

# Triplet-Aware Scene Graph Embeddings

Brigit Schroeder\*  
 University of California, Santa Cruz  
 brschroe@ucsc.edu

Subarna Tripathi  
 Intel AI Lab  
 subarna.tripathi@intel.com

Hanlin Tang  
 Intel AI Lab  
 hanlin.tang@intel.com

## Abstract

*Scene graphs have become an important form of structured knowledge for tasks such as for image generation, visual relation detection, visual question answering, and image retrieval. While visualizing and interpreting word embeddings is well understood, scene graph embeddings have not been fully explored. In this work, we train scene graph embeddings in a layout generation task with different forms of supervision, specifically introducing triplet supervision and data augmentation. We see a significant performance increase in both metrics that measure the goodness of layout prediction, mean intersection-over-union (mIoU) (52.3% vs. 49.2%) and relation score (61.7% vs. 54.1%), after the addition of triplet supervision and data augmentation. To understand how these different methods affect the scene graph representation, we apply several new visualization and evaluation methods to explore the evolution of the scene graph embedding. We find that triplet supervision significantly improves the embedding separability, which is highly correlated with the performance of the layout prediction model.*

## 1. Introduction

Scene graphs are a structured data format which encodes semantic relationships between objects [22]. Objects are represented as nodes in the graph and are connected by edges that expresses relationship, in the form of triplets. Each triplet is comprised of a  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$  which semantically describes the relationship between two objects such as  $\langle \text{car}, \text{on}, \text{road} \rangle$  or  $\langle \text{dog}, \text{left of}, \text{person} \rangle$ .

Scene graphs can be processed by graph convolutional neural networks (GCNN) which are able to pass information along graph edges[7]. Layout prediction models, such as those used image generation [7], are multi-stage networks (see Figure 1) which predict scene layout masks and object localization (bounding boxes), based upon the structure of the scene graph. The first stage of the layout predic-

tion model is a GCNN that is used to learn a scene graph embedding. Object embeddings are passed to the next stage of the network to predict a bounding box and mask for each object node in the scene graph. We are interested in layout prediction models in particular because a scene graph embedding is learned in the process of training a model. The performance metrics used for evaluating the quality of scene layout, namely mean intersection-over-union (mIoU) and relation score [15], can be used to correlate with the relative “quality” of the scene graph embedding. Relation score is particularly relevant as it measures the compliance of the generated layout to the scene graph, possibly giving insight into how well the embedding is structured, especially in terms of class separability.

In our work, we introduce several new supervisory signals that are conditioned upon triplet embeddings to train scene layout prediction models. We also apply data augmentation by using heuristic-based relationships [15] to maximize the number of triplets during training. The goal is to learn a triplet-aware scene graph embedding with the hypothesis additional supervision and data augmentation will enrich the embedding representation. To this end, we also introduce several methods of embedding visualization and evaluation that help understand the degree of separability that is achieved from our improved layout prediction models.

## 2. Related Work

Visual Genome [9], a human-annotated scene graph dataset, has known issues of incomplete and incorrect annotations. Additionally this only provides the bounding boxes for object instances but not their segmentation masks. On the other hand, the synthetic scene graphs generated from COCO stuff [2] are limited to simple geometric relationships (above, below, left, right, inside, surrounding) but are not hampered by incorrect annotations. COCO stuff also provides segmentation masks for instances. In this paper we use synthetic scene graphs from COCO stuff [2] for our experiments.

Generating scene graphs from visual features [22, 24,

\*Work done as a resident in Intel AI Lab.

23, 4, 3, 14, 26, 20, 13, 5] is a relatively explored task. Wan *et al.* [19] specifically predicts new triplets for scene graph completion using existing scene graphs and visual features. Another line of work that emerged recently [7, 15, 11, 16, 8, 12, 18] take scene graphs as input and produce final RGB images. All of these methods perform an intermediate layout prediction by learning embeddings of nodes. In closely related work, Belilovsky *et al.* [1] learn a joint visual-scene graph embedding for use in image retrieval. None of the above explored utilizing  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$  triplets as additional supervisory signal for more effective structured prediction.

Few papers that aim to detect visual relations from images also analyze the learned embeddings qualitatively. Dornadula *et al.* [3] visualizes learned object category embeddings using 2-dimensional t-SNE plots. In GraphRCNN [23] Yang *et al.* showed how common sense of  $\langle \text{object-predicate} \rangle$  co-occurrence [6] emerges as the top scorers based upon the extracted weights that use visual features. They show similar common sense also appears for  $\langle \text{object-object} \rangle$  co-occurrences. In VTransE paper [25] utilizes t-SNE for visualizing the embeddings of predicate vectors and speculate when the model learns only co-occurrence vs. the actual meaning of the relationship by inspecting the neighbors in the t-SNE space. Note all these described models utilize visual features. In contrast, our model is trained from scene graphs without access to visual features. We provide quantitative and qualitative measures on *why* the embeddings trained by the proposed model is superior.

