

Learning to Relate from Captions and Bounding Boxes

Sarthak Garg*, Joel Ruben Antony Moniz*, Anshu Aviral*, Priyatham Bollimpalli*

School of Computer Science

Carnegie Mellon University

{sarthakg, jrmoniz, aanshu, pbollimp}@cs.cmu.edu

Abstract

In this work, we propose a novel approach that predicts the relationships between various entities in an image in a weakly supervised manner by relying on image captions and object bounding box annotations as the sole source of supervision. Our proposed approach uses a top-down attention mechanism to align entities in captions to objects in the image, and then leverage the syntactic structure of the captions to align the relations. We use these alignments to train a relation classification network, thereby obtaining both grounded captions and dense relationships. We demonstrate the effectiveness of our model on the Visual Genome dataset by achieving a recall@50 of 15% and recall@100 of 25% on the relationships present in the image. We also show that the model successfully predicts relations that are not present in the corresponding captions.

1 Introduction

Scene graphs serve as a convenient representation to capture the entities in an image and the relationships between them, and are useful in a variety of settings (for example, Johnson et al. (2015); Anderson et al. (2016); Liu et al. (2017)). While the last few years have seen considerable progress in classifying the contents of an image and segmenting the entities of interest without much supervision (He et al., 2017), the task of identifying and understanding the way in which entities in an image interact with each other without much supervision remains little explored.

Recognizing relationships between entities is non-trivial because the space of possible relationships is immense, and because there are $\mathcal{O}(n^2)$ relationships possible when n objects are present in an image. On the other hand, while image captions are easier to obtain, they are often not

completely descriptive of an image (Krishna et al., 2017). Thus, simply parsing a caption to extract relationships from them is likely to not sufficiently capture the rich content and detailed spatial relationships present in an image.

Since different images have different objects and captions, we believe it is possible to get the information that is not present in the caption of one image from other similar images which have the same objects and their captions. In this work, we thus aim to learn the relationships between entities in an image by utilizing only image captions and object locations as the source of supervision. Given that generating a good caption in an image requires one to understand the various entities and the relationships between them, we hypothesize that an image caption can serve as an effective weak supervisory signal for relationship prediction.

2 Related work

The task of Visual Relationship Detection has been the main focus of several recent works (Lu et al., 2016; Li et al., 2017a; Zhang et al., 2017a; Dai et al., 2017; Hu et al., 2017; Liang et al., 2017; Yin et al., 2018). The goal is to detect a generic \langle subject, predicate, object \rangle triplet present in an image. Various techniques have been proposed to solve this task, such as by using language priors (Lu et al., 2016; Yatskar et al., 2016), deep network models (Zhang et al., 2017a; Dai et al., 2017; Zhu and Jiang, 2018; Yin et al., 2018), referring expressions (Hu et al., 2017; Cirik et al., 2018) and reinforcement learning (Liang et al., 2017). Recent work has also studied the closely related problem of Scene Graph Generation, (Li et al., 2017b; Newell and Deng, 2017; Xu et al., 2017; Yang et al., 2017, 2018). The major limitation of the aforementioned techniques is that they are supervised, and require the presence of ground truth scene graphs or rela-

