

# Fully Convolutional Scene Graph Generation

Hengyue Liu<sup>1\*</sup>

Ning Yan<sup>2</sup>

Masood Mortazavi<sup>2</sup>

Bir Bhanu<sup>1</sup>

<sup>1</sup>University of California, Riverside

<sup>2</sup>Futurewei Technologies Inc.

{hliu087, bir.bhanu}@ucr.edu, {nyan, masood.mortazavi}@futurewei.com

## Abstract

This paper presents a fully convolutional scene graph generation (FCSGG) model that detects objects and relations simultaneously. Most of the scene graph generation frameworks use a pre-trained two-stage object detector, like Faster R-CNN, and build scene graphs using bounding box features. Such pipeline usually has a large number of parameters and low inference speed. Unlike these approaches, FCSGG is a conceptually elegant and efficient bottom-up approach that encodes objects as bounding box center points, and relationships as 2D vector fields which are named as Relation Affinity Fields (RAFs). RAFs encode both semantic and spatial features, and explicitly represent the relationship between a pair of objects by the integral on a sub-region that points from subject to object. FCSGG only utilizes visual features and still generates strong results for scene graph generation. Comprehensive experiments on the Visual Genome dataset demonstrate the efficacy, efficiency, and generalizability of the proposed method. FCSGG achieves highly competitive results on recall and zero-shot recall with significantly reduced inference time.

## 1. Introduction

Philosophers, linguists and artists have long wondered about the semantic content of what the mind perceives in images and speech [1, 8, 29, 43]. Many have argued that images carry layers of meaning [2, 45]. Considered as an engineering problem, semantic content has been modeled either as latent representations [11, 17, 25, 38], or explicitly as structured representations [37, 54, 56]. For a computer vision system to explicitly represent and reason about the detailed semantics, Johnson [24] *et al.* adopt and formalize scene graphs from computer graphics community. A scene graph serves as a powerful representation that enables many down-stream high-level reasoning tasks such as image captioning [59, 62], image retrieval [23, 24], Visual Question answering [20, 49] and image generation [23, 57].

\*Work done in part as an intern at Futurewei Technologies Inc.

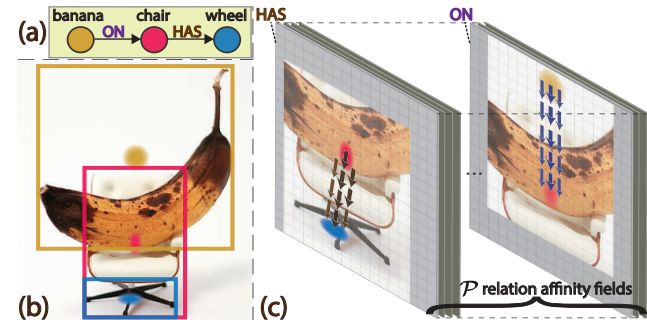


Figure 1: An example of scene graph generation. (a) The ground-truth scene graph of an image. (b) The ground-truth bounding boxes and their centers. (c) Our proposed relationship representation called relation affinity fields. (The image is 2353896.jpg from Visual Genome [27].)

A scene graph is considered as an explicit structural representation for describing the semantics of a visual scene. The nodes in a scene graph represent the object classes and the edges represent the relationships between the objects. Figure 1(a) shows a simple example of a scene graph that represents the underlying semantics of an image. Each relationship between two objects is denoted as a triplet of  $\langle \text{subject}, \text{PREDICATE}, \text{object} \rangle$ , i.e.,  $\text{banana} \xrightarrow{\text{ON}} \text{chair}$  and  $\text{chair} \xrightarrow{\text{HAS}} \text{wheel}$  in Figure 1(a). Most of the SGG work [5, 47, 48, 56, 66] is build as a two-stage pipeline: object detection then scene graph generation. For the first stage of object detection, existing object detectors, i.e., Fast/Faster R-CNN [12, 41], are used for object feature extraction from region proposals. For the second stage of scene graph generation, various approximation methods [48, 56, 66] for graph inference have been used. Some work [5, 63, 64, 65] have also investigated how to utilize external knowledge for improving the results. However, most previous work suffer from not only the long-tailed distribution of relationships [6, 10, 66], but also the highly biased prediction conditioned on object labels [15, 21, 26, 47]. Consequently, frequent predicates will prevail over less frequent ones, and unseen relationships can not be identified. Moreover, the extensibility and inference speed of a SGG

framework is crucial for accelerating down-stream tasks. Although few researchers have studied the efficiency and scalability in SGG [13, 31, 58], the high computational complexity impedes the practicality towards real-world applications. A natural question that arises is: *can we solve scene graph generation in a per-pixel prediction fashion?* Recently, anchor-free object detectors [28, 50, 61, 68] have become popular due to their simplicity and low cost. These methods treat an object as a single or many, pre-defined or self-learned keypoints. Relating object detection to human pose estimation, if an object can be modeled as a point (human “keypoint”), is it possible to represent a binary relationship as vectors (human “limb”)?

In this paper, we propose a novel fully convolutional scene graph generation model, *i.e.*, FCSGG, with state-of-the-art object detection results on Visual Genome dataset [27], as well as compelling SGG results compared with visual-only methods. We present a bottom-up representation of objects and relationships by modeling objects as points and relationships as vectors. Each relationship is encoded as a segment in a 2D vector field called *relation affinity field* (RAF). Figure 1(c) shows an illustration of RAFs for predicates ON and HAS. Both objects and relationships are predicted as dense feature maps without losing spatial information. For the first time, scene graphs can be generated from a single convolutional neural network (CNN) with significantly reduced model size and inference speed. Specifically, we make the following contributions:

- We propose the first fully convolutional scene graph generation model that is more compact and computationally efficient compared to previous SGG models.
- We introduce a novel relationship representation called *relation affinity fields* that generalizes well on unseen visual relationships. FCSGG achieves strong results on zero-shot recall.
- Our proposed model outperforms most of the visual-only SGG methods, and achieves competitive results compared to methods boosted by external knowledge.
- We conduct comprehensive experiments and benchmark our proposed method together with several previous work on model efficiency, and FCSGG achieves near real-time inference.

## 2. Related Work

We categorize the related work of SGG into the following directions: refinement of contextual feature, adaptation of external knowledge, and others.