From an embedding introspection perspective, the closest work is [21]. In this concurrent work, Xu *et al.* introduced generic embeddings for common visual objects. They visualized embeddings of most common objects using t-SNE, and qualitatively studied results on vector decomposition and projection on specific axes anecdotally. We note that these embeddings are learned solely based on co-occurrence and thus they are not able to deal with relationships between objects.

### 3. Method

**Layout Prediction with Triplet Supervision.** As part of our scene graph embedding introspection, we use a layout prediction network inspired by the image generation pipeline in [7]. Figure 1 gives an overview of the network architecture. A GCNN processes an input scene graph to produce embeddings corresponding to object nodes in the graph. Singleton object embeddings are passed to the next stage of the layout prediction network per [7].

We utilize the object embeddings to form a set of triplet embeddings where each is composed of a  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$  embedding. We pass these through a triplet mask prediction network. Figure 2 highlights the details of

the triplet mask prediction process. Rather than just learn individual class labels, the network learns to label objects as either subject or object, enforcing both an ordering and relationship between objects. We also pass triplet embeddings through a triplet superbox regression network, where we train the network to do joint localization over subject and object bounding boxes.

Ultimately, all of the outputs of the second stage of the layout prediction model are used to compose a scene layout mask with object localization.

**Data Augmentation.** We applied the method in [15] which uses heuristics to augment scene graphs with new spatial relations that induce a richer learned representation. The depth order between objects from an observer’s viewpoint is quasi-exhaustively determined, and for 2D images, determining this order is non-trivial. Linear perspective-based heuristics are utilized for augmenting spatial relationship vocabulary.

**Training.** We train the layout prediction network to minimize two additional triplet-based losses in addition to those used in [7]:

- Triplet mask loss,  $L_{\text{triplet-mask}}$ , penalizing differences between ground truth triplet masks and predicted triplet masks with pixelwise cross-entropy loss.
- Triplet superbox loss,  $L_{\text{superbox}}$ , penalizing the  $L_2$  difference between the ground truth and predicted triplet superboxes.

## 4. Experiments

### 4.1. Dataset

**COCO:** We performed experiments on the 2017 COCO-Stuff dataset [2], which augments the COCO dataset [10] with additional stuff categories. The dataset annotates 40K train and 5K val images with bounding boxes and segmentation masks for 80 thing categories (people, cars, etc.) and 91 stuff categories (sky, grass, etc.). Similar to [7], we used thing and stuff annotations to construct synthetic scene graphs based on the 2D image coordinates of the objects, encoding six mutually exclusive geometric relationships: *left of*, *right of*, *above*, *below*, *inside*, *surrounding*. We ignored objects covering less than 2% of the image, and used images with 3 to 8 objects.

Table 1. Layout Prediction Model.

Model	mIoU	Relation Score
Baseline [7]	49.2%	54.1%
Triplet Supervision	50.3%	59.4%
Triplet Supervision + DA	52.3%	61.7%

