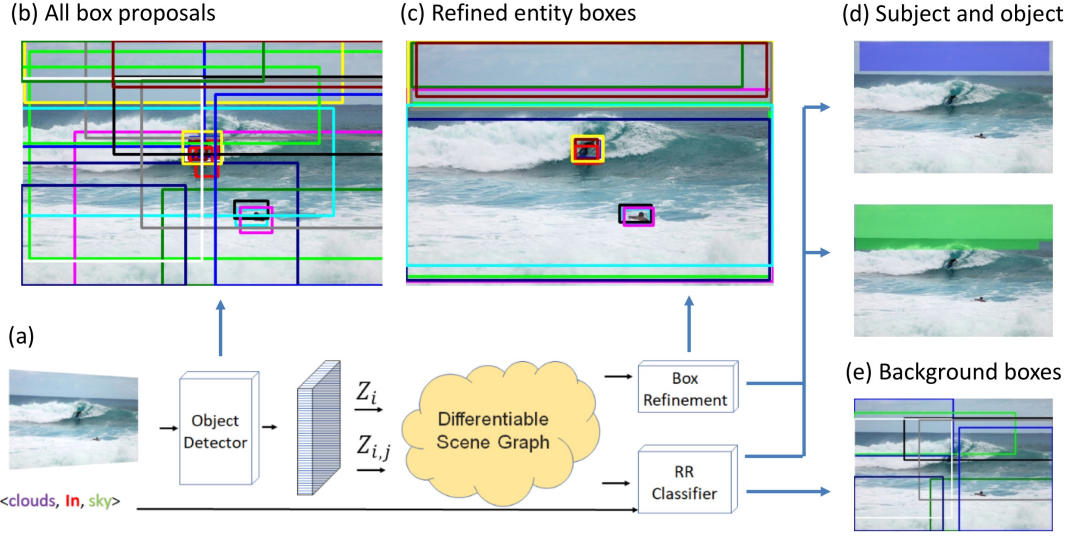


Figure 3. The effect of box refinement and RR classification. (a) The DSG network is applied to an input image. (b) The *object detector* component generates box proposals for entities in the image. (c) The *RR classifier* component uses information from the DSG to label candidate boxes as *object* or *subject* entities. Then, the *box refinement* component also uses DSG information, this time to improve box locations for those boxes labeled as entities by RR classifier. Here, boxes are tuned to focus on the most relevant entities in the image: the two “men”, the “surfboard”, the “sky” and the “ocean”. (d) Once the RR classifier labels entity boxes, it can correctly refer to the entities in the query $\langle \text{cloud, in, sky} \rangle$ (sky in green, clouds in violet). (e) Examples of candidate boxes classified as *background*.

b_i , as well as representations z'_i for each box. We next use a prediction model $F_{RRC}(z'_i, q)$ that takes as input features describing the bounding box and the query, and outputs one of four labels $\{\text{Subject, Object, Other, Background}\}$ (see Sec. 2). Denote the logits generated by this classifier for the i^{th} box by $r_i \in \mathbb{R}^4$. The output set \mathcal{B}_s (or \mathcal{B}_o) is simply the set of bounding boxes classified as *Subject* (or *Object*). See supplemental materials for further implementation details.

(B) Box Refiner. The DSG is also used for further refinement of the bounding-boxes generated by the RPN network. The idea is that additional knowledge about image context can be used to improve the coordinates of a given entity. This is done via a network $F_{BR}(b_i, z'_i)$ that takes as input the RPN box coordinates and a differentiable representation z'_i for box i , and outputs new bounding box coordinates. See Fig. 3 for an illustration of box refinement, and the supplemental material for further details.

(C) Optional auxiliary losses: Scene-Graph Labeling. In addition to the *Box Refiner* and *Referring Relationships Classifier* modules described above, one can also use supervision about labels of entities and relations if these are available at training time. Specifically, we train an object-recognition classifier operating on boxes, which predicts the label of every box for which a label is available. This classifier is trained as an auxiliary loss, in a multi-task fashion, and is described in detail below.