*Equal Contribution

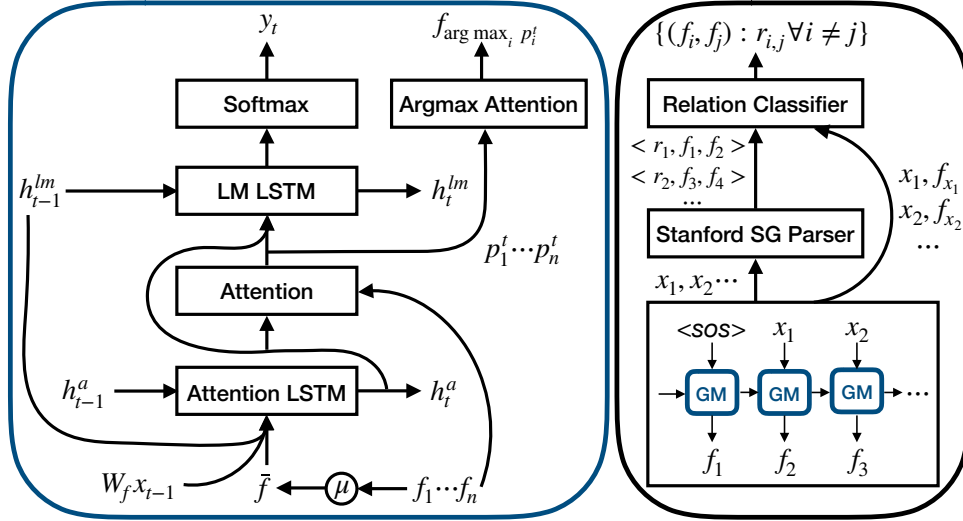


Figure 1: C-GEARD Architecture (left) and its integration with the relation classifier (right). C-GEARD acts as the Grounding Module (GM) in our relation classifier.

tion annotations. Obtaining these annotations can be an extremely tedious and time consuming process that often needs to be done manually. Our model in contrast does the same task through weak supervision, which makes this annotation significantly easier.

Most similar to our current task is work in the domain of Weakly Supervised Relationship Detection. [Peyre et al. \(2017\)](#) uses weak supervision to learn the visual relations between the pairs of objects in an image using a weakly supervised discriminative clustering objective function ([Bach and Harchaoui, 2008](#)), while [Zhang et al. \(2017b\)](#) uses a region-based fully convolutional network neural network to perform the same task. They both use $\langle \text{subject, predicate, object} \rangle$ annotations without any explicit grounding in the images as the source of weak supervision, but require these annotations in the form of image-level triplets. Our task, however, is more challenging, because free-form captions can potentially be both extremely unstructured and significantly less informative than annotated structured relations.

3 Proposed Approach

Our proposed approach consists of three sequential modules: a feature extraction module, a grounding module and a relation classifier module. Given the alignments found by the grounding module, we train the relation classifier module, which takes in a pair of object features and classifies the relation between them.

3.1 Feature Extraction

Given an image I with n objects and their ground truth bounding boxes $\{b_1, b_2, \dots, b_n\}$, the feature extraction module extracts their feature representations $\mathcal{F} = \{f_1, f_2, \dots, f_n\}$. To avoid using ground truth instance-level class annotations that would be required to train an object detector, we use a ResNet-152 network pre-trained on ImageNet as our feature extractor. For every object i , we crop and resize the portion of the image I corresponding to the bounding box b_i and feed it to the ResNet model to get its feature representation f_i . f_i is a dense d -dimensional vector capturing the semantic information of the i^{th} object. Note that we do not fine-tune the ResNet architecture.

3.2 Grounding Caption Words to Object Features

Given an image I , its caption consisting of words $\mathcal{W} = \{w_1, w_2, \dots, w_k\}$ and the feature representations \mathcal{F} obtained above, the grounding module aligns the entities and relations found in the captions with the objects' features and the features corresponding to pairs of objects in the image. It thus aims to find the subset of words in the caption corresponding to entities $\mathcal{E} \subseteq \mathcal{W} \mid \mathcal{E} = \{e_{i_1}, e_{i_2}, \dots, e_{i_m}\}$, and to ground each such word with its best matching object feature f_{i_j} . It also aims to find the subset of relational words $\mathcal{R} \subseteq \mathcal{W} \mid \mathcal{R} = \{r_{i_1}, r_{i_2}, \dots, r_{i_l}\}$ and to ground each relation to a pair of object features $\{f_{i_{\text{subj}}}, f_{i_{\text{obj}}}\}$ which correspond to the subject and object of that relation.