*Contextual feature refinement.* Xu *et al.* [56] proposed an iterative message passing mechanism based on Gated Recurrent Units [7], where the hidden states are used for predictions. Followers [48, 66] studied better recurrent neural

networks [19, 42, 46] for encoding object and edge context. Others trying to incorporate more spatial features into SGG. Li *et al.* [32] proposed the MSDN that merges features across multiple semantic levels, and later achieved message passing constrained on visual phrase [30]. Dai *et al.* [9] proposed a spatial module by learning from bounding box masks. Woo *et al.* [53] introduced the geometric embeddings by directly encoding the bounding box offsets between objects. Wang *et al.* [52] further studied the effects of relative positions between objects for extracting more discriminating features. Our method is fundamentally different from these methods as the relationships are grounded semantically and spatially directly into CNN features. Without any explicit iterative information exchange between nodes and edges, our model is able to predict objects and relationships in a single forward pass.

*External knowledge adaptation.* Beyond visual features, linguistic knowledge can serve as additional features for SGG [14, 37, 40, 63]. By adopting statistical correlations of objects, Chen *et al.* [5] utilized graph neural networks [44] to infer relationships. Gu *et al.* [16] and Zareian *et al.* [64, 65] explored the usefulness of knowledge or commonsense graphs for SGG. Tang *et al.* [47] proposed an de-biasing method by causal interventions of predictions. Lin *et al.* [36] investigated the graph properties and mitigated the long-tailed distributions of relationships. Compared with these methods, our proposed model relies only on visual features but still yields a strong performance.

Very few researchers have investigated alternatives either for object feature or relationship feature representations. Newell [39] and Zhang *et al.* [67] tried latent-space embeddings for relationship and achieved improvements. Different from most of the previous work, FCSGG reformulates and generalizes relationship representations from only visual-based features in near real-time, which is much faster than specifically designed SGG models for efficiency [31, 58].

## 3. Object Detection as Keypoint Estimation

In this section, we provide the preliminaries of modeling object detection as keypoint estimation in a single-scale dense feature map prediction fashion.

Our model is built upon a one-stage anchor-free detector, namely CenterNet [68]. Different from commonly used anchor-based R-CNN approaches for generating object proposals and features, it predicts three dense features that represent centers of object bounding boxes, center offsets, and object sizes. More specifically, an input image  $I \in \mathbb{R}^{3 \times W \times H}$  will go through a backbone CNN generating feature maps of size  $w \times h = \lfloor \frac{W}{\tau} \rfloor \times \lfloor \frac{H}{\tau} \rfloor$  where  $\tau$  is the total stride until the last layer, and we set  $\tau = 4$  unless specified otherwise. Then these features will be fed into three prediction heads, each of which consists of several convo-



Our model is conceptually simple: in addition to the outputs that are produced by the object detection network described in Section 3, we add another branch that outputs a novel feature representation for relationships called *relation affinity fields* (RAFTs). Specifically, the RAFTs are a set of 2D vector fields  $\mathbf{F} = \{\mathbf{F}_p\} \in \mathbb{R}^{\mathcal{P} \times 2 \times h \times w}$ , where  $p \in \mathbb{R}^{\mathcal{P}}$  and  $\mathcal{P}$  is the number of predicate classes in a dataset. Each 2D vector field  $\mathbf{F}_p$  represents the relationships among all the object pairs of predicate  $p$ . Given our definition of objects as center points, the ground-truth RAFTs are defined as vectors flow from the center of subject to the center of object. More formally, we define the binary relationships among objects  $\mathbf{B}$  in the input image as  $\mathbf{R} = \{r^{i \rightarrow j}\}$ , where  $r^{i \rightarrow j} = (b^i, p^{i \rightarrow j}, b^j)$  is the relationship triplet from subject  $b^i$  to object  $b^j$  with predicate  $p^{i \rightarrow j}$ . We define a “path”  $\pi_p^{i \rightarrow j}$  that “propagates”  $p^{i \rightarrow j}$  from subject center  $\mathbf{o}^i$  to object center  $\mathbf{o}^j$ . For a point  $\mathbf{p} = (x, y)$ , its ground-truth relation affinity field vector  $\mathbf{F}_{p,x,y}$  is given as

$$\mathbf{F}_{p,x,y} = \begin{cases} \mathbf{e}^{i \rightarrow j} = \frac{\mathbf{o}^j - \mathbf{o}^i}{\|\mathbf{o}^j - \mathbf{o}^i\|_2} & \text{if } \mathbf{p} \in \pi_p^{i \rightarrow j} \\ \mathbf{0} & \text{otherwise,} \end{cases} \quad (3)$$

and the path  $\pi_p^{i \rightarrow j}$  is defined on a set of points between object centers forming a rectangular region:

$$\pi_p^{i \rightarrow j} = \{\mathbf{p} \mid 0 \leq \mathbf{e}^{i \rightarrow j} \cdot (\mathbf{p} - \mathbf{o}^i) \leq \epsilon_{\mathbf{e}^{i \rightarrow j}} \text{ and } |\mathbf{e}_{\perp}^{i \rightarrow j} \cdot (\mathbf{p} - \mathbf{o}^i)| \leq \epsilon_{\mathbf{e}_{\perp}^{i \rightarrow j}}\}, \quad (4)$$

where  $\epsilon_{\mathbf{e}^{i \rightarrow j}} = \|\mathbf{o}^j - \mathbf{o}^i\|_2$  as the relationship “length” along the direction  $\mathbf{e}^{i \rightarrow j}$ , and  $\epsilon_{\mathbf{e}_{\perp}^{i \rightarrow j}} = \min(a^i, b^i, a^j, b^j)$  as the relationship “semi-width” along  $\mathbf{e}_{\perp}^{i \rightarrow j}$  (orthogonal to  $\mathbf{e}^{i \rightarrow j}$ ) being the minimum of object centers’ radii. Since vectors may overlap at the same point, the ground-truth RAFT  $\mathbf{F}_p$  averages the fields computed for all the relationship triplets containing that particular predicate  $p$ . It is given as  $\mathbf{F}_p = \frac{1}{n_c(x,y)} \sum_{x,y} \mathbf{F}_{p,x,y}$ , where  $n_c(x,y)$  is the number of non-zero vectors at point  $(x,y)$ . With the definition of ground-truth RAFTs, we can train our network to regress such dense feature maps. The RAFT regression loss  $\mathcal{L}_{raf}$  can be estimated using a normal regression losses  $\mathcal{L}_{reg}$  such as L1, L2 or smooth L1 [12]. Given the predicted RAFTs  $\hat{\mathbf{F}}$ , the loss is defined as per-pixel weighted regression loss as

$$\mathcal{L}_{raf} = \mathbf{W} \cdot \mathcal{L}_{reg}(\hat{\mathbf{F}}, \mathbf{F}), \quad (5)$$

where  $\mathbf{W}$  is a pixel-wise weight tensor of the same shape of  $\mathbf{F}$ . The weights  $\mathbf{W}$  are determined and divided into three cases (Figure 3): a)  $\mathbf{W}_{p,x,y} = 1$  if  $(x,y)$  is exactly on the line segment between objects having the relationship  $p$ , b)  $\mathbf{W}_{p,x,y} \in (0, 1)$  if the distance between  $(x,y)$  and the line segment is small and the value is negative correlated to the distance and c) otherwise where  $\mathbf{F}_{p,x,y} = \mathbf{0}$ . We provide ablation study on the choice of losses and the weight