4. Training with Multiple Losses

We next explain how our model is trained for the RR task, and how we can also use the RR training data for supervising the DSG component. We train with a weighted sum of three losses: (1) Referring Relationships Classifier (2) Box Refiner (3) Optional Scene-Graph Labeling loss. We now describe each of these components. Additional details are provided in the supplemental material.

4.1. Referring Relationship Classification Loss

The *Referring Relationships Classifier* (Sec. 3.2) outputs logits r_i for each box, corresponding to its prediction (subject, object, etc.). To train these logits, we need to extract their ground-truth values from the training data. Recall that a given image in the training data may have multiple queries, and so may have multiple boxes that have been tagged as subject or object for the corresponding queries. To obtain the ground-truth for box i and query $q = \langle s, r, o \rangle$ we take the following steps. First, we find the ground-truth box that has maximal overlap with box i . If this box is either a subject or object for the query q , we set r_i^{gt} to be *Subject* or *Object* respectively. Otherwise, if the overlap with a ground-truth box for a different image-query is greater than 0.5, we set $r_i^{gt} = \text{Other}$ (since it means there is some other entity in the box), and we set $r_i^{gt} = \text{Background}$ if the overlap is less than 0.3. If the overlap is in $[0.3, 0.5]$ we do not use the box for training. For instance, given a

query $\langle \text{woman}, \text{feeding}, \text{giraffe} \rangle$ with ground-truth boxes for “woman” and “giraffe”, consider the box in the RPN that is closest to the ground-truth box for “woman”. Assume the index of this box is 7. Similarly, assume that the box closest to the ground-truth for “giraffe” has index 5. We would have $\mathbf{r}_7^{gt} = \text{Subject}$, $\mathbf{r}_5^{gt} = \text{Object}$ and the rest of the \mathbf{r}_i^{gt} values would be either `Other` or `Background`. Given these ground-truth values, the Referring Relationship Classifier Loss is simply the sum of cross entropies between the logits \mathbf{r}_i and the one-hot vectors corresponding to \mathbf{r}_i^{gt} .

4.2. Box Refiner Loss

To train the Box Refiner, we use a smooth L_1 loss between the coordinates of the refined (predicted) boxes and their ground truth ones.

4.3. Scene-Graph Labeling Loss

When ground-truth data about entity labels is available, we can use it as an additional source of supervision to train the DSG. Specifically, we train two classifiers. A classifier from features of entity boxes \mathbf{z}'_i to the set of entity labels, and a classifier from features of relation boxes $\mathbf{z}'_{i,j}$ to relation labels. We then add a loss to maximize the accuracy of these classifiers with respect to the ground truth box labels.

4.4. Tuning the Object Detector

In addition to training the DSG and its downstream visual-reasoning predictors, the object detector RPN is also trained. The output of the RPN is a set of bounding boxes. The ground-truth contains boxes that are known to contain entities. The goal of this loss is to encourage the RPN to include these boxes as proposals. Concretely, we use a sum of two losses: First, an RPN classification loss, which is a cross entropy over RPN anchors where proposals of 0.8 overlap or higher with the ground truth boxes were considered as positive. Second, an RPN box regression loss which is a smooth L1 loss between the ground-truth boxes and proposal boxes.

5. Experiments

In the following sections we provide details about the datasets, training, baselines models, evaluation metrics, model ablations and results. Implementation details of the model are provided in the supplemental material.

5.1. Datasets

We evaluate the model in the task of referring relationships on three datasets, each exhibiting a unique set of characteristics and challenges.