To identify and ground the relations between entities in an image, we propose C-GEARD (Captioning-Grounding via Entity Attention for Relation Detection). C-GEARD passes the caption through the Stanford Scene Graph Parser (Schuster et al., 2015) to get the set of triplets $\mathcal{T} = \{(s_1, p_1, o_1), (s_2, p_2, o_2), \dots, (s_g, p_g, o_g)\}$. Each triplet corresponds to one relation present in the caption. For $(s_i, p_i, o_i) \in \mathcal{T}$, s_i , p_i and o_i denote *subject*, *predicate* and *object* respectively. The entity and relation subsets are then constructed as:

$$\mathcal{E} = \bigcup_{(s_i, p_i, o_i) \in \mathcal{T}} \{s_i, o_i\} \quad \mathcal{R} = \bigcup_{(s_i, p_i, o_i) \in \mathcal{T}} \{p_i\}$$

Captioning using visual attention has proven to be very successful in aligning the words in a caption to their corresponding visual features, such as in Anderson et al. (2018). As shown in Figure 1, we adopt the two-layer LSTM architecture in Anderson et al. (2018); our end goal, however, is to associate each word with the closest object feature rather than producing a caption.

The lower Attention LSTM cell takes in the words and the global image context vector (\bar{f} , the mean of all features \mathcal{F}), and its hidden state h_t^a acts as a query vector. This query vector is used to attend over the object features $\mathcal{F} = \{f_1, f_2, \dots, f_n\}$ (serving as both key and value vectors) to produce an attention vector which summarizes the key visual information needed for predicting the next word. The Attention module is parameterized as in Bahdanau et al. (2014). The concatenation of the query vector and the attention vector is passed as an input to the upper LM-LSTM cell, which predicts the next word of the caption.

The model is trained by minimizing the standard negative log-likelihood loss.

$$\mathcal{L}_{NLL} = -\frac{1}{k} \sum_{i=1}^k \log(\mathcal{P}(w_i | w_1 \dots w_{i-1}))$$

Let $p_x^{w_i}$ denote the attention probability over feature x when *previous word* w_{i-1} is fed into the LSTM. C-GEARD constructs alignments of the entity and relation words as follows:

$$w_e \xrightarrow{\text{align}} \arg \max_{f \in \mathcal{F}} (p_f^{w_e}) \quad \forall w_e \in \mathcal{E}$$

$$p_i \xrightarrow{\text{align}} (\arg \max_{f \in \mathcal{F}} (p_f^{s_i}), \arg \max_{f \in \mathcal{F}} (p_f^{o_i})) \quad \forall p_i \in \mathcal{R}$$

3.3 Relation Classifier

We run the grounding module C-GEARD over the training captions to generate a “grounded” relationship dataset consisting of tuples $\{((f_i, f_j), p_{i,j})\}$, where f_i and f_j are two object features and $p_{i,j}$ refers to the corresponding aligned predicates. These predicates occur in free form; however, the relations in the test set are restricted to only the top 50 relation classes. We manually annotate the correspondence between the 300 most frequent parsed predicates and their closest relation class. For example, we map the parsed predicates *dress in*, *sitting in* and *inside* to the canonical relation class *in*. Using this mapping we get tuples of the form $\{((f_i, f_j), c_{i,j})\}$ where $c_{i,j}$ denotes the canonical class corresponding to $p_{i,j}$.

Since this dataset is generated by applying the grounding module on the set of all images and the corresponding captions, it pools the relation information from across the whole dataset, which we then use to train our relation classifier.

We parameterize the relation classifier with a 2-layer MLP. Given the feature vectors of any two objects f_i and f_j , the relation classifier is trained to classify the relation $c_{i,j}$ between them.

3.4 Model at Inference

During inference, the features extracted from each pair of objects is passed through the relation classifier to predict the relation between them.

