

Self-Supervised Learning for Visual Relationship Detection through Masked Bounding Box Reconstruction

Zacharias Anastasakis^{1,4}

zaxarisanastasakis@gmail.com

Dimitrios Mallis²

dimitrios.mallis@uni.lu

Markos Diomataris³

mdiomataris@student.ethz.ch

George Alexandridis⁴

gealexandri@islab.ntua.gr

Stefanos Kollias⁴

stefanos@cs.ntua.gr

Vassilis Pitsikalis¹

vpitsik@deeplab.ai

¹ Deeplab, Athens

² SnT, University of Luxembourg

³ ETH, Zürich

⁴ National Technical University of Athens

Abstract

We present a novel self-supervised approach for representation learning, particularly for the task of Visual Relationship Detection (VRD). Motivated by the effectiveness of Masked Image Modeling (MIM), we propose Masked Bounding Box Reconstruction (MBBR), a variation of MIM where a percentage of the entities/objects within a scene are masked and subsequently reconstructed based on the unmasked objects. The core idea is that, through object-level masked modeling, the network learns context-aware representations that capture the interaction of objects within a scene and thus are highly predictive of visual object relationships. We extensively evaluate learned representations, both qualitatively and quantitatively, in a few-shot setting and demonstrate the efficacy of MBBR for learning robust visual representations, particularly tailored for VRD. The proposed method is able to surpass state-of-the-art VRD methods on the Predicate Detection (PredDet) evaluation setting, using only a few annotated samples. We make our code available at <https://github.com/deeplab-ai/SelfSupervisedVRD>.

1. Introduction

Developing machines endowed with the ability to interpret and make decisions based on visual inputs is a critical Computer Vision objective. Research in the field of scene understanding aims to analyze an entire scene or image, in a way similar to that of a human observer. Visual Relationship Detection (VRD) provides an effective approach to scene understanding and constitutes an important component of larger vision pipelines for tasks such as image captioning [19] and visual question answering [26, 32, 34].

VRD goes beyond identifying/classifying individual

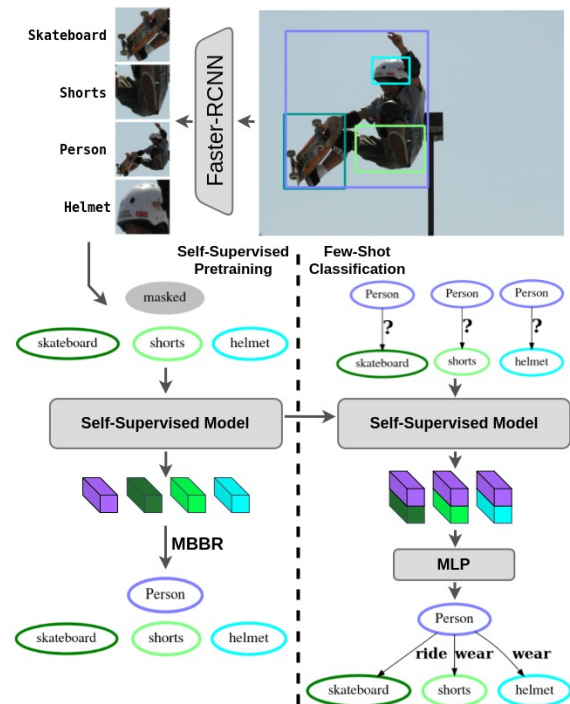


Figure 1. High-level overview of Masked Bounding Box Reconstruction (MBBR), a novel pretext task for self-supervised VRD. Visual representations for entities within a scene are first extracted through a Faster-RCNN object detector. These visual features are then randomly masked and fed into a transformer encoder. Masked representations are reconstructed, conditioned on the unmasked ones. Learned representations are highly predictive of visual relationships and result in strong Predicate Detection performance in a few-shot classification setting.

components or objects, to extracting the relationships between detected entities. A relationship can be defined as a triplet in the form of $\langle S, P, O \rangle$, which indicates that the subject S correlates with the object O through the pred-

icate P , which is usually selected from a finite set of possible predicate classes. In the example of Fig. 1, subject [PERSON] correlates with object [SKATEBOARD] through the predicate [RIDE]. Thus, a scene can be represented in a structured way, *i.e.*, a directed graph where the nodes are the detected entities, while the edges represent the relations among them with a direction from subject to the object.

The emergence of large datasets providing relationship-level annotations has enabled supervised learning as the dominant approach for visual relationship detection [4, 20, 21, 37, 39, 42, 44]. However, labeled data for VRD is difficult to obtain while annotating large-size datasets is a quite time-consuming and expensive procedure. Consider that potential per-image relationships can grow exponentially to the number of objects in the scene, requiring multiple annotations to accurately capture the complex web of possible interactions and associations. Given the apparent cost of data annotation, existing datasets only capture the relationships between a few object categories mostly centered around the human-sensing domain, thus limiting a wider range of potential applications on novel object categories. Finally, defining relationships through distinct predicate classes can introduce various ambiguities (multiple predicate classes can be representative of a visual relationship), human annotators are heavily biased and predicate classes commonly demonstrate a long tail distribution on existing datasets [38]. Given the identified limitations of the supervised approach, this work focused on the self-supervised learning (SSL) paradigm for visual relationship detection. SSL involves training models using a pretext task where supervision is provided from the data without requiring any manual annotations. Despite impressive performance in both natural language processing [2, 5] and various computer vision tasks [1, 9, 18, 23, 31], self-supervised approaches for VRD have attracted less attention.