CLEVR [21]. A synthetic dataset generated from scene-graphs with four spatial relations: “left”, “right”, “front”

	Visual Genome		Average IOU		CLEVR	
	subject	object	VRD subject	VRD object	subject	object
SS [29]	0.399	0.469	0.320	0.371	0.740	0.740
CO [10]	0.414	0.490	0.347	0.389	0.691	0.691
VRD [35]	0.417	0.480	0.345	0.387	0.734	0.732
SASS [27]	0.421	0.482	0.369	0.410	0.778	0.778
NO-DSG	0.412	0.47	0.333	0.366	0.937	0.937
DSG	0.489	0.539	0.4	0.435	0.963	0.963

Table 1. **Comparison with baselines.** Test-set mean IOU in the referring relationship task for the baselines in Sec. 5.3 and the Differentiable Scene Graph (DSG) model. Results are also reported for a NO-DSG model (see Sec. 6.2) which classifies the referring relationship directly from the RPN output.

	Average IOU	
	subject	object
TWO STEP	0.430 ± 0.0014	0.491 ± 0.0014
NO-DSG	0.405 ± 0.0013	0.461 ± 0.0013
DSG -SGL	0.455 ± 0.0014	0.511 ± 0.0013
DSG -BR	0.469 ± 0.0014	0.519 ± 0.0014
DSG	0.477 ± 0.0014	0.528 ± 0.0014

Table 2. **Model ablations:** Results (including standard errors) for DSG variants on the validation set of the Visual Genome dataset. DSG values slightly differ from Table 1 which reports IOU on the test set. The various models are described in Sec. 6.2.

and “behind”, and 48 entity categories. It has over 5M relationships where 33% are ambiguous entities (multiple entities of the same type in an image).

VRD [35]. The Visual Relationship Detection dataset contains 5,000 images with 100 entity categories and 70 relation categories. In total, VRD contains 37,993 relationship annotations with 6,672 unique relationship types and 24.25 relations per entity category. 60.3% of these relationships refer to ambiguous entities.

Visual Genome [28]. VG is the largest public corpus for visual relationships in real images, with 108,077 images annotated with bounding boxes, entities and relations. On average, images have 12 entities and 7 relations per image. In total, there are over 2.3M relationships where 61% of those refer to ambiguous entities.

For a proper comparison with previous results [27], we used the data from [27] including the same entity and relation categories, query relationships and data splits.

5.2. Evaluation Metrics

We compare our model to previous work using the average IOU for subjects and for objects. To compute the average subject IOU, we first generate two $L \times L$ binary attention maps: one that includes all the ground truth boxes labeled as `Subject` (recall that few entities might be labeled as `Subject`) and the other includes all the box proposals

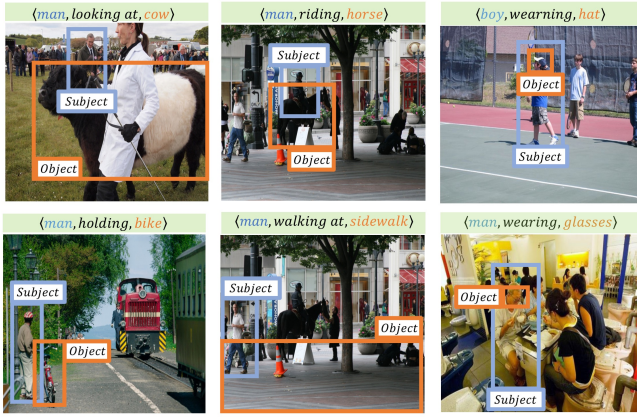


Figure 4. Qualitative examples demonstrating successful predictions of the DSG model (six left panels) and errors (six right panels). The right panels illustrate common failure cases for each error type. **a.** *Missed detection*: the detector missed the glasses on the table. **b,c.** *Misclassified object*:, the cake is detected but classified as a background. **d.** *Misclassified relation*: The box classified as *Subject* is indeed a man but it is not the man that has the required relation with the skate. **e,f.** *Multiplicity*, Either too few or too many GT boxes are classified as *Subject* or *Object*.

predicted as *Subject*. If no box is predicted as *Subject*, the box with the highest score for the label *Subject* is included in the predicted attention map. We then compute the Intersection-Over-Union between the binary attention maps. For a proper comparison with previous work [27], we use $L = 14$. The object boxes are evaluated similarly.