4 Experiments

4.1 Dataset

We use the MS COCO (Lin et al., 2014) dataset for training and the Visual Genome (Krishna et al., 2017) dataset for evaluation. MS COCO has images and their captions, and Visual Genome contains images and their associated scene graphs. The Visual Genome dataset consists in part of MS COCO images, and since we require ground truth captions and bounding boxes during training, we filter the Visual Genome dataset by considering only those images which are part of the original MS COCO dataset. Similar to Xu et al. (2017), we manually remove poor quality and overlapping bounding boxes with ambiguous object names, and filter to keep the 150 most frequent object categories and 50 most frequent predicates. Our final dataset thus comprises of 41,731 images with 150 unique objects and 50 unique relations. We use a 70-30 train-validation split. We use the same

Recall@	IMP (Xu et al., 2017)	Pixel2Graph (Newell and Deng, 2017)	Graph-RCNN (Yang et al., 2018)	Parsed caption (baseline)	C-GEARD (ours)
50	44.8	68.0	54.2	4.1	15.3
100	53.0	75.2	59.1	4.1	25.2

Table 1: Comparison with respect to Recall@50 and Recall@100 on PredCls metric, in %.

test set as Xu et al. (2017), so that the results are comparable with other supervised baselines.

4.2 Baselines

Since, to the best of our knowledge, this work is the first to introduce the task of weakly supervised relationship prediction solely using captions and bounding boxes, we do not have any directly comparable baselines, i.e., all other work is either completely supervised or relies on all ground truth entity-relation triplets being present at train time. Consequently, we construct baselines relying solely on captions and ground truth bounding box locations that are comparable to our task. In particular, running the Stanford Scene Graph Parser (Schuster et al., 2015) on ground truth captions constructs a scene graph just from the image captions (which almost never capture all the information present in an image). We use this baseline as a lower bound, and to obtain insight into the limitations of scene graphs directly generated from captions. On the other hand, we use supervised scene graph generation baselines (Yang et al., 2018; Newell and Deng, 2017) to upper bound our performance, since we rely on far less information and data.

4.3 Evaluation Metric

As our primary objective is to detect relations between entities, we use the PredCls evaluation metric (Xu et al., 2017), defined as the performance of recognizing the relation between two objects given their ground truth locations. We only use the entity bounding boxes’ locations without knowing the ground truth objects they contain. We show results on Recall@ k (the fraction of top k relations predicted by the model contained in the ground truth) for $k = 50$ and 100. The predicted relations are ranked over all objects pairs for all relation classes by the relation classifier’s model confidence.

5 Results and Discussion

5.1 Performance

We show the performance of C-GEARD in Table 1. We compare its performance with various super-

vised baselines, as well as a baseline which parses relations from just the caption using Stanford Scene Graph Parser (Schuster et al., 2015) (caption-only baseline), on the PredCls metric. Our proposed method substantially outperforms the caption-only baseline. This shows that our model predicts relationships more successfully than by purely relying on captions, which contain limited information. This in turn supports our hypothesis that it is possible to detect relations by pooling information from captions across images, without requiring all ground truth relationship annotations for every image.

Note that our model is at a significant disadvantage when compared to supervised approaches. First, we use pre-trained ResNet features (trained on a classification task) without any fine-tuning; supervised methods, however, use Faster RCNN (Ren et al., 2015), whose features are likely much better suited for multiple objects. Second, supervised methods likely have a better global view than C-GEARD, because Faster RCNN provides a significantly larger number of proposals, while we rely on ground truth regions which are far fewer in number. Third, and most significant, we have no ground truth relationship or class information, relying purely on weak supervision from captions to provide this information. Finally, since we require captions, we use significantly less data, training on the subset of Visual Genome overlapping with MS COCO (and has ground truth captions as a result).

5.2 Relation Classification

We train the relation classifier on image features of entity pairs and using the relations found in the caption as the only source of supervision. On the validation set, we obtain a relation classification accuracy of 22%.

We compute the top relations that the model gets most confused about, shown in Table 2. We observe that even when the predictions are not correct, they are semantically close to the ground truth relation class.