Our proposed method comprises of a 2-stage pipeline. Initially, a model is pre-trained in a self-supervised manner on large-scale datasets and is able to learn meaningful representations which are later fine-tuned and demonstrate strong performance in few-shot relationship classification. At the core of our architecture lies a transformer encoder that learns to capture object relationships without requiring any manually annotated labels. The transformer is fed object-level visual features, extracted from bounding boxes using Faster-RCNN [29]. We propose *Masked Bounding Box Reconstruction* (MBBR), a pretext task particularly for VRD, where a percentage of the input features is randomly masked, and the model is charged with predicting the masked representations solely based on the context provided by the unmasked object features. The key insight is that through MBBR, the model learns representations that encapsulate the complex relationships between objects. These context-aware representations, garnered from the in-

terplay of objects within a scene, prove to be highly predictive of predicates, thereby enabling effective VRD (framework overview in Fig. 1).

The performance of learned representations is evaluated on few-shot classification. A VRD classifier is trained for a k -shot setting, on top of the features derived from our trained transformer encoder. We show that a simple 2-layer Multi-Layer Perceptron (MLP) can achieve state-of-the-art results for few-shot VRD, thus demonstrating the effectiveness of the proposed SSL framework.

In summary, our contributions are: (1) We propose MBBR, a novel self-supervised pretext task particularly tailored for self-supervised VRD pre-training, that does not require manual relationship-level annotation. (2) Our proposed method can learn rich relationship-aware object representations that can be used for downstream predicate detection in a few-shot setting. (3) We extensively evaluate our proposed approach both quantitatively and qualitatively on both VRD and VG200 datasets, achieving large performance improvements for few-shot Visual Relationship Detection.

2. Related Work

2.1. Visual Relationship Detection

The VRD task was originally introduced in [30], where authors treat each triplet $\langle S, P, O \rangle$ as an individual class for which a separate classifier is trained. This early formulation required training of $O(N^2K)$ classifiers (where N, K is the number of entities and predicates respectively), which was later [22] reduced to $O(NK)$ by training a single classifier per predicate class. Lu *et al.* [22] proposed the addition of linguistic priors to exploit entity semantics, leading to enhanced predicate classification. In [41], authors conduct a statistical analysis on the Visual Genome (VG) dataset [38] and discover valuable motifs, *i.e.*, repeated structures and recurring patterns across large scene graphs, which exploit later using bidirectional LSTMs.

VTransE, proposed in [43], is a visual translation embedding network which projects the subject S_{emb} , object O_{emb} and predicate P_{emb} representations into a low dimensional embedding vector space where $S_{emb} + P_{emb} \approx O_{emb}$. Following [43], authors in [15] introduced UVTransE based on the observation that the subtraction of S_{emb} and O_{emb} from the embedding of the union of subject and object, results in the predicate embedding, *i.e.*, $U_{emb} - S_{emb} - O_{emb} \approx P_{emb}$. In [11] multi-head attention is employed, with a separate head for each predicate class, thus the model is able to attend on multiple visual regions of the input image. Inspired by Faster-RCNN [29] propose Graph-RCNN [40], a graph network which utilizes a relationship proposal network (RePN) and an attentional graph convolutional network (aGCN), to al-




Figure 2. Overview of the proposed Masked Bounding Box Reconstruction framework. Given an input image, we first extract visual features for each entity in a scene, using a pre-trained Faster-RCNN detector. A percentage of the object embeddings is randomly masked. The network is trained to reconstruct masked objects, from the context provided by unmasked entities.

low contextual information sharing among objects. Compared to these methods, our proposed approach does not rely on relationship-level supervision for model training. Relationship-aware object representations are learned instead in a self-supervised manner through Masked Bounding Box Reconstruction.

Only a few works explore visual relationship detection under reduced manual supervision. Authors in [33] introduce an energy-based loss for VRD and conduct a few-shot evaluation on relationship triplets. For our few-shot formulation in contrast, a classifier is trained with k samples per predicate class (instead of relationship triplets). To enhance performance under long-tailed distribution for predicate classes of popular VRD datasets [38], authors in [8] propose a two-stage approach. Pre-training is originally performed on the 25 most frequent classes, followed by fine-tuning through few-shot VRD classification on the 25 remaining classes. Predicates are learned as functions, which are then utilized as message-passing mechanisms within a graph convolutional network (GCN). Since [8] focuses on the 25 most frequent predicates, it remains a supervised approach, in contrast to our self-supervised pre-training task that requires no manual annotations. To the best of our knowledge, we are the first to propose a fully self-supervised pre-training pretext task based on bounding box reconstruction, specifically tailored for few-shot Visual Relationship Detection.