5.3. Baselines

The Referring Relationship task was introduced recently [27], and the SSAS model was proposed as a possible approach (see below). We report the results for the baseline models in [27]. When evaluating our Differentiable Scene-Graph model, we use exactly the evaluation setting as in [27] (i.e., same data splits, entity and relation categories). The baselines reported are:

1. SYMMETRIC STACKED ATTENTION SHIFTING (SSAS): [27] An iterative model that localizes the relationship entities using attention shift component learned for each relation.
2. SPATIAL SHIFTS [29]: Same as SSAS, but with no iterations and by replacing the shift attention mechanism with a simple statistical spatial shift for each relation label.
3. CO-OCCURRENCE [10]: Uses an embedding of the subject and object pair for attending over the image features.
4. VISUAL RELATIONSHIP DETECTION (VRD) [35]: Similar to Co-Occurrences model, but with an additional relationship embedding.




Figure 5. A typical image from the CLEVR [21] dataset. The image was trimmed to focus on areas with visual content.

6. Results

Table 1 provides average IOU for *Subject* and *Object* over the three datasets described in Sec. 5.1. We compare our model to four baselines described in Sec. 5.3. Our Differentiable Scene-Graph approach outperforms all baselines in terms of the average IOU.

Our results for the CLEVR dataset are significantly better than those in [27]. Because CLEVR objects have a small set of distinct colors (Fig 5), object detection in CLEVR is much easier than in natural images, making it easier to achieve high IOU. The baseline model without the DSG layer (NO-DSG) is an end-to-end model with a two-stage detector in contrast to [27] and already improves strongly over prior work with 93.7%, and our novel DSG approach further improves to 96.3% (reducing error by 50%).

6.1. Analysis of Success and Failure Cases.

Fig. 4 shows example of success cases and failure cases. We further analyzed the types of common mistakes and

their distribution. Since DSGs depend on box proposals, they are sensitive to the quality of the object detectors. Manual inspection of images revealed four main error types: **(1)** 30%: Detector failed: the relevant box is missing from the box proposal list. **(2)** 23.3% *Subject* or *Object* detected but classified as *Other* or as *Background*. **(3)** 16.6%: Relation misclassified. The entities classified as *Subject* and *Object* match the query, but without the required relation. **(4)** 16.6%: *Multiplicity*. Either too few or too many of the GT boxes are classified as *Subject* or *Object*. **(5)** 13.3%: Other, including incorrect GT, and hard-to-decide cases.

6.2. Model Ablations

To gain further insight into the performance of the DSG model we performed the following ablations. First, since the model is trained with three loss components, we quantify the contribution of the Box Refinement loss and the Scene-Graph Labeling loss (it is not possible to omit the Referring Relationships Classifier loss). We further evaluate the contribution of the DSG compared with a two-step approach which first predicts an SG, and then reasons over it. We compare the following models:

1. DSG: The Differentiable Scene-Graph model described in Sec. 3.2 and trained as described in Sec. 4.
2. TWO STEPS: Two-step model. We first predict a scene-graph, and then match the query with the SG. The SG predictor consists of the same components used in the DSG: A box detector, DSG dense descriptors, and an SG labeler. It is trained with the same set of SG labels used for training the DSG. Details in the supplemental material.
3. DSG -SGL: DSG without the Scene-Graph Labeling component described in Sec. 4.3).
4. DSG -BR: DSG where the *Box Refiner* component of Section 4.2 is replaced with fine tuning the coordinates of the box proposal using the visual features f_i extracted by the Object Detector. This variant allows us to quantify the benefit of refining the box proposals based on the differentiable representation of the scene.
5. NO-DSG: A baseline model that does not use the DSG representations. Instead, the model includes only an Object Detector and a RR classifier. The RR classifier uses the f_i features extracted by the Object Detector instead of the z'_i features. This model allows us to quantify the benefit of the differentiable scene representation for RR classification.