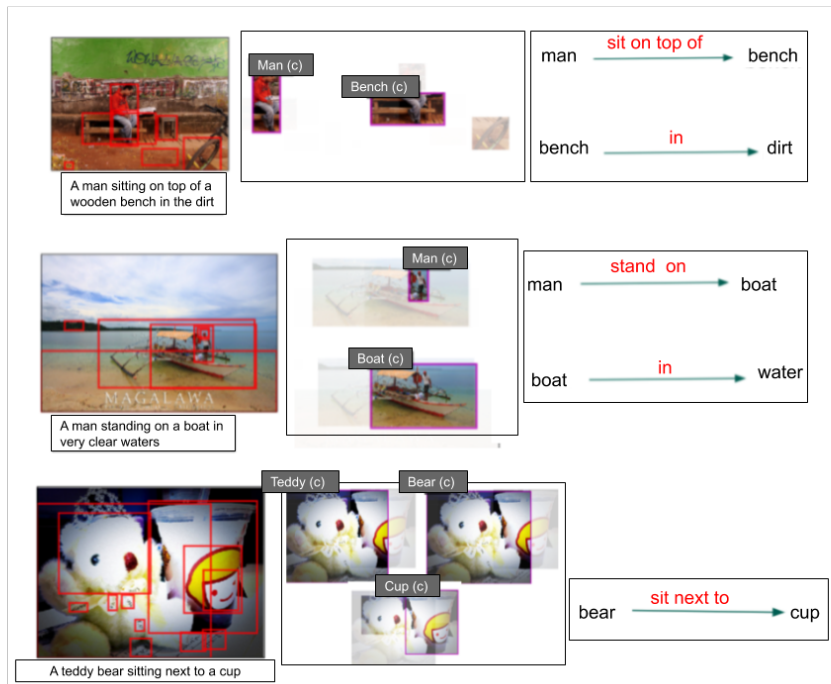


Figure 2: Attention masks for each of the entities in the caption for C-GEARD. The output of the Stanford Scene Graph Parser is given on the right.

Relation	Confusion with Relations
above	on, with, sitting on, standing on, of
carrying	holding, with, has, carrying, on
laying on	on, lying on, in, has
mounted on	on, with, along, at, attached to

Table 2: Relations the classification model gets most confused about

5.3 Visualizations

Three images with their captions are given in Figure 2. We can see that C-GEARD generates precise entity groundings, and that the Stanford Scene Graph Parser generates correct relations. This results in the correct grounding of the entities and relations which yields accurate training samples for the relation classifier.

6 Conclusion

In this work, we propose a novel task of weakly-supervised relation prediction, with the objective of detecting relations between entities in an image purely from captions and object-level bounding box annotations without class information. Our proposed method builds upon top-down attention (Anderson et al., 2018), which generates captions and grounds word in these captions to entities in images.

We leverage this along with structure found from the captions by the Stanford Scene Graph Parser (Schuster et al., 2015) to allow for the classification of relations between pairs of objects without having ground truth information for the task. Our proposed approaches thus allow weakly-supervised relation detection.

There are several interesting avenues for future work. One possible line of work involves removing the requirement of ground truth bounding boxes altogether by leveraging a recent line of work that does weakly-supervised object detection (such as (Oquab et al., 2015; Bilen and Vedaldi, 2016; Zhang et al., 2018; Bai and Liu, 2017; Arun et al., 2018)). This would reduce the amount of supervision required even further. An orthogonal line of future work might involve using a Visual Question Answering (VQA) task (such as in Krishna et al. (2017)), either on its own replacing the captioning task, or in conjunction with the captioning task with a multi-task learning objective.

Acknowledgements

We would like to thank Louis-Philippe Morency, Carla Viegas, Volkan Cirik and Barun Patra for helpful discussions and feedback. We would also like to thank the anonymous reviewers for their insightful comments and suggestions.