2.2. Self-Supervised Learning

SSL is the main learning approach used in this work for learning relationship-aware entity representations. It leverages inherent structures and patterns of the data to

learn meaningful features [7, 10, 25, 27], thus circumventing the need for manual annotations. Several techniques have been proposed in the literature, with contrastive learning and Masked Image Modeling (MIM) recently attracting the most attention. Contrastive methods like SimCLR [3] or MoCo [13] can learn strong representation by comparing augmented versions of the same image (positive pairs) with distinct images (negative samples). MIM based methods learn through masking a part of an image that is subsequently reconstructed. He *et al.* [12] propose to mask a significant portion of an image (75% of the image patches) and use autoencoders for patch reconstruction. In BEit [1], the model is pre-trained on discrete visual tokens from a randomly masked section of the image, that are obtained from the latent representations of discrete variational autoencoders [28]. Related to ours is the pretext task proposed in [6] for image manipulation instead of predicate detection. Compared to our work, authors in [6] explore a complex network architecture that also leverages semantic features and perform reconstruction in the pixel space instead of the feature space, utilizing generative networks. In this work, motivated by the recent success of MIM for representation learning based on the transformer architecture, we propose an SSL pipeline based on Masked Bounding Box Reconstruction, particularly for VRD representation learning.

3. Approach

3.1. Problem Definition

Let $D = \{I_i, b_{\langle i,j \rangle}^S, b_{\langle i,j \rangle}^O, S_{\langle i,j \rangle}, O_{\langle i,j \rangle}, P_{\langle i,j \rangle}\}$ be a dataset of N images with $i \in [1, N]$, where $b_{\langle i,j \rangle}^S$ and $b_{\langle i,j \rangle}^O$ are the subject / object bounding boxes and $S_{\langle i,j \rangle}, O_{\langle i,j \rangle}$,

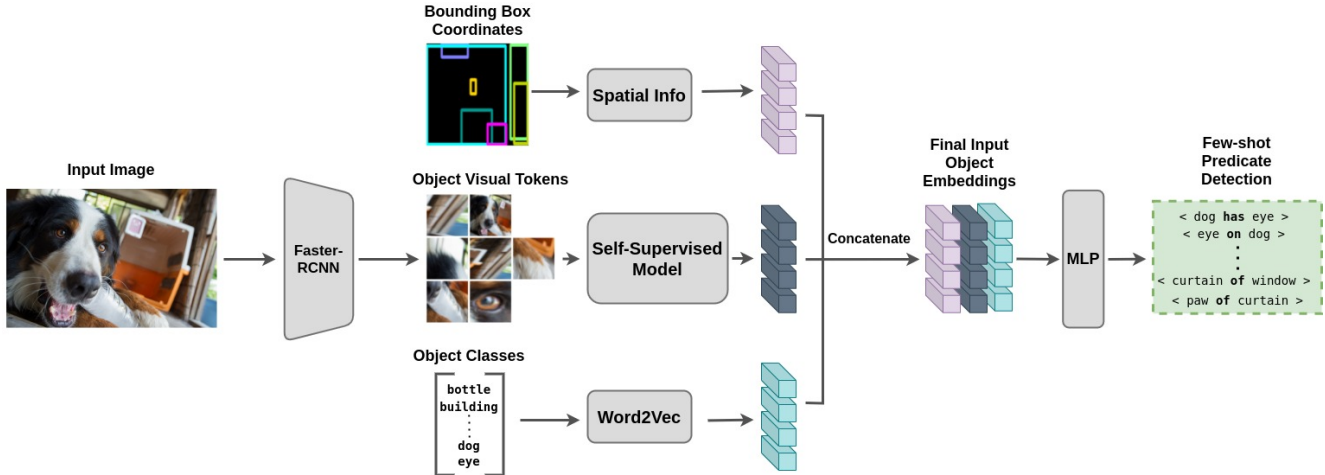


Figure 3. Architecture of the VRD k -shot classifier. Initially, an input image is passed through a pre-trained Faster-RCNN detector, extracting entity-level visual features. Our self-supervised model maps visual features into context-aware representations that capture complex object interactions. These representations along with spatial [11] and linguistic features (word2vec embeddings [24] for entity classes) are used to train k -shot classifiers for predicate detection.

$P_{\langle i,j \rangle}$ denote the subject labels, object labels and predicate labels for the j^{th} relationship in image I_i .

In this work, we focus on the Predicate Detection task (PredDet), where the goal is to learn a mapping function

$$\Psi : \{I, b^S, b^O, S, O\} \rightarrow P_k \quad (1)$$

with $k \in [1, K]$ being the set of possible predicate classes.

We propose to address the above problem through a 2-stage architecture. Firstly (Eq. 2), a deep neural network with weights θ_f learns to extract a representation z for each object on a scene in a self-supervised manner without using any object class or predicate class annotations.

$$\Psi_{\theta_f} : \{I, b\} \rightarrow z \quad (2)$$

Then in the second stage (Eq. 3), a small 2-layer MLP with parameters θ_s is used to predict

$$\Psi_{\theta_s} : \{I, z^S, z^O, S, O, b^S, b^O\} \rightarrow P_k \quad (3)$$

powered by the representations extracted for z^S and z^O during stage-one pre-training. We will describe our proposed architecture for training both Ψ_{θ_f} and Ψ_{θ_s} in the following subsection.