Table 2 provides results of ablation experiments for the Visual Genome dataset [28] on the validation set. All model variants based on scene representation perform better than the model that does not use the DSG representation (i.e., DSG -SG), demonstrating the power of contextu-




Figure 6. Inferring a Scene Graph from a DSG. Applying the RPN to this image results in 28 boxes. In (a) we show five of these, which received the largest weight in the attention model (details in the supplemental material.) within the DSG generator (Sec. 3.2). As mentioned in Sec. 4.3 in “Scene-Graph Labeling Loss” we can use the DSG for generating a labeled scene graph, corresponding to a fixed set of entities and relations. (b) shows this scene graph (i.e., the output of the classifiers predicting entity labels and relations), restricted to the largest confidence relations. It can be seen that most relations are correct, despite not having trained this model on complete scene graphs.

alized scene representation. The DSG model outperforms all model ablations, illustrating the improvements achieved by using partial supervision for training the differentiable scene-graph. Fig. 7 illustrates the effect of ablating various components of the model.

6.3. Inferring SGs from DSGs

The DSG is designed as a dense representation of objects and relations in the scene. It is thus natural to use it to predict these. This is easy to do in our context, since in Sec. 4.3 we in fact train such classifiers as an auxiliary task. Thus, for a given image we can construct a scene-graph out of the outputs of these objects and relation classifiers.

Fig. 6 illustrates the result of this process, showing a Scene-Graph inferred from the DSG. The predicted graph is indeed largely correct, even though it was not directly trained for this task (but rather from partial supervision). We further analyzed the accuracy of predicted SGs by comparing to ground-truth SG on visual genome (complete SGs were not used for training, only for analysis). SGs decoded from DSGs achieve accuracy of 76% for object labels and 70% for relations (calculated for proposals with $\text{IOU} \geq 0.8$).

7. Related Work

Graph Neural Networks. Recently, major progress has been made in constructing graph neural networks (GNN). These refer to a class of neural networks that operate directly on graph-structured data by passing local messages [11, 30]. Variants of GNNs have been shown to be highly effective at relational reasoning tasks [46], classification of graphs [3, 6, 39, 7] and classification of nodes in large




Figure 7. Comparing failures of ablation models with DSG predictions. The top row shows DSG results, while the bottom row shows results from different ablation models as specified in Sec. 6.2. In the first column, TWO STEP model the SG did not include the shirt of one of the the men, therefore this "subject" prediction was missed. In the second column, the DSG -SGL predicted failed to distinct between few entity classes "woman" and "child". In the third column, the DSG refine the box of "sky" to cover all of the sky area. In the last column, the NO-DSG didn't classify the "object" box correctly.

graphs [25, 14]. The expressive power of GNNs has also been studied in [17, 56]. GNNs have also been applied to visual understanding in [17, 51, 49, 16] and control [45, 2]. Similar aggregation schemes have also been applied to object detection [18]. Our goal here is to generate DSG such that each object descriptor encompasses not only the local information about the object, but also information about its context within the scene. To achieve this we use the GNN proposed by [17].

Visual Relationships and Scene Graphs. Earlier work aimed to leverage visual relationships for improving detection [44], action recognition [16], few shot [5, 9], pose estimation [8], semantic image segmentation [13] or detection of human-object interactions [52, 40, 31]. Lu *et al.* [35] were the first to formulate detection of visual relationships as a separate task. They learn a likelihood function that uses a language prior based on word embeddings for scoring visual relationships and constructing SGs. SGs provide a compact representation of the semantics of an image. Previous SG prediction works used attention [41, 12] or neural message passing [50]. [38] suggested to predict graphs directly from pixels in an end-to-end manner. [57] considers global context using an RNN by reading sequentially the independent predictions for each entity and relation and then refines those predictions. SGs have been shown to be useful for semantic-level interpretation and reasoning about a visual scene [20, 1, 15, 47]. Extracting SGs from images provides a semantic representation that can later be used for reasoning, question answering [54, 19, 32], and image retrieval [22, 42]. Using SGs for reasoning tasks is challenging. Instead, we propose an intermediate representation which captures the relational information as in SGs but can

be trained end-to-end in a task-specific manner.