Figure 3: An example of GT relation affinity field of predicate LAYING ON based on equations 3 and 4. A non-zero unit vector is only defined on locations inside  $\pi_{\text{LAYING ON}}^{\text{cat} \rightarrow \text{table}}$ .

tensor in Section 5.3. Finally, the complete loss for training our proposed model can be written as  $\mathcal{L} = \mathcal{L}_{det} + \mathcal{L}_{raf}$ .

Our proposed RAFTs encode rich information of both spatial and semantic features as dense feature maps, and enable end-to-end joint training of object detection and relation detection. To extract relationship from predictions, path integral over RAFTs is performed which will be described below.

#### 4.1. Inference

We compute path integrals over RAFTs along the line segments connecting pairs of detected object centers as the scores of relationships. Specifically, for two candidate object centers  $\hat{\mathbf{o}}^i$  and  $\hat{\mathbf{o}}^j$  with predicted class scores  $\hat{h}^i$  and  $\hat{h}^j$ , we gather the predicted RAFTs  $\hat{\mathbf{F}}$  along the path between  $\hat{\mathbf{o}}^i$  and  $\hat{\mathbf{o}}^j$ , and compute the mean of their projections onto  $\hat{\mathbf{e}}^{i \rightarrow j} = (\hat{\mathbf{o}}^j - \hat{\mathbf{o}}^i) / \|\hat{\mathbf{o}}^j - \hat{\mathbf{o}}^i\|_2$ . The path integral scores  $\mathbf{E}^{i \rightarrow j}$  are identified as the confidences of existence of relationships:

$$\mathbf{E}^{i \rightarrow j} = \frac{\hat{h}^i \cdot \hat{h}^j}{m^{i \rightarrow j}} \sum_{p \in \mathbb{R}^{\mathcal{P}}} \sum_{(x,y) \in \pi^{i \rightarrow j}} \mathbf{F}_{p,x,y} \cdot \hat{\mathbf{e}}^{i \rightarrow j}, \quad (6)$$

where  $m^{i \rightarrow j} = |\pi^{i \rightarrow j}|$  is the number of points in  $\pi^{i \rightarrow j}$ . Since our RAFTs are object-class-agnostic, we multiply the class scores of the objects and the path integral score as the overall classification score for the relationship predicate. The integral will be performed spatially for each predicate channel, so  $\mathbf{E}^{i \rightarrow j}$  represents the confidences of predicted relationship triplet  $\hat{r}^{i \rightarrow j}$  for all predicates. Note that the integral could be negative that indicates an opposite relationship of the object pairs, and those negative integral values can be simply negated as  $\mathbf{E}^{j \rightarrow i} = -\mathbf{E}^{i \rightarrow j}$ . Finally, both scores of  $\mathbf{E}^{i \rightarrow j}$  and  $\mathbf{E}^{j \rightarrow i}$  will be used for ranking the predicted relationships. We also experiment with a simple re-weighting step known as frequency bias [66] by multiplying  $\mathbf{E}^{i \rightarrow j}$  with  $1.0001^{n_c(\hat{r}^{i \rightarrow j})}$ , where  $n_c(\hat{r}^{i \rightarrow j})$  counts the occurrence of triplet  $r^{i \rightarrow j}$  in training set. The path integral



---

**Algorithm 1** Path Integral over Relation Affinity Fields

---

**Input:** Object centers  $\{\hat{o}^i\}$ , Relation Affinity Fields  $\hat{\mathbf{F}}$

**Output:** Relations  $\hat{\mathbf{R}} = \{\hat{r}^{i \rightarrow j} = (i, \mathbf{E}^{i \rightarrow j}, j) | i \neq j\}$

```
1:  $\hat{\mathbf{R}} \leftarrow \{\}$ 
2: for each center pair  $(\hat{o}^i, \hat{o}^j) | i \neq j$  do
3:    $\pi^{i \rightarrow j} \leftarrow \text{Linspace}(\hat{o}^i, \hat{o}^j) \in \mathbb{R}^{m^{i \rightarrow j} \times 2}$ 
4:    $\hat{\mathbf{F}}_{\pi^{i \rightarrow j}} \leftarrow \text{INDEX}(\hat{\mathbf{F}}, \pi^{i \rightarrow j}) \in \mathbb{R}^{\mathcal{P} \times m^{i \rightarrow j} \times 2}$ 
5:    $\mathbf{E}^{i \rightarrow j}, \mathbf{E}^{j \rightarrow i} \in \mathbb{R}^{\mathcal{P}} \leftarrow \text{Equation 6}$ 
6:    $\hat{\mathbf{R}} \leftarrow \hat{\mathbf{R}} \cup \{\hat{r}^{i \rightarrow j} = (i, \mathbf{E}^{i \rightarrow j}, j)\}$ 
7:    $\hat{\mathbf{R}} \leftarrow \hat{\mathbf{R}} \cup \{\hat{r}^{j \rightarrow i} = (j, \mathbf{E}^{j \rightarrow i}, i)\}$ 
8: end for
9: return SORTED( $\hat{\mathbf{R}}$ )
```

---

procedure is presented in Algorithm 1. In practice, the operations are performed using matrix multiplication in stead of FOR LOOP for fast inference.

## 4.2. Multi-scale Prediction

Since object centers are downsampled to feature level, their centers could collide onto the same pixel location. Regression ambiguity may rise due to single-scale feature representations. In this subsection, we address this problem by utilizing multi-scale prediction and shared detection heads.