3.2. Self-supervised representations for VRD

Ψ_{θ_f} is trained to reconstruct the feature representation of masked objects, through the context provided by the rest of the objects in the scene (Fig. 2). Thus, learned features are highly predictive of the relationships between entities and demonstrate robust performance for few-shot predicate detection.

Learning Target Formation: To form the learning target that will be used for MBBR, we initially pass an input image through an off-the-shelf, pre-trained Faster-RCNN to extract visual features. We then apply multi-scale feature pooling for all features inside each entity bounding box b , thus extracting a single visual representation $f_b \in \mathbb{R}^{256}$ per entity of the scene.

Visual geometry embeddings: Visual geometry embeddings are also used to capture the arrangements of visual embedding within the scene. Following Hu *et al.* [14], we represent the position of each image entity by a 4-d vector $(\frac{X_{LT}}{W}, \frac{Y_{LT}}{H}, \frac{X_{RB}}{W}, \frac{Y_{RB}}{H})$, where (X_{LT}, Y_{LT}) and (X_{RB}, Y_{RB}) are the coordinates of the top left and right bottom corners of the bounding box, respectively, of the entity and W, H are the width and height of the input image. This vector is then projected into a high-dimensional space $f_{pos} \in \mathbb{R}^{256}$ by computing sine and cosine functions of different wavelengths.

Masked Bounding Box Reconstruction: We randomly mask each entity feature f_b with a probability of 50%. Features are then concatenated with geometry embeddings f_{pos} and projected through a linear layer to an entity embedding $f_e \in \mathbb{R}^{256}$. For this work, Ψ_{θ_f} (Eq. 2) is modelled as a standard feed-forward transformer encoder of [35]. The output of the transformer is a representation per object z_i that is then projected through a linear layer to form the reconstructed entity embeddings $y_{rec,i}$. Our model is trained through a mean square error (MSE) loss between reconstructed entity embeddings and the input embeddings $f_{b,i}$ from the pre-trained Faster-RCNN, for all N entities in a scene. The embeddings z_i can be used later for training a simple MLP in a few-show setting, as discussed next.

Method	Graph Constraints		No Graph Constraints	
	10-shot	20-shot	10-shot	20-shot
Faster-RCNN [29]	2.6 \pm 2.81	14.03 \pm 5.68	8.13 \pm 7.04	21.30 \pm 9.48
Motifs [41]	2.48 \pm 3.28	2.91 \pm 6.08	4.51 \pm 5.05	3.77 \pm 6.34
VTransE [43]	9.75 \pm 2.55	14.66 \pm 3.54	18.61 \pm 3.25	26.98 \pm 4.07
UVTransE [15]	10.41 \pm 3.29	15.98 \pm 2.55	20.7 \pm 5.44	30.67 \pm 0.88
ATR-Net [11]	2.05 \pm 0.58	17.88 \pm 1.94	6.53 \pm 1.24	31.38 \pm 2.21
Our method	20.87 \pm 2.46	21.52 \pm 1.34	30.75 \pm 3.66	34.01 \pm 2.51

Table 1. Comparison of our method with state-of-the-art on $\{10, 20\}$ -shot predicate detection on VRD dataset [22]. We report $R@20$ and show the mean value and standard deviation of 5 random initializations.

Method	Graph Constraints		No Graph Constraints	
	10-shot	20-shot	10-shot	20-shot
Faster-RCNN [29]	1.82 \pm 1.76	3.83 \pm 2.80	7.91 \pm 3.80	10.25 \pm 6.79
Motifs [41]	2.62 \pm 4.60	2.39 \pm 3.73	6.83 \pm 3.6	7.99 \pm 4.15
VTransE [43]	7.07 \pm 3.3	7.90 \pm 2.35	13.70 \pm 4.18	17.64 \pm 4.30
UVTransE [15]	5.45 \pm 2.91	8.87 \pm 4.29	12.74 \pm 5.5	20.24 \pm 8.55
ATR-Net [11]	0.52 \pm 0.43	3.29 \pm 4.72	2.79 \pm 1.82	8.46 \pm 6.73
Our method	8.02 \pm 1.32	15.37 \pm 2.27	16.71 \pm 3.76	28.87 \pm 2.30

Table 2. Comparison of our method with state-of-the-art on $\{10, 20\}$ -shot predicate detection on VG200 dataset [38]. We report $R@20$ and show the mean value and standard deviation of 5 random initializations.

3.3. Few-shot Classification

After self-supervised pre-training, we keep only the encoder without the projection head, and treat the derived representations z_i as visual features for training a classifier on few-shot VRD (Fig. 3). In addition to the visual features, we employ linguistic features, *i.e.*, the word2Vec [24] embeddings of each entity’s label, and spatial features [11] for modeling the respective location of the subject-object pairs. The few-shot classifier takes as input the concatenation of the above features and is trained for predicate class prediction using a standard cross-entropy loss.

4. Experiments

4.1. Datasets and Metrics

Evaluation is performed on two commonly used publicly available datasets, namely VRD [22] and VG [38]. The first is a widely used for visual relationship detection. It contains 5,000 images with 100 object and 70 relationship categories. We use the same split as [22], *i.e.*, 4,000 training images and 1,000 test images. The total number of annotated triplets is 203,284 in the training set and 7,624 in the test set.