References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, pages 382–398. Springer.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, volume 3, page 6.
- Aditya Arun, CV Jawahar, and M Pawan Kumar. 2018. Dissimilarity coefficient based weakly supervised object detection. *arXiv preprint arXiv:1811.10016*.
- Francis R Bach and Zaïd Harchaoui. 2008. Diffrac: a discriminative and flexible framework for clustering. In *Advances in Neural Information Processing Systems*, pages 49–56.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Peng Tang Xinggang Wang Xiang Bai and Wenyu Liu. 2017. Multiple instance detection network with on-line instance classifier refinement.
- Hakan Bilen and Andrea Vedaldi. 2016. Weakly supervised deep detection networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2846–2854.
- Volkan Cirik, Taylor Berg-Kirkpatrick, and Louis-Philippe Morency. 2018. Using syntax to ground referring expressions in natural images. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Bo Dai, Yuqi Zhang, and Dahua Lin. 2017. Detecting visual relationships with deep relational networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 3298–3308. IEEE.
- Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask r-cnn. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2980–2988. IEEE.
- Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. 2017. Modeling relationships in referential expressions with compositional modular networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 4418–4427. IEEE.
- Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. 2015. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73.
- Yikang Li, Wanli Ouyang, Xiaogang Wang, and Xiao’ou Tang. 2017a. Vip-cnn: Visual phrase guided convolutional neural network. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 7244–7253. IEEE.
- Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. 2017b. Scene graph generation from objects, phrases and region captions. In *ICCV*.
- Xiaodan Liang, Lisa Lee, and Eric P Xing. 2017. Deep variation-structured reinforcement learning for visual relationship and attribute detection. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 4408–4417. IEEE.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Siqi Liu, Zhenhai Zhu, Ning Ye, Sergio Guadarrama, and Kevin Murphy. 2017. Improved image captioning via policy gradient optimization of spider. In *Proc. IEEE Int. Conf. Comp. Vis*, volume 3, page 3.
- Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. 2016. Visual relationship detection with language priors. In *European Conference on Computer Vision*, pages 852–869. Springer.
- Alejandro Newell and Jia Deng. 2017. Pixels to graphs by associative embedding. In *Advances in neural information processing systems*, pages 2171–2180.
- Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. 2015. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 685–694.
- Julia Peyre, Josef Sivic, Ivan Laptev, and Cordelia Schmid. 2017. Weakly-supervised learning of visual relations. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5179–5188.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99.
- Sebastian Schuster, Ranjay Krishna, Angel Chang, Li Fei-Fei, and Christopher D Manning. 2015. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In

Proceedings of the fourth workshop on vision and language, pages 70–80.

Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. 2017. Scene graph generation by iterative message passing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2.

Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. 2018. Graph r-cnn for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 670–685.

Michael Ying Yang, Wentong Liao, Hanno Ackermann, and Bodo Rosenhahn. 2017. On support relations and semantic scene graphs. *ISPRS journal of photogrammetry and remote sensing*, 131:15–25.

Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5534–5542.

Guojun Yin, Lu Sheng, Bin Liu, Nenghai Yu, Xiaogang Wang, Jing Shao, and Chen Change Loy. 2018. Zoom-net: Mining deep feature interactions for visual relationship recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 322–338.

Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. 2017a. Visual translation embedding network for visual relation detection. In *CVPR*, volume 1, page 5.

Hanwang Zhang, Zawlin Kyaw, Jinyang Yu, and Shih-Fu Chang. 2017b. Ppr-fcn: weakly supervised visual relation detection via parallel pairwise r-fcn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4233–4241.

Xiaopeng Zhang, Jiashi Feng, Hongkai Xiong, and Qi Tian. 2018. Zigzag learning for weakly supervised object detection. *arXiv preprint arXiv:1804.09466*.

Yaohui Zhu and Shuqiang Jiang. 2018. Deep structured learning for visual relationship detection. In *Thirty-Second AAAI Conference on Artificial Intelligence*.