Though Zhou *et al.* [68] argued that only a very small fraction ( $<0.1\%$  in COCO [35] dataset) of objects have center collision problem at stride of 4, the size and offset regression targets need better assignment strategy since there is only one valid target per pixel. We follow the work of FPN [33], RetinaNet [34] and FCOS [50], and assign the ground-truth bounding boxes to different levels based on scales. Building upon the backbone features, we construct multi-level feature maps  $\{P_k\}$  where  $P_k$  is of stride  $2^k$ . We refer the network component of generating multi-scale features as the “neck” (the green box in Figure 2), such as FPN [33]. We define a valid range  $[l_k, u_k] \subset [0, L_{max}]$  for objects in each scale, where  $L_{max}$  is the maximum size of longer edge allowed for training and testing. Only bounding boxes of area within  $[l_k^2, u_k^2]$  are qualified for the  $k$ -th scale training. We experiment different number of scale levels and input image size. For smaller input image of  $L_{max} = 512$ , we build 4-scale features [33]  $\{P_2, P_3, P_4, P_5\}$  (as shown in Figure 2) with valid ranges  $\{[0, 32^2], [32^2, 64^2], [64^2, 128^2], [128^2, 512^2]\}$ ; for larger input image of  $L_{max} = 1024$  (shorter edge is at least  $L_{min} = 640$ ), we use 5-scale features [34, 50]  $\{P_3, P_4, P_5, P_6, P_7\}$  with area ranges  $\{[0, 64^2], [64^2, 128^2], [128^2, 256^2], [256^2, 512^2], [512^2, 1024^2]\}$ . If there is still more than one target at the same location, we simply choose the smallest object for regression.

In terms of the multi-scale RAFs training, the GT assignment is based on the distances between object centers. For high-level semantic features like  $P_5$ , the feature map

can capture large objects, as well as relationships among distant objects. We select the relationships of “length”  $\epsilon_{e^{i \rightarrow j}} \in [l_k, u_k]$  as valid samples for training the  $k$ -th scale. The exact ranges for 4-scale or 5-scale RAFs are the same as the settings for bounding boxes. Finally, the weights of detection heads are shared across different feature scales for efficiencies and performance improvements. During inference, we gather outputs from each scale based on the corresponding valid range, then merge and rank all relationship triplets. Figure 2 illustrates the details of our proposed architecture using a four-scale feature setup as an example with shared detection heads. Our experiments (Section 5) show that the multi-scale GT and scale-aware training resolve the aforementioned ambiguity problem thus improve the results over single-scale prediction.

## 5. Experiments

*Dataset.* We use the Visual Genome (VG) [27] dataset to train and evaluate our models. We followed the widely-used preprocessed subset of VG-150 [56] which contains the most frequent 150 object categories ( $\mathcal{C} = 150$ ) and 50 predicate categories ( $\mathcal{P} = 50$ ). The dataset contains approximately 108k images, with 70% for training and 30% for testing. Different from previous works [4, 48, 66], we do not filter non-overlapping triplets for evaluation.

*General settings.* We experiment on two settings, one for small input size ( $L_{max} = 512$ ) and one for larger size ( $L_{max} = 1024$ ). The model is trained end-to-end using SGD optimizer with the batch size of 16 for 120k iterations. The initial learning rate is set to 0.02 and decayed by the factor of 10 at 80k<sup>th</sup> and 100k<sup>th</sup> iteration. We adopt standard image augmentations of horizontal flip and random crop with multi-scale training. During testing, we keep the top 100 detected objects for path integral.

*Metrics.* We conduct comprehensive analysis following three standard evaluation tasks: Predicate Classification (PredCls), Scene Graph Classification (SGCls), and Scene Graph Detection (SGDet). We report results of recall@K (R@K) [37], no-graph constraint recall@K (ng-R@K) [39, 66], mean recall@K (mR@K) [5, 48], no-graph constraint mean recall@K (ng-mR@K), zero-shot recall@K (zsR@K) [37] and no-graph constraint zero-shot recall@K (ng-zsR@K) [47] for all three evaluation tasks. We do not train separate models for different tasks.

### 5.1. Implementation Details

We conduct experiments on different backbone and neck networks. Each of the detection heads consists of four  $3 \times 3$  convolutions followed by batch normalization and ReLU, and one  $1 \times 1$  convolution with the desired number of output channels in all our experiments unless specified otherwise. For convenience, our models are named as BACKBONE - # OF OUTPUT SCALES - NECK - OTHER OPTIONS.

Recall @K / No-graph Constraint Recall @K		AP <sub>50</sub>	Predicate Classification		Scene Graph Classification		Scene Graph Detection	
			R@20/50/100	ng-R@20/50/100	R@20/50/100	ng-R@20/50/100	R@20/50/100	ng-R@20/50/100
External Knowledge	VCtree [48]	-	60.1/ 66.4/ 68.1	-	35.2/ 38.1/ 38.8	-	22.0/ 27.9/ 31.3	-
	KERN [5]	-	-/ 65.8/ 67.6	-/ 81.9/ 88.9	-/ 36.7/ 37.4	-/ 45.9/ 49.0	-/ 27.1/ 29.8	-/ 30.9/ 35.8
	GPS-Net [36]	-	67.6/ 69.7/ 69.7	-	41.8/ 42.3/ 42.3	-	22.3/ 28.9/ 33.2	-
	MOTIFS-TDE [47, 66]	28.1	33.6/ 46.2/ 51.4	-	21.7/ 27.7/ 29.9	-	12.4/ 16.9/ 20.3	-
	GB-Net- $\beta$ [64]	-	-/ 66.6/ 68.2	-/ 83.5/ 90.3	-/ 37.3/ 38.0	-/ 46.9/ 50.3	-/ 26.3/ 29.9	-/ 29.3/ 35.0
Visual Only	VTransE* [67]	-	-	-	-	-	-/ 5.5/ 6.0	-
	FactorizableNet* [31]	-	-	-	-	-	-/ 13.1/ 16.5	-
	IMP† [56, 66]	20.0	58.5/ 65.2/ 67.1	-	31.7/ 34.6/ 35.4	-	14.6/ 20.7/ 24.5	-
	Pixels2Graphs [39]	-	-	-/ 68.0/ 75.2	-	-/ 26.5/ 30.0	-	-/ 9.7/ 11.3
	Graph R-CNN [58]	23.0	-/ 54.2/ 59.1	-	-/ 29.6/ 31.6	-	-/ 11.4/ 13.7	-
	VRF [10]	-	-/ 56.7/ 57.2	-	-/ 23.7/ 24.7	-	-/ 13.2/ 13.5	-
FCSGG (Ours)	CISC [52]	-	42.1/ 53.2/ 57.9	-	23.3/ 27.8/ 29.5	-	7.7/ 11.4/ 13.9	-
	HRNetW32-1S	21.6	27.6/ 34.9/ 38.5	32.2/ 46.3/ 56.6	12.3/ 15.5/ 17.2	13.5/ 19.3/ 23.6	11.0/ 15.1/ 18.1	12.4/ 18.2/ 23.0
	HRNetW48-1S	25.0	24.2/ 31.0/ 34.6	28.1/ 40.3/ 50.0	13.6/ 17.1/ 18.8	14.2/ 19.6/ 24.0	11.5/ 15.5/ 18.4	12.7/ 18.3/ 23.0
	ResNet50-4S-FPN $\times 2$	23.0	28.0/ 35.8/ 40.2	31.6/ 44.7/ 54.8	13.9/ 17.7/ 19.6	14.8/ 20.6/ 25.0	11.4/ 15.7/ 19.0	12.2/ 18.0/ 22.8
	HRNetW48-5S-FPN $\times 2$	28.5	28.9/ 37.1/ 41.3	34.0/ 48.1/ 58.4	16.9/ 21.4/ 23.6	18.6/ 26.1/ 31.6	13.5/ 18.4/ 22.0	15.4/ 22.5/ 28.3
	HRNetW48-5S-FPN $\times 2$ -f	28.5	33.4/ 41.0/ 45.0	37.2/ 50.0/ 59.2	19.0/ 23.5/ 25.7	19.6/ 26.8/ 32.1	16.1/ 21.3/ 25.1	16.7/ 23.5/ 29.2