Visual Genome is one of the largest datasets in visual relationship detection. It contains 108,077 images, 3.8 million annotated objects and 2.3 million annotated triplets. We follow the same train/test split as in [38], *i.e.*, 75,651 training images and 32,422 testing images with 150 object classes and 50 relation classes.

In this paper, we focus on the PredDet task. Our proposed SSL pipeline is evaluated in a few-shot setting. Note that in our formulation, k -shot refers to k samples per predicate class. Thus, for 10-shot evaluation, in VRD we will use 70×10 relationships, where 70 is the number of predicate classes. We use $Recall_k@N$ ($R_k@N$) as our evaluation metric. Given an input image with x subject-object pairs, $R_k@N$ considers only the top- k predictions for each pair and then selects the N most confident out of a total of $x \cdot k$ predictions. Following [11], we refer to evaluation with $k = 1$ as *graph constraints*, indicating that only one edge between entities is allowed. Larger values of k are signified as *no graph constraints*, allowing multiple edges between entities. In this work, when referring to *no graph constraints*, k is set to 50 and 70 for the VG200 and VRD, respectively.

4.2. Implementation details

The transformer encoder we use is comprised of 8 attention heads, 6 layers and a feature dimension of 256. We use the Adam optimizer [17] with a base learning rate 2×10^{-3} and weight decay of 10^{-4} . The model is trained for 30 epochs with batch size 16. After pre-training, we further fine-tune our model on a k -shot setting for 20 more epochs. For all reported experiments, self-supervised pre-training is performed on VG200 [38] as in VRD there are not enough images for our model to learn useful representations. All of our models are implemented in PyTorch.