Referring Relationships. The RR task is closely related to the task of referring expressions, where an entity in an image needs to be identified given a natural language expression. Several recent works considered using context for this task [24, 26, 4, 48, 53, 34]. [36] described a model that has two parts: one for generating expressions that point to an entity in a discriminative fashion and a second for understanding these expressions and detecting the referred entity. [55] explored the role of context and visual comparison with other entities in referring expressions. Modelling context was also the focus of [37], using a multi-instance-learning objective. RR [27] as opposed to referring expression, focuses on the vision side rather than the language side by forming a simple structured query that requires modeling interactions between the image entities. [27] also introduce an explicit iterative model that localizes the two entities in the RR task, conditioned on one another. We use the RR task to demonstrate the power of our semantic latent representation, resulting in a new state of the art results on three vision datasets that contains visual relationships.

8. Conclusion

This work is motivated by the assumption that accurate reasoning about images may require access to a detailed representation of the image. While scene graphs provide a natural structure for representing relational information, it is hard to train very dense SGs in a fully supervised manner, and for any given image, the resulting SGs may not be appropriate for downstream reasoning tasks. Here we advocate DSGs, an alternative representation that captures the information in SGs, which is continuous and can be trained jointly with downstream tasks. Our results, both qualitative (Fig 4) and quantitative (Table 1,2), suggest that DSGs effectively capture scene structure, and that this can be used for down-stream tasks such as referring relationships.

One natural next step is to study such representations in additional downstream tasks that require integrating information across the image. Some examples are caption generation and visual question answering. DSGs can be particularly useful for VQA, since many questions are easily answerable by scene graphs (e.g., counting questions and questions about relations). Another important extension to DSGs would be a model that captures high-order interactions, as in a hyper-graph. Finally, it will be interesting to explore other approaches to training the DSG, and in particular finding ways for using unlabeled data for this task.

Acknowledgments: This project was funded by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant ERC HOLI 819080).