Table 1: Recall and no-graph constraint recall @K evaluation results on VG-150.  $\star$  denotes the methods evaluated on other datasets, such that VTransE is evaluated on VG-200 [67] and FactorizableNet on a smaller set following [32].  $\dagger$  denotes the methods with updated re-implementation results. - denotes the results that are not reported in the corresponding work.

Mean Recall @K / Ng Mean Recall @K		Predicate Classification		Scene Graph Classification		Scene Graph Detection	
		mR@20/50/100	ng-mR@20/50/100	mR@20/50/100	ng-mR@20/50/100	mR@20/50/100	ng-mR@20/50/100
VCtree [48]		14.0/ 17.9/ 19.4	-	8.2/ 10.1/ 11.8	-	5.2/ 6.9/ 8.0	-
KERN [5]		-/ 17.7/ 19.4	-	-/ 9.4/ 10.0	-	-/ 6.4/ 7.3	-
GPS-Net [36]		-/ -/ 22.8	-	-/ -/ 12.6	-	-/ -/ 9.8	-
MOTIFS-TDE [47, 66]		18.5/ 25.5/ 29.1	-	9.8/ 13.1/ 14.9	-	5.8/ 8.2/ 9.8	-
GB-Net- $\beta$ [64]		-/ 22.1/ 24.0	-	-/ 12.7/ 13.4	-	-/ 7.1/ 8.5	-
HRNetW32-1S		4.0/ 5.5/ 6.3	5.4/ 9.7/ 13.6	1.9/ 2.5/ 2.8	2.7/ 4.4/ 6.2	1.7/ 2.4/ 2.9	2.2/ 3.6/ 4.9
HRNetW48-1S		3.7/ 5.2/ 6.1	5.2/ 9.5/ 14.7	2.2/ 2.9/ 3.4	3.5/ 6.3/ 9.4	1.8/ 2.6/ 3.1	2.7/ 4.7/ 6.9
ResNet50-4S-FPN $\times 2$		4.2/ 5.7/ 6.7	6.5/ 11.3/ 16.6	2.2/ 2.9/ 3.3	3.6/ 6.0/ 8.3	1.9/ 2.7/ 3.3	3.0/ 4.9/ 6.8
HRNetW48-5S-FPN $\times 2$		4.3/ 5.8/ 6.7	6.1/ 10.3/ 14.2	2.6/ 3.4/ 3.8	4.1/ 6.4/ 8.4	2.3/ 3.2/ 3.8	3.7/ 5.7/ 7.4
HRNetW48-5S-FPN $\times 2$ -f		4.9/ 6.3/ 7.1	6.6/ 10.5/ 14.3	2.9/ 3.7/ 4.1	4.2/ 6.5/ 8.6	2.7/ 3.6/ 4.2	3.8/ 5.7/ 7.5

Table 2: The SGG results on mean recall@K and no-graph constraint mean recall@K.

*ResNet* [18, 33]. We start by using ResNet-50 as our backbone and build a 4-scale FPN for multi-scale prediction. Since the tasks of object detection and RAFs prediction are jointly trained, the losses from the two tasks could compete with each other. We implement a neck named “FPN $\times 2$ ” with two parallel FPNs, such that one FPN is used for constructing features for object detection heads (center, size and offset), and the other is for producing features for RAFs. We name this model as ResNet50-4S-FPN $\times 2$ .

*HRNet* [51]. We then experiment on a recent proposed backbone network called HRNet that consists of parallel convolution branches with information exchange across different scales. For single-scale experiments, we use HRNetV2-W32 and HRNetV2-W48; and for multi-scale prediction, we adopt its pyramid version called HRNetV2p. We omit their version number for the rest of the paper. We test several models: HRNetW32-1S, HRNetW48-1S and HRNetW48-5S-FPN $\times 2$ . We also experiment on adding frequency bias for inference as discussed in Section 4.1, and the model used is called HRNetW48-5S-FPN $\times 2$ -f.