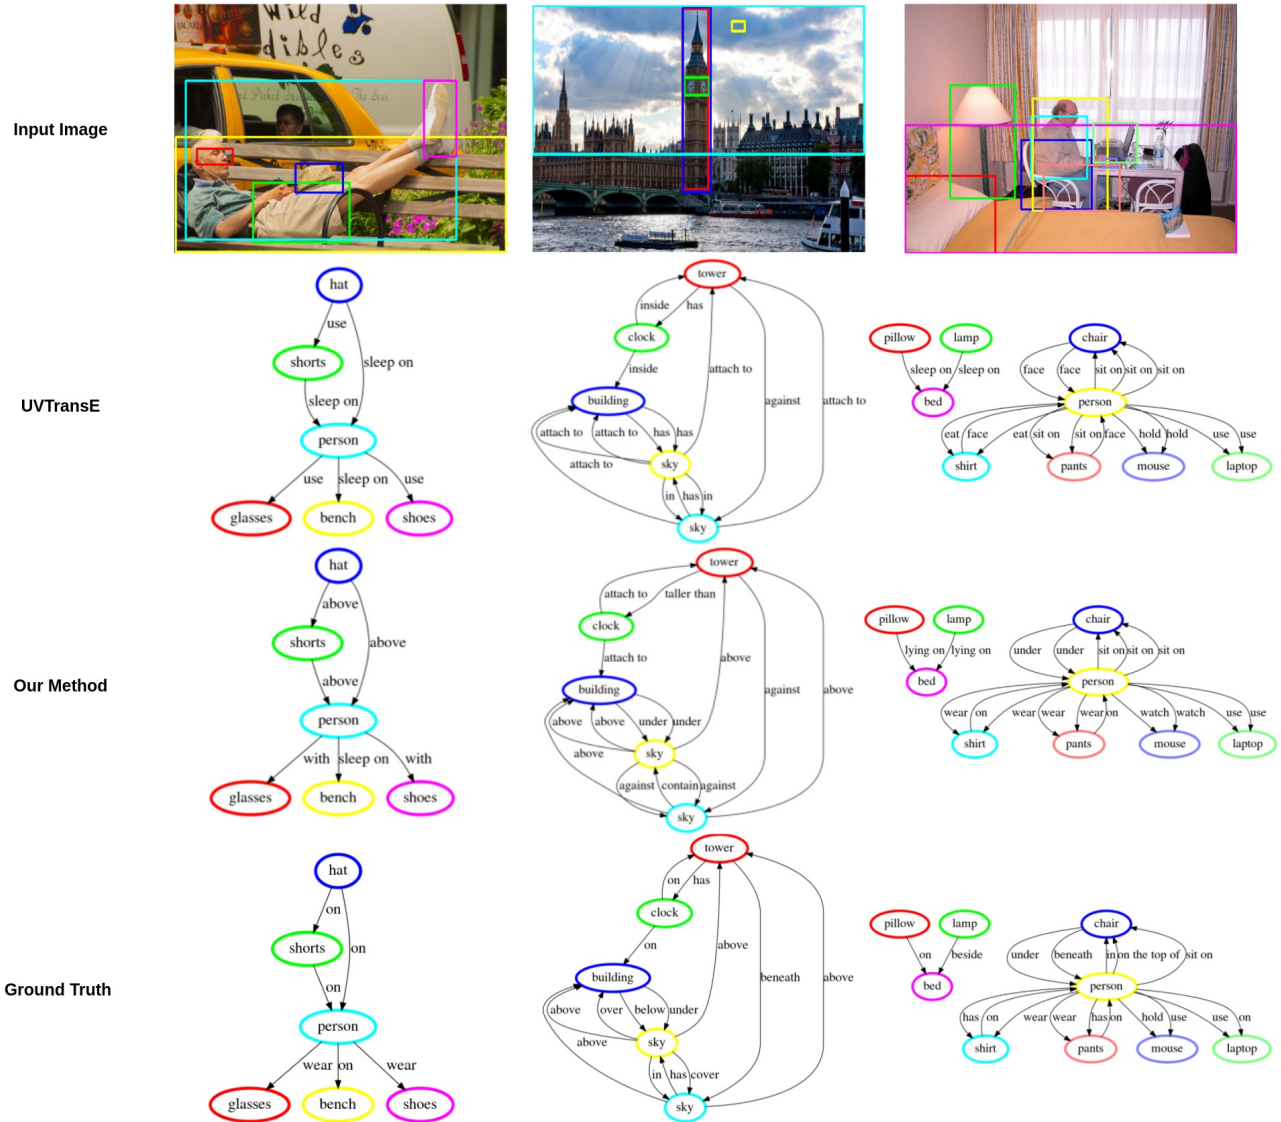


Figure 4. Directed graphs extracted from 10-shot classifiers on the VRD dataset [22]. We provide a qualitative comparison of our method and UVtransE [15]. Even in cases where our method does not recover the ground truth triplet, predicted predicates remain semantically accurate. For example see groundtruth relationship $\langle \text{BUILDING}, \text{BELOW}, \text{SKY} \rangle$ (*second column*). Our method recovers a plausible triplet $\langle \text{BUILDING}, \text{UNDER}, \text{SKY} \rangle$ compared to $\langle \text{BUILDING}, \text{HAS}, \text{SKY} \rangle$ for UVTransE.

Method	Graph Constraints			No Graph Constraints		
	1-shot	2-shot	5-shot	1-shot	2-shot	5-shot
Faster-RCNN [29]	4.23 \pm 2.87	4.24 \pm 2.4	5.2 \pm 5.34	9.65 \pm 2.36	10.52 \pm 2.12	12.6 \pm 4.55
Motifs [41]	0.1 \pm 0.1	1.15 \pm 2.35	9.62 \pm 7.97	2.28 \pm 2.54	2.36 \pm 2.55	11.76 \pm 8.17
VTransE [43]	7.46 \pm 1.55	9.43 \pm 2.69	12.10 \pm 1.45	14.43 \pm 1.05	16.6 \pm 3.90	23.44 \pm 1.63
UVTransE [15]	4.47 \pm 2.84	7.52 \pm 2.88	12.6 \pm 0.98	10.95 \pm 5.95	14.88 \pm 4.59	23.62 \pm 1.37
ATR-Net [11]	1.43 \pm 1.30	1.17 \pm 0.62	2.67 \pm 0.84	3.96 \pm 1.92	5.28 \pm 2.56	8.75 \pm 0.69
Our method	9.53\pm1.88	11.98\pm2.52	19.90\pm1.00	16.92\pm3.14	19.22\pm3.18	29.92\pm2.83

Table 3. Comparison of our method with state-of-the-art on $\{1, 2, 5\}$ -shot predicate detection on VRD dataset [22]. For this experiment, samples used for few-shot learning are selected manually to ensure effective training.

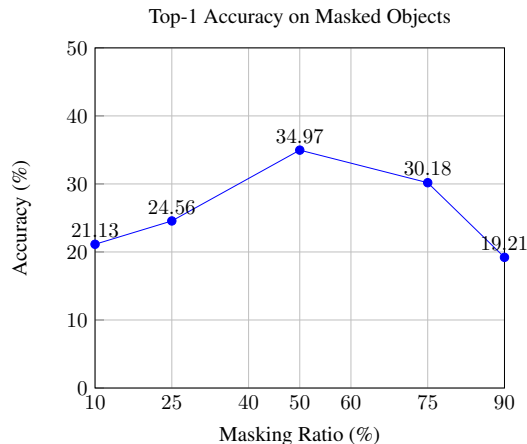


Figure 5. *Top-1* classification accuracy of reconstructed objects (*masked*), learned by our self-supervised encoder, for various masking ratio values.

4.3. Results

In this section, we analyze the effectiveness of our proposed approach, both quantitatively and qualitatively. Since self-supervised representation learning for predicate detection is a previously unexplored area, comparisons are performed w.r.t. recent supervised methods, trained on a few-shot setting [11, 15, 41, 43]. We also include a *Faster-RCNN* baseline where instead of using z_i derived from our pre-trained SSL model, we use the object-level visual features that are extracted from the *Faster-RCNN*.

Tables 1 and 2 summarize the obtained results on both datasets. Evaluation for PredDet is performed in a 10 and 20-shot setting. We observe that through self-supervised pertaining, our encoder can learn robust and generalizable representations that surpass both recent supervised methods and the *Faster-RCNN* baseline in the few-shot setting, thus demonstrating the effectiveness of MBBR.