References

- [1] O. Ashual and L. Wolf. Specifying object attributes and relations in interactive scene generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4561–4569, 2019.
- [2] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. F. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, Ç. Gülçehre, F. Song, A. J. Ballard, J. Gilmer, G. E. Dahl, A. Vaswani, K. Allen, C. Nash, V. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, and R. Pascanu. Relational inductive biases, deep learning, and graph networks. *CoRR*, abs/1806.01261, 2018.
- [3] J. Bruna and S. Mallat. Invariant scattering convolution networks. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):1872–1886, 2013.
- [4] D.-J. Chen, S. Jia, Y.-C. Lo, H.-T. Chen, and T.-L. Liu. See-through-text grouping for referring image segmentation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [5] V. Chen, P. Varma, R. Krishna, M. Bernstein, C. Re, and L. Fei-Fei. Scene graph prediction with limited labels. In *International Conference on Computer Vision*, 2019.
- [6] H. Dai, B. Dai, and L. Song. Discriminative embeddings of latent variable models for structured data. In *Int. Conf. Mach. Learning*, pages 2702–2711, 2016.
- [7] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Neural Inform. Process. Syst.*, pages 3837–3845, 2016.
- [8] C. Desai and D. Ramanan. Detecting actions, poses, and objects with relational phraselets. In *ECCV*, pages 158–172, 2012.
- [9] A. Dornadula, A. Narcomey, R. Krishna, M. Bernstein, and L. Fei-Fei. Visual relationships as functions: Enabling few-shot scene graph prediction. In *ArXiv*, 2019.
- [10] C. Galleguillos, A. Rabinovich, and S. Belongie. Object categorization using co-occurrence, location and appearance. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2008.
- [11] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- [12] N. Gkanatsios, V. Pitsikalis, P. Koutras, and P. Maragos. Attention-translation-relation network for scalable scene graph generation. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- [13] A. Gupta and L. S. Davis. Beyond nouns: Exploiting prepositions and comparative adjectives for learning visual classifiers. In *ECCV*, pages 16–29, 2008.
- [14] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Neural Inform. Process. Syst.*, pages 1024–1034, 2017.
- [15] R. Herzig, A. Bar, H. Xu, G. Chechik, T. Darrell, and A. Globerson. Learning canonical representations for scene graph to image generation. *arXiv preprint arXiv:1912.07414*, 2019.
- [16] R. Herzig, E. Levi, H. Xu, H. Gao, E. Brosh, X. Wang, A. Globerson, and T. Darrell. Spatio-temporal action graph networks. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- [17] R. Herzig, M. Raboh, G. Chechik, J. Berant, and A. Globerson. Mapping images to scene graphs with permutation-invariant structured prediction. In *Advances in Neural Information Processing Systems (NIPS)*, 2018.
- [18] H. Hu, J. Gu, Z. Zhang, J. Dai, and Y. Wei. Relation networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3588–3597, 2018.
- [19] R. Hu, A. Rohrbach, T. Darrell, and K. Saenko. Language-conditioned graph networks for relational reasoning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [20] J. Johnson, A. Gupta, and L. Fei-Fei. Image generation from scene graphs. *arXiv preprint arXiv:1804.01622*, 2018.
- [21] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. B. Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. *CoRR*, 2016.
- [22] J. Johnson, R. Krishna, M. Stark, L. Li, D. A. Shamma, M. S. Bernstein, and F. Li. Image retrieval using scene graphs. In *Proc. Conf. Comput. Vision Pattern Recognition*, pages 3668–3678, 2015.
- [23] J. Johnson, R. Krishna, M. Stark, L.-J. Li, D. Shamma, M. Bernstein, and L. Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3668–3678, 2015.
- [24] S. Kazemzadeh, V. Ordonez, M. Matten, and T. Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014.
- [25] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907, 2016.
- [26] E. Krahmer and K. Van Deemter. Computational generation of referring expressions: A survey. *Computational Linguistics*, 38(1):173–218, 2012.
- [27] R. Krishna, I. Chami, M. Bernstein, and L. Fei-Fei. Referring relationships. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [28] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. Bernstein, and L. Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *ArXiv e-prints*, 2016.
- [29] D. LaBerge, R. Carlson, J. Williams, and B. Bunney. Shifting attention in visual space: tests of moving-spotlight models versus an activity-distribution model. *Journal of experimental psychology. Human perception and performance*, 23(5):1380–1392, October 1997.
- [30] Y. Li, D. Tarlow, M. Brockschmidt, and R. S. Zemel. Gated graph sequence neural networks. *CoRR*, abs/1511.05493, 2015.