## 5.2. Quantitative Results

We first compare results of R@K and ng-R@K with various of state-of-the-art (SOTA) models, and divide them into two categories: 1) models that only use visual features derived from the input image like our proposed model and 2) models that not only use visual features, but also use features like language embeddings, dataset statistics or counterfactual causality, *etc.* The results are shown in Table 1. Our best model achieves 28.5 average precision at IoU = 0.5 (AP<sub>50</sub>) for object detection. Though our SGG results do not outperform the SOTA approaches, we achieve the best scene graph detection results among visual-only models. Specifically, Pixels2Graphs [39] and our models are the only models without using Faster R-CNN [41] as object detector, and we achieves 13.8 / 17.9 gain on SGGDet ng-R@50 / 100 compared with their results using RPN [41]. We then report mean recall and no-graph constraint mean recall results shown in Table 2. We still obtain competitive results, especially on ng-mR. By comparing (m) R@K and ng-(m) R@K directly, we observe more gain on our results

Zero-shot Recall @K	PredCls		SGCls		SGDet	
Method	zsR@50/100		zsR@50/100		zsR@50/100	
MOTIFS-TDE [47]	14.4 / 18.2		3.4 / 4.5		2.3 / 2.9	
VTransE-TDE [47]	13.3 / 17.6		2.9 / 3.8		2.0 / 2.7	
VCTree-TDE [47]	14.3 / 17.6		3.2 / 4.0		2.6 / 3.2	
Knyazev <i>et al.</i> [26]	- / 21.5		- / 4.2		- / -	
FCSGG (Ours)	zsR@50/100	ng-zsR@50/100	zsR@50/100	ng-zsR@50/100	zsR@50/100	ng-zsR@50/100
HRNetW32-1S	8.3 / 10.7	12.9 / 19.2	1.0 / 1.2	2.3 / 3.5	0.6 / 1.0	1.2 / 1.6
HRNetW48-1S	8.6 / 10.9	12.8 / 19.6	1.7 / 2.1	2.9 / 4.4	1.0 / 1.4	1.8 / 2.7
ResNet50-4S-FPN $\times 2$	8.2 / 10.6	11.7 / 18.1	1.3 / 1.7	2.4 / 3.8	0.8 / 1.1	1.0 / 1.7
HRNetW48-5S-FPN $\times 2$	7.9 / 10.1	11.5 / 17.7	1.7 / 2.1	2.8 / 4.8	0.9 / 1.4	1.4 / 2.4
HRNetW48-5S-FPN $\times 2-f$	7.8 / 10.0	11.4 / 17.6	1.6 / 2.0	2.8 / 4.8	0.8 / 1.4	1.4 / 2.3

Table 3: Comparisons of SGG results on zero-shot Recall@K, and our results on no-graph constraint zero-shot Recall@K.

than other methods’.

Zero-shot recall (zsR) [15, 21, 26, 37, 40, 47, 60] is a proper metric for evaluating the model’s robustness and generalizability for generating scene graphs. It computes recall on those subject-predicate-object triplets that do not present during training. There are in total of 5971 unique zero-shot triplets from the testing set of VG-150. The results and comparisons are listed in Table 3. We also compute the per-predicate recall@100 for predicate classification task using HRNetW48-5S-FPN $\times 2-f$ , and show the comparisons in Figure 4. We observe similar behavior with our results on recall, such that the results on no-graph constraint zero-shot recall are significantly better than zero-shot recall. Even for unseen triplets, purely based on visual features, FCSGG is still capable of predicting meaningful RAFs which proves its generalization capability. In other words, when constructing scene graphs from RAFs, our approach does not highly depend on the object classes but only focuses on the context features between the objects. When comparing with other reported results on zsR, we achieve slightly lower results than those much larger models. For example of PredCls task, ResNet50-4S-FPN $\times 2$  achieves 10.6 zsR@100 with only 36 million (M) number of parameters and inference time of 40 milliseconds (ms) per image, while VCTree-TDE [47] achieves 17.6 zsR@100 with 360.8M number of parameters and inference time of 1.69 s per image. For comparison, ResNet50-4S-FPN $\times 2$  is 10 times smaller and 42 times faster than VCTree-TDE.

*Limitations.* It should be noticed that FCSGG also has some “disadvantages” over Faster R-CNN-based methods on easier tasks such as PredCls and SGCls. During evaluations with given GT bounding boxes or classes, our RAFs features will not change, while R-CNN extracted object/union-box features will change which leads to better results. When using visual-only representation of relationships, it is hard for the network to distinguish predicates between WEARS / WEARING (by comparing R and ng-R in Figure 4) or LAYING ON / LYING ON, which is common in VG dataset. In this sense, incorporating exter-




Figure 4: FCSGG per-predicate PredCls@100 results for selected predicates using HRNetW48-5S-FPN $\times 2-f$ .

nal knowledge gives FCSGG a large potential in improving results. Comparing the model HRNetW48-5S-FPN $\times 2$  and its frequency-biased counterpart HRNetW48-5S-FPN $\times 2-f$ , we find noticeable improvement by using training set statistics. This simple cost-free operation can improve R@20 by 2.6, and we expect better results from fine-tuning hyperparameters. However, the focus of this work is not perfectly fitting on a dataset, but improving generalization of relationship based on visual features. More sophisticated ensemble methods or extensions are beyond the scope of this paper.

### 5.3. Ablation Study

#### 5.3.1 RAF regression Loss

We experience the same difficulty of training from sparsely annotated scene graphs as discussed by Newell *et al.* [39]. The network has the potential of generating reasonable triplets not covered in the ground-truth, and our results on zero-shot recall prove the argument. To reduce the penalty on these detections, we investigate the design methodology of RAF regression loss  $\mathcal{L}_{raf}$  (Equation 5).

We refer the loss applied at locations having GT RAFs defined as positive loss  $\mathcal{L}_{raf}^+$ , and we test different regression losses. As for locations where  $\mathbf{F}_{p,x,y} = \mathbf{0}$ , we apply so called negative loss  $\mathcal{L}_{raf}^-$  using L1 for regression.  $\mathcal{L}_{raf}^-$

$\mathcal{L}_{raf}^+$	$\beta$	AP <sub>50</sub>	R@50	zR@50	mR@50
L1	0	21.57	6.22	0.40	2.28
L1	1	21.52	9.80	0.56	2.33
L1	10	21.56	15.05	0.60	2.36
Smooth L1	0	20.15	5.00	0.30	2.51
Smooth L1	1	19.65	7.46	0.57	2.45
Smooth L1	10	20.63	11.82	0.61	2.83
L2	0	19.62	4.82	0.26	2.34
L2	1	21.60	10.76	0.68	2.57
L2	10	21.62	2.89	0.57	2.50

Table 4: Ablations on losses used for positive samples and regularization factor on negative samples of RAFs. AP<sub>50</sub> and SGDet results are reported using HRNetW32-1S.

will be multiplied by a factor  $\beta$  for adjusting the penalty. Spatially,  $\mathcal{L}_{raf}$  can be re-written as  $\mathcal{L}_{raf} = \mathcal{L}_{raf}^+ + \beta \mathcal{L}_{raf}^-$ . Table 4 shows the effects of different losses and penalty factor on the performance. We observe better performance using L1 and  $\beta = 10$ . However, when only supervise on  $\mathcal{L}_{raf}^+$  loss, the model has comparable mean recall results and it can detect more semantic and rare relationships. On the other hand, adding  $\mathcal{L}_{raf}^-$  loss will push the model more biased to dominating predicates like ON and HAS.