Additional few-shot results are provided on Table 3 for {1, 2, 5}-shots. For this evaluation, we opt to manually select the few accurate relationships that are used to train our classifiers. The reason is that relationship tuples in both VRD and VG200 can be highly noisy [43], e.g., $\langle \text{SKY}, \text{HAS}, \text{SKY} \rangle$ and learned classifiers might fail to generalize when trained on a very small number of noisy examples. As previously mentioned, we find that our pre-trained model surpasses all related methods by a large margin. Interestingly, the 5-shot setting in Table 3 results in a similar performance to 20-shots in Table 1, thus further demonstrating the importance of selecting accurate relationships for few-shot classification.

5. Ablation Studies

We conduct an ablation study to further investigate MBBR as a pretext task for VRD representation learning.

Method	Graph Constraints		No Graph Constraints	
	10-shot	20-shot	10-shot	20-shot
L + S	13.68 \pm 2.27	16.02 \pm 3.39	20.39 \pm 1.22	26.30 \pm 3.87
L + S + V	20.87 \pm 2.46	21.52 \pm 1.34	30.75 \pm 3.66	34.01 \pm 2.51

Table 4. Investigation on the impact of linguistic / spatial features. We compare a classifier trained with linguistic / spatial features (**L+S**) only, to our full model (**L+S+V**) utilising self-supervised visual representations (**V**). Results for {10, 20}-shot predicate detection on the VRD dataset [22]. We report $R@20$ and show the mean value and standard deviation of 5 random initializations.

When evaluating the effectiveness of entity representations learned through unsupervised pretraining, we report *top-1* classification accuracy, for a classifier trained on the reconstructed embeddings $y_{rec,i}$ of masked input objects. Comparison is performed against a baseline classifier trained on *Faster-RCNN* pooled features $f_{b,i}$.

Masking Ratio. We start by investigating the effect of the entity masking ratio used during self-supervised pretraining (Fig. 5). We find that a masking ratio of 50% results in optimal performance measured in terms of classification accuracy of reconstructed masked objects. Intuitively, a much larger ratio, i.e., 75% or 90% degrades performance since only a few objects remain unmasked and thus provide context for masked object reconstruction. Interestingly, a very small masking ratio of 10% also results in limited performance. Since our encoder is learned through the reconstruction of all object features (not only the masked ones), masking only a small percentage of input objects enables the network to focus on unmasked object reconstruction. Note that as a baseline, the classification of *Faster-RCNN* representation leads to a *top-1* accuracy of 74.9% (compared to the 34.9% achieved by our method). Even though our model cannot reach perfect reconstruction of masked objects (compared to an unmasked baseline), we see that learned representations are highly predictive of object relationships and thus achieve strong performance for downstream predicate detection (Tables 1 and 2).

Image vs Feature Masking. For our proposed method, an image is first passed through a *Faster-RCNN* to extract input features, a percentage of which is then masked for MBBR. This is in contrast to *VLBert* [31] where masking is performed on the original input image, i.e., before visual feature extraction. We opt for masking after *Faster-RCNN* given that for common VRD databases [38] large parts of objects on the scene tend to overlap (see qualitative results in Fig. 4). Thus, masking a larger object before feature extraction will also affect the visual features of all overlapping objects. In our experiments, we observed a 13.63% *top-1* accuracy, when a 15% masking ratio was applied before *Faster-RCNN* and 22.78% when it was applied after. Larger masking ratios applied before *Faster-RCNN* rapidly degrade the effectiveness of MBBR.