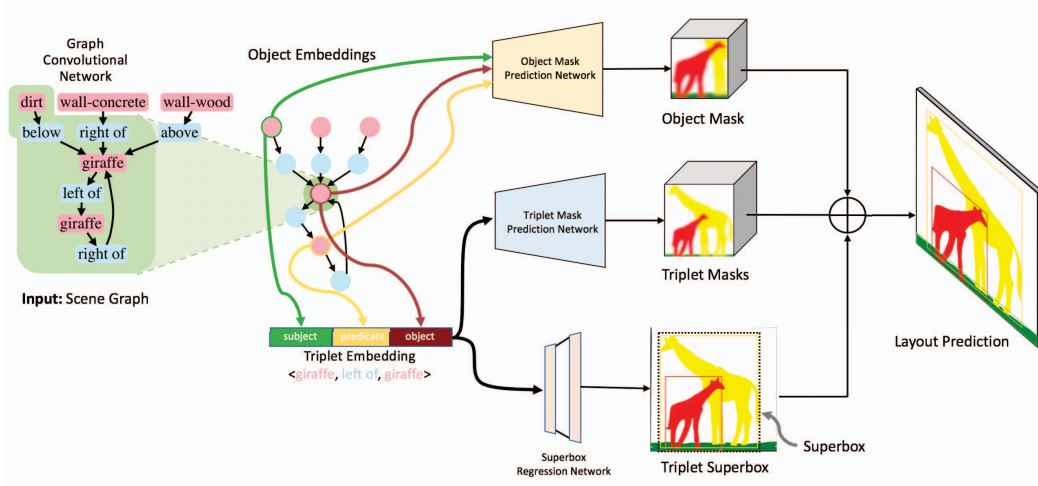


Figure 1. **Triplet-Supervised Layout Prediction Network.** Here we show a multi-stage network which takes as input a scene graph and outputs a scene layout prediction. A graph convolutional neural network processes the scene graph to produce object embeddings, which can be used to form triplet embeddings composed of  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ . Both single object and triplet embeddings are passed to the next stage to predict object masks, triplet masks and triplet superboxes. Ultimately, these are combined to predict a scene layout.




Figure 2. **Triplet Mask Prediction.** Triplets containing a  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$  found in a scene graph are used to predict corresponding triplet masks, labelling pixels either as subject and object. The mask prediction is used as supervisory signal during training.

## 4.2. Scene Layout Prediction

In order to evaluate the quality of scene layout prediction, we compared the predicted layout with the ground truth using intersection-over-union (mIoU) and relation score. The relation score measures the percentage of scene graph relations that are satisfied in the generated layout. The additional supervision and data augmentation introduced significantly improves over the baseline model [7] as can be seen in Table 1.

## 4.3. Embedding Visualization and Separability

To visualize the results, we first visualized the scene graph embedding using a t-SNE plot [17] (Figure 3) for the three models. For clarity, we only show the five most frequent classes in COCO-stuff dataset in the visualization.

With the triplet supervision and data augmentation, the individual classes become more separable, as seen in the progression of models. We quantified this observation by using an SVM with a linear kernel to classify the class categories (Table 2). The classification accuracy of the ten most frequent classes along with the mean classification accuracy across all classes is shown. The model trained with triplet supervision and data augmentation shows significantly better class separability over the baseline, as measured by classification accuracy. This is especially apparent in lower frequency classes such as wall, building and pavement.

In Figure 4, the mean scene embedding nodes for the fifty most frequent classes in COCO-stuff are visualized. We can see that several semantically similar objects clusters form, such as ones containing “animals” and “ground coverings” (indicated by colored boxes). Higher-level clusters such as “indoor/outdoor structural” (blue box) and “materials” (yellow box) which are also semantically similar at a higher level, are found closer together in embedding space. This kind of hierarchical clustering can also be seen in the 2d dendrogram in Figure 5 which visualizes a clustered heatmap of mean embedding distances between the top fifty most frequent classes. Distinct regions of blue (e.g. near-zero distance, other than along diagonal) indicate clusters of object classes that form correlated logical groupings in embedding space, such as outdoor objects (e.g. tree, bush, fence) and outdoor environments (e.g. sky, clouds, ground).

Table 2. Scene Graph Embedding Classification

Model	person	tree	sky	grass	metal	wall	bldg	pavement	road	clothes	mean accur
Baseline	97.1%	75.6%	85.9%	64.6%	77.6%	51.5%	30.6%	32.5%	52.3%	79.5%	48.3%
Triplet Supervision + DA	98.1%	73.1%	89.4%	70.8%	84.8%	61.9%	50.9%	59.1%	58.8%	79.5%	59.4%




Figure 3. **Triplet Supervision and Data Augmentation.** A t-SNE plot of the top 5 highest frequency classes in the COCO-Stuff dataset which shows (from left to right) that as supervision and data augmentation are added, clusters representing each class are tighter and more distinct.




Figure 4. **Mean Embedding Representation.** A t-SNE plot of mean embeddings for the top 50 highest frequency classes in the COCO-Stuff dataset. The embedding forms logical clusters of classes that are correlated.

## 5. Conclusion

In the course of our scene graph embedding introspection, we have identified many key findings. We find that




Figure 5. **Scene Graph Embedding Dendrogram** A dendrogram represented as distance heatmap between the mean embeddings of top 50 highest frequency in the COCO-Stuff dataset.

introducing triplet supervision and applying data augmentation significantly improves the performance of layout prediction models. We understand in turn that this correlates with improved scene embedding separability, quantified by measuring the classification accuracy of all nodes in the scene graph embedding. We have identified that scene



graph embedding structure, sometimes expressed as separability, can be influenced by both method of training and richness of scene graph representation. We find overall there is a strong correlation between scene graph embedding structure and scene layout prediction model performance. The metrics for this kind of model (or other models trained with scene graph embeddings) may be useful in understanding future scene-graph related tasks.