### 5.3.2 Architecture Choices

Neck	Norm	AP <sub>50</sub>	R@50	zR@50	mR@50
FPN	GN	22.75	11.29	0.71	2.95
FPN	MS-BN	22.10	13.23	0.75	2.67
FPN <sub>×2</sub>	GN	22.74	11.96	0.78	3.00
FPN <sub>×2</sub>	MS-BN	22.60	12.01	0.80	2.88

Table 5: Comparisons of FPN vs. FPN<sub>×2</sub>, and Multi-scale batch normalization vs. group normalization. AP<sub>50</sub> and SgDet results are reported using ResNet50-4S.

By comparing our results between single-scale and multi-scale models, we see substantial performance gain on both object detection and scene graph generation from Table 1 2 3. We also observe HRNet has better results over ResNet due to its multi-scale feature fusions. For investigating the entanglement of object features and contextual features producing RAFs, we compare the results of FPN and FPN<sub>×2</sub> using ResNet-50 as backbone shown in Table 5. As observed in [50], the regression range differs across different levels. Therefore, to improve the performance of shared fully-convolutional heads, we replace each batch normalization (BN) [22] layer in the head with a set of BN layers, each of which is only applied for the corresponding scale. We name this modified BN as multi-scale batch normalization (MS-BN). We also experiment on group normalization (GN) [55] which stabilizes the training as well. We show the comparisons of MS-BN and GN (Section 4.2) in the same table. We observe better mR by using FPN<sub>×2</sub> and better recall and zSR by using MS-BN. We expect more performance improvement if using larger batch size with MS-BN.

### 5.3.3 Model Size and Speed

We also conduct experiments on the model size and inference speed. Few work benchmarked on efficiency of scene graph generation previously [31, 64]. Though scene graphs are powerful, it is almost not possible to perform SGG and down-stream tasks in real-time due to significantly increased model complexity. FCSGG alleviates the computational complexity effectively. Our experiments are performed on a same NVIDIA GeForce GTX 1080 Ti GPU with inference batch size of 1. For comparisons,

Method	#Params (M)	Input Size	s / image
Pixels2Graphs [67]	94.8	512 × 512	3.55
VCtree-TDE [47]	360.8	600 × 1000	1.69
MOTIFS-TDE [47]	369.5	600 × 1000	0.87
KERN [5]	405.2	592 × 592	0.79
MOTIFS [47]	367.2	600 × 1000	0.66
FactorizableNet [31]	40.4	600 × 1000	0.59
VTransE-TDE [47]	311.6	600 × 1000	0.55
GB-NET- $\beta$ [64]	444.6	592 × 592	0.52
Graph R-CNN [58]	80.2	800 × 1024	0.19
<b>FCSGG (Ours)</b>			
HRNetW32-1S	47.3	512 × 512	0.07
HRNetW48-1S	86.1	512 × 512	0.08
ResNet50-4S-FPN <sub>×2</sub>	36.0	512 × 512	0.04
HRNetW48-5S-FPN <sub>×2</sub>	87.1	640 × 1024	0.12
HRNetW48-5S-FPN <sub>×2-f</sub>	87.1	640 × 1024	0.12

Table 6: Model size and speed comparisons for SgDet.

we include several previous work by running corresponding open-source codes under the same settings. The results are shown in Table 6. Both the number of parameters and inference time are considerably lower for FCSGG models. It is worth noting that the computation overhead is from the backbone network. The path integral (Algorithm. 1) is performed pair-wisely for all 100 kept objects across five scales, which results in  $\binom{100}{2} \times 2 \times 5 = 49500$  maximum number of candidate relationships for an image. The inference time for path integral is almost invariant over the number of instances as analyzed by Cao *et al.* [3]. We believe that object relationships exist universally, especially geometric ones. By grounding the full graph in RAFs as intermediate features, richer semantics can be retained for down-stream tasks. More importantly, convolution is hardware-friendly, and the model size is kept small for deployment on edge devices. We anticipate that real-time mobile SGG can be performed in the near future.

## 6. Conclusions

Scene graph generation is a critical pillar for building machines to visually understand scenes and perform high-level vision and language tasks. In this paper, we introduce a fully convolutional scene graph generation framework that is simple yet effective with fast inference speed. The proposed relation affinity fields serve as a novel representation for visual relationship and produce strong generalizability for unseen relationships. By only using visual features, our exploratory method achieves competitive results over object detection and SGG metrics on the VG dataset. We expect that FCSGG can serve as a general and strong baseline for SGG task, as well as a vital building block extending to down-stream tasks.

## Acknowledgement

This work was partially supported by Bourns Endowment funds at the University of California, Riverside.