- [31] Y.-L. Li, S. Zhou, X. Huang, L. Xu, Z. Ma, H.-S. Fang, Y. Wang, and C. Lu. Transferable interactiveness knowledge for human-object interaction detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [32] Y. Liang, Y. Bai, W. Zhang, X. Qian, L. Zhu, and T. Mei. Vrr-vg: Refocusing visually-relevant relationships. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [33] W. Liao, M. Y. Yang, H. Ackermann, and B. Rosenhahn. On support relations and semantic scene graphs. *arXiv preprint arXiv:1609.05834*, 2016.
- [34] X. Liu, L. Li, S. Wang, Z.-J. Zha, D. Meng, and Q. Huang. Adaptive reconstruction network for weakly supervised referring expression grounding. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [35] C. Lu, R. Krishna, M. S. Bernstein, and F. Li. Visual relationship detection with language priors. In *European Conf. Comput. Vision*, pages 852–869, 2016.
- [36] J. Mao, J. Huang, A. Toshev, O. Camburu, A. L. Yuille, and K. Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016.
- [37] V. K. Nagaraja, V. I. Morariu, and L. S. Davis. Modeling context between objects for referring expression understanding. In *European Conference on Computer Vision*, pages 792–807. Springer, 2016.
- [38] A. Newell and J. Deng. Pixels to graphs by associative embedding. In *Advances in Neural Information Processing Systems 30 (to appear)*, pages 1172–1180. Curran Associates, Inc., 2017.
- [39] M. Niepert, M. Ahmed, and K. Kutzkov. Learning convolutional neural networks for graphs. In *Int. Conf. Mach. Learning*, pages 2014–2023, 2016.
- [40] B. A. Plummer, A. Mallya, C. M. Cervantes, J. Hockenmaier, and S. Lazebnik. Phrase localization and visual relationship detection with comprehensive image-language cues. In *ICCV*, pages 1946–1955, 2017.
- [41] M. Qi, W. Li, Z. Yang, Y. Wang, and J. Luo. Attentive relational networks for mapping images to scene graphs. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [42] D. Raposo, A. Santoro, D. Barrett, R. Pascanu, T. Lillicrap, and P. Battaglia. Discovering objects and their relations from entangled scene representations. *arXiv preprint arXiv:1702.05068*, 2017.
- [43] S. Ren, K. He, R. B. Girshick, and J. Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *CoRR*, abs/1506.01497, 2015.
- [44] M. A. Sadeghi and A. Farhadi. Recognition using visual phrases. In *Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [45] A. Sanchez-Gonzalez, N. Heess, J. T. Springenberg, J. Merel, M. A. Riedmiller, R. Hadsell, and P. Battaglia. Graph networks as learnable physics engines for inference and control. In *ICML*, pages 4467–4476, 2018.
- [46] A. Santoro, D. Raposo, D. G. Barrett, M. Malinowski, R. Pascanu, P. Battaglia, and T. Lillicrap. A simple neural network module for relational reasoning. In *Neural Inform. Process. Syst.*, pages 4967–4976, 2017.
- [47] B. Schroeder, S. Tripathi, and H. Tang. Triplet-aware scene graph embeddings. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- [48] M. Tanaka, T. Itamochi, K. Narioka, I. Sato, Y. Ushiku, and T. Harada. Generating easy-to-understand referring expressions for target identifications. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [49] X. Wang and A. Gupta. Videos as space-time region graphs. In *ECCV*, 2018.
- [50] D. Xu, Y. Zhu, C. B. Choy, and L. Fei-Fei. Scene Graph Generation by Iterative Message Passing. In *Proc. Conf. Comput. Vision Pattern Recognition*, pages 3097–3106, 2017.
- [51] J. Yang, J. Lu, S. Lee, D. Batra, and D. Parikh. Graph R-CNN for scene graph generation. In *European Conf. Comput. Vision*, pages 690–706, 2018.
- [52] M. Y. Yang, W. Liao, H. Ackermann, and B. Rosenhahn. On support relations and semantic scene graphs. *ISPRS journal of photogrammetry and remote sensing*, 131:15–25, 2017.
- [53] S. Yang, G. Li, and Y. Yu. Dynamic graph attention for referring expression comprehension. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [54] K. Yi, J. Wu, C. Gan, A. Torralba, P. Kohli, and J. Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In *Advances in Neural Information Processing Systems*, pages 1039–1050, 2018.
- [55] L. Yu, P. Poirson, S. Yang, A. C. Berg, and T. L. Berg. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer, 2016.
- [56] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. R. Salakhutdinov, and A. J. Smola. Deep sets. In *Advances in Neural Information Processing Systems 30*, pages 3394–3404. Curran Associates, Inc., 2017.
- [57] R. Zellers, M. Yatskar, S. Thomson, and Y. Choi. Neural motifs: Scene graph parsing with global context. *arXiv preprint arXiv:1711.06640*, abs/1711.06640, 2017.
- [58] J. Zhang, K. J. Shih, A. Tao, B. Catanzaro, and A. Elgammal. An interpretable model for scene graph generation. *CoRR*, abs/1811.09543, 2018.