## References

- [1] E. Belilovsky, M. Blaschko, J. R. Kiros, R. Urtasun, and R. Zemel. Joint Embeddings of Scene Graphs and Images. International Conference On Learning Representations - Workshop, 2017. Poster. [2](#)
- [2] H. Caesar, J. Uijlings, and V. Ferrari. Coco-stuff: Thing and stuff classes in context. In *Computer vision and pattern recognition (CVPR), 2018 IEEE conference on*. IEEE, 2018. [1](#), [2](#)
- [3] A. Dornadula, A. Narcomey, R. Krishna, M. Bernstein, and L. Fei-Fei. Visual relationships as functions: Enabling few-shot scene graph prediction. *CoRR*, abs/1906.04876, 2019. [2](#)
- [4] B. L. N. Y. X. W. J. S. C. C. L. Guojun Yin, Lu Sheng. Zoomnet: Mining deep feature interactions for visual relationship recognition. In *ECCV*, 2018. [2](#)
- [5] 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. [2](#)
- [6] Jianwei Yang. Scene graph generation and its application to vision-and-language tasks. [https://drive.google.com/file/d/loH5njQtMZkD39WFuXyM7uw1Px\\_pyLsaq/view](https://drive.google.com/file/d/loH5njQtMZkD39WFuXyM7uw1Px_pyLsaq/view), 2019. [CVPR 2019 Tutorial, Accessed: 2019-08-14]. [2](#)
- [7] J. Johnson, A. Gupta, and L. Fei-Fei. Image generation from scene graphs. *CVPR*, 2018. [1](#), [2](#), [3](#)
- [8] A. A. Jyothi, T. Durand, J. He, L. Sigal, and G. Mori. Layoutvae: Stochastic scene layout generation from a label set. *ICCV*, 2019. [2](#)
- [9] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L. Li, D. A. Shamma, M. S. Bernstein, and F. Li. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *CoRR*, abs/1602.07332, 2016. [1](#)
- [10] T.-Y. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [2](#)
- [11] G. Mittal, S. Agrawal, A. Agarwal, S. Mehta, and T. Marwah. Interactive image generation using scene graphs. *CoRR*, abs/1905.03743, 2019. [2](#)
- [12] Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation. <https://www.youtube.com/watch?v=V2v0qEPsjr0tm>, 2019. [ICCV 2019, Accessed: 2019-08-14]. [2](#)
- [13] 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. [2](#)
- [14] K. Tang, H. Zhang, B. Wu, W. Luo, and W. Liu. Learning to compose dynamic tree structures for visual contexts. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [15] S. Tripathi, A. Bhiwandiwala, A. Bastidas, and H. Tang. Heuristics for image generation from scene graphs. *ICLR LLD workshop*, 2019. [1](#), [2](#)
- [16] S. Tripathi, S. N. Sridhar, S. Sundaresan, and H. Tang. Compact scene graphs for layout composition and patch retrieval. *CVPRW*, 2019. [2](#)
- [17] L. van der Maaten and G. Hinton. Visualizing data using t-sne, 2008. [3](#)
- [18] D. M. Vo and A. Sugimoto. Visual-relation conscious image generation from structured-text, 2019. [2](#)
- [19] H. Wan, Y. Luo, B. Peng, and W.-S. Zheng. Representation learning for scene graph completion via jointly structural and visual embedding. pages 949–956, 07 2018. [2](#)
- [20] S. Woo, D. Kim, D. Cho, and I.-S. Kweon. Linknet: Relational embedding for scene graph. In *NeurIPS*, 2018. [2](#)
- [21] C. Xu, Z. Chen, and C. Li. Obj-glove: Scene-based contextual object embedding. *CoRR*, abs/1907.01478, 2019. [2](#)
- [22] D. Xu, Y. Zhu, C. Choy, and L. Fei-Fei. Scene graph generation by iterative message passing. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. [1](#), [2](#)
- [23] J. Yang, J. Lu, S. Lee, D. Batra, and D. Parikh. Graph R-CNN for scene graph generation. In *The European Conference on Computer Vision (ECCV)*, September 2018. [2](#)
- [24] R. Zellers, M. Yatskar, S. Thomson, and Y. Choi. Neural motifs: Scene graph parsing with global context. *CoRR*, abs/1711.06640, 2017. [2](#)
- [25] H. Zhang, Z. Kyaw, S. Chang, and T. Chua. Visual translation embedding network for visual relation detection. *CVPR*, abs/1702.08319, 2017. [2](#)
- [26] J. Zhang, K. J. Shih, A. Elgammal, A. Tao, and B. Catanzaro. Graphical contrastive losses for scene graph generation. *CoRR*, abs/1903.02728, 2019. [2](#)