## References

- [1] W. Benjamin and E. Leslie. *On Photography*. Reaktion Books, 2015. 1
- [2] J. Berger. *Ways of Seeing*. Penguin Modern Classics. Penguin Books Limited, 2008. 1
- [3] Zhe Cao, Gines Hidalgo Martinez, Tomas Simon, Shih-En Wei, and Yaser A Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 3, 8
- [4] Long Chen, Hanwang Zhang, Jun Xiao, Xiangnan He, Shiliang Pu, and Shih-Fu Chang. Counterfactual critic multi-agent training for scene graph generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4613–4623, 2019. 5
- [5] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6163–6171, 2019. 1, 2, 5, 6, 8
- [6] Vincent S Chen, Paroma Varma, Ranjay Krishna, Michael Bernstein, Christopher Re, and Li Fei-Fei. Scene graph prediction with limited labels. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2580–2590, 2019. 1
- [7] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder–decoder approaches. *Syntax, Semantics and Structure in Statistical Translation*, page 103, 2014. 2
- [8] Noam Chomsky. *Language and Mind*. Cambridge University Press, 3 edition, 2006. 1
- [9] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *Proceedings of the IEEE conference on computer vision and Pattern recognition*, pages 3076–3086, 2017. 2
- [10] Apoorva Dornadula, Austin Narcomey, Ranjay Krishna, Michael Bernstein, and Fei-Fei Li. Visual relationships as functions: Enabling few-shot scene graph prediction. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 1, 6
- [11] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pages 2121–2129, 2013. 1
- [12] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 1, 4
- [13] Nikolaos Gkanatsios, Vassilis Pitsikalis, Petros Koutras, and Petros Maragos. Attention-translation-relation network for scalable scene graph generation. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 2
- [14] Nikolaos Gkanatsios, Vassilis Pitsikalis, Petros Koutras, and Petros Maragos. Attention-translation-relation network for scalable scene graph generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. 2
- [15] Nikolaos Gkanatsios, Vassilis Pitsikalis, and Petros Maragos. From saturation to zero-shot visual relationship detection using local context. In *Proceedings of the British Machine Vision Conference*, 2020. 1, 7
- [16] Jiuxiang Gu, Handong Zhao, Zhe Lin, Sheng Li, Jianfei Cai, and Mingyang Ling. Scene graph generation with external knowledge and image reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2019. 2
- [17] David Harwath, Wei-Ning Hsu, and James Glass. Learning hierarchical discrete linguistic units from visually-grounded speech. In *International Conference on Learning Representations*, 2020. 1
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 6
- [19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [20] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6700–6709, 2019. 1
- [21] Zih-Siou Hung, Arun Mallya, and Svetlana Lazebnik. Contextual translation embedding for visual relationship detection and scene graph generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 1, 7
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456, 2015. 8
- [23] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1219–1228, 2018. 1
- [24] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015. 1
- [25] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015. 1
- [26] Boris Knyazev, Harm de Vries, Cătălina Cangea, Graham W Taylor, Aaron Courville, and Eugene Belilovsky. Graph density-aware losses for novel compositions in scene graph generation. *arXiv preprint arXiv:2005.08230*, 2020. 1, 7
- [27] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 1, 2, 5

- [28] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 734–750, 2018. 2, 3
- [29] C.D. Lewis. *The Poetic Image (Clark Lectures)*. Cambridge, 1946. 1
- [30] Yikang Li, Wanli Ouyang, Xiaogang Wang, and Xiao’ou Tang. Vip-cnn: Visual phrase guided convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1347–1356, 2017. 2
- [31] Yikang Li, Wanli Ouyang, Bolei Zhou, Jianping Shi, Chao Zhang, and Xiaogang Wang. Factorizable net: an efficient subgraph-based framework for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 335–351, 2018. 2, 6, 8
- [32] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1261–1270, 2017. 2, 6
- [33] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3, 5, 6
- [34] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 3, 5
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [36] Xin Lin, Changxing Ding, Jinquan Zeng, and Dacheng Tao. Gps-net: Graph property sensing network for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3746–3753, 2020. 2, 6
- [37] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *European conference on computer vision*, pages 852–869. Springer, 2016. 1, 2, 5, 7
- [38] Masood Mortazavi. Speech-image semantic alignment does not depend on any prior classification tasks. In *Proceedings of InterSpeech*, 2020. 1
- [39] Alejandro Newell and Jia Deng. Pixels to graphs by associative embedding. In *Advances in neural information processing systems*, pages 2171–2180, 2017. 2, 3, 5, 6, 7
- [40] Bryan A Plummer, Arun Mallya, Christopher M Cervantes, Julia Hockenmaier, and Svetlana Lazebnik. Phrase localization and visual relationship detection with comprehensive image-language cues. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1928–1937, 2017. 2, 7
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 1, 6
- [42] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986. 2
- [43] Gilbert Ryle. *The Concept of Mind*. U of Chicago Press, 1949. 1
- [44] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008. 2
- [45] S. Sontag. *On Photography*. Kushiell’s Legacy. Picador, 2001. 1
- [46] Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1556–1566, 2015. 2
- [47] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3716–3725, 2020. 1, 2, 5, 6, 7, 8
- [48] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6619–6628, 2019. 1, 2, 5, 6
- [49] Damien Teney, Lingqiao Liu, and Anton van Den Hengel. Graph-structured representations for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2017. 1
- [50] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9627–9636, 2019. 2, 5, 8
- [51] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 6
- [52] Wenbin Wang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Exploring context and visual pattern of relationship for scene graph generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2019. 2, 6
- [53] Sanghyun Woo, Dahun Kim, Donghyeon Cho, and In So Kweon. Linknet: Relational embedding for scene graph. In *Advances in Neural Information Processing Systems*, pages 560–570, 2018. 2
- [54] Hao Wu, Jiayuan Mao, Yufeng Zhang, Yuning Jiang, Lei Li, Weiwei Sun, and Wei-Ying Ma. Unified visual-semantic embeddings: Bridging vision and language with structured meaning representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6609–6618, 2019. 1

- [55] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. [8](#)
- [56] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5410–5419, 2017. [1](#), [2](#), [5](#), [6](#)
- [57] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1316–1324, 2018. [1](#)
- [58] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 670–685, 2018. [2](#), [6](#), [8](#)
- [59] Xu Yang, Kaihua Tang, Hanwang Zhang, and Jianfei Cai. Auto-encoding scene graphs for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10685–10694, 2019. [1](#)
- [60] Xu Yang, Hanwang Zhang, and Jianfei Cai. Shuffle-then-assemble: Learning object-agnostic visual relationship features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 36–52, 2018. [7](#)
- [61] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9657–9666, 2019. [2](#)
- [62] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 684–699, 2018. [1](#)
- [63] Ruichi Yu, Ang Li, Vlad I Morariu, and Larry S Davis. Visual relationship detection with internal and external linguistic knowledge distillation. In *Proceedings of the IEEE international conference on computer vision*, pages 1974–1982, 2017. [1](#), [2](#)
- [64] Alireza Zareian, Svebor Karaman, and Shih-Fu Chang. Bridging knowledge graphs to generate scene graphs. *arXiv preprint arXiv:2001.02314*, 2020. [1](#), [2](#), [6](#), [8](#)
- [65] Alireza Zareian, Haoxuan You, Zhecan Wang, and Shih-Fu Chang. Learning visual commonsense for robust scene graph generation. *arXiv preprint arXiv:2006.09623*, 2020. [1](#), [2](#)
- [66] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018. [1](#), [2](#), [4](#), [5](#), [6](#)
- [67] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5532–5540, 2017. [2](#), [6](#), [8](#)
- [68] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. [2](#), [3](#), [5](#)