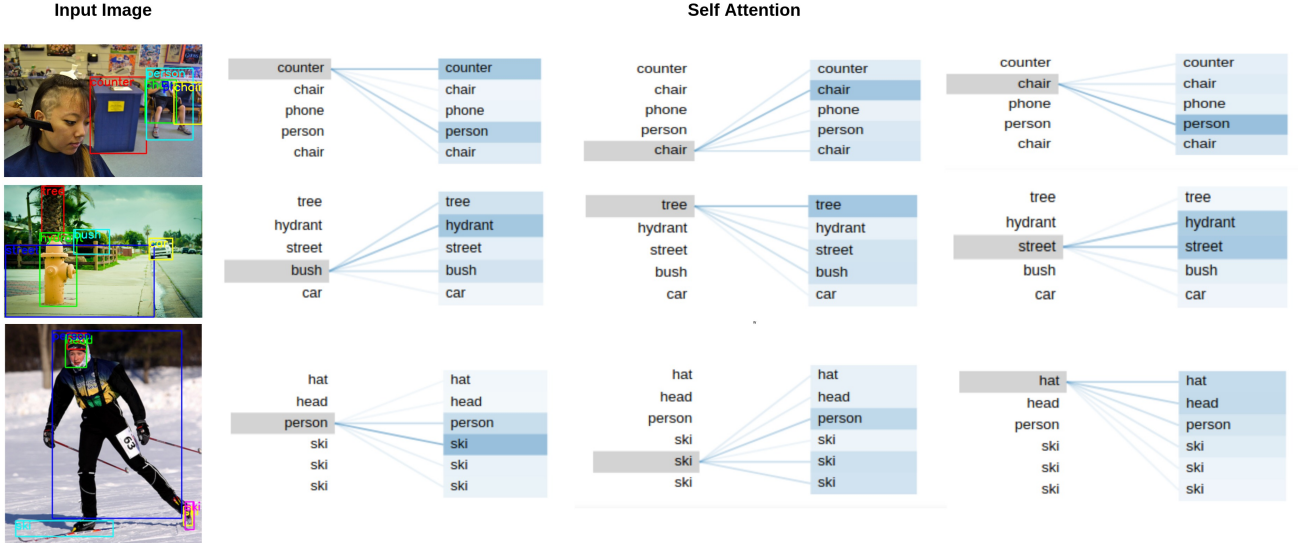


Figure 6. Visualization of self-attention scores (as in [36]) for our self-supervised transformer model trained with MBBR. We observe that for the effective reconstruction of a specific object, the model attends to other entities with which the reconstructed object is associated through a visual relationship. For instance, the object [PERSON] (*third row*) directs its attention to the object [SKI] with which has a visual relationship, while it shows no attention towards the objects [HAT] or [HEAD] with which it has no association.

Reconstruction vs Classification Loss. Our model is trained to reconstruct inputs through an MSE loss on feature space (similar to the commonly used perceptual loss [16]), achieving a $R@20$ score of $20.87_{\pm 2.46}$ for 10-shot predicate detection with graph constraints on the VRD dataset [22] (Table 1 for 5 random initializations). An alternative approach could be learning representations through the prediction of the masked object’s class, as in [31]. We find that a classification loss degrades representation learning performance, with a $R@20$ score of $16.7_{\pm 1.51}$ on the same task.

Impact of linguistic / spatial features. For few-shot predicate classification, our method utilises not only visual representations \mathbf{V} (learned through MBBR), but also linguistic \mathbf{L} and spatial \mathbf{S} features (as in [11]). To quantify the impact of linguistic and spatial features in overall model performance, we compare our complete model ($\mathbf{L+S+V}$) to a variant that only utilises language / spatial features (denoted as $\mathbf{L+S}$). In Table 4, we show that the addition of visual representation leads to superior performance, with an approximate 8% increase in 10-shot and 5.5% increase in 20-shot in $R@20$ compared to the variant that does not perceive visual information.

Self-Attention Scores. Through MBBR we learn object representations that are highly predictive of object relationships. To gain further insight into the entity relationships discovered by our self-supervised approach, we provide self-attention visualizations. In Fig. 6 we observe entities tend to attend to other subjects/objects in the scene on which they are connected with a visual relationship. For

example in row 3 of Fig. 6, the [SKI]’s attend more to the [PERSON] and other [SKI]’s for representation prediction and less to the person’s [HEAD] or [HAT].

Qualitative Results. We also include a qualitative comparison with the UVTransE method of [15] of Fig. 4 on 10-shot setting. Even in cases where our model’s predicate prediction differs from ground-truth, we can see that the detected predicate is semantically more accurate compared to the corresponding predictions of UVTransE. For example see $\langle \text{HAT, SLEEP ON, PERSON} \rangle$ for UVTransE compared to our prediction of $\langle \text{HAT, ABOVE, PERSON} \rangle$ for a ground-truth relationship of $\langle \text{HAT, ON, PERSON} \rangle$. Overall, our proposed method can produce accurate predicate predictions even when fine-tuned with a very small number of relationship samples.

6. Conclusion

We have introduced a self-supervised learning method for visual relationship detection based on Masked Bounding Box Reconstruction (MBBR). MBBR efficiently learns rich context-aware representations that are highly predictive of object visual relationships. Our experiments demonstrate that learned representations can surpass existing methods for downstream VRD in a few-shot setting.

Acknowledgements This project has been funded by `deeplab.ai`, as part of its research activities, *i.e.*, funding of student research-training and, collaborations with academic institutions. Work was conducted in part, while Z.Anastasakis was an intern with `deeplab.ai`.

References

- [1] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: BERT pre-training of image transformers. In *ICLR*, 2022. 2, 3
- [2] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. 2
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 3
- [4] Yuren Cong, Hanno Ackermann, Wentong Liao, Michael Ying Yang, and Bodo Rosenhahn. *NODIS: Neural Ordinary Differential Scene Understanding*. 11 2020. 2
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *NAACL*, 2019. 2
- [6] H. Dhama, A. Farshad, I. Laina, N. Navab, G. D. Hager, F. Tombari, and C. Rupprecht. Semantic image manipulation using scene graphs. In *CVPR*, Los Alamitos, CA, USA, jun 2020. IEEE Computer Society. 3
- [7] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 3
- [8] Apoorva Dornadula, Austin Narcomey, Ranjay Krishna, Michael S. Bernstein, and Li Fei-Fei. Visual relationships as functions: Enabling few-shot scene graph prediction. In *ICCV Workshops*, 2019. 3
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2
- [10] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 3
- [11] Nikolaos Gkanatsios, Vassilis Pitsikalis, Petros Koutras, and Petros Maragos. Attention-translation-relation network for scalable scene graph generation. In *ICCV Workshops*, 2019. 2, 4, 5, 6, 7, 8
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 3
- [13] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. *CVPR*, 2019. 3
- [14] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. 11 2017. 4
- [15] Zih-Siou Hung, Arun Mallya, and Svetlana Lazebnik. Contextual translation embedding for visual relationship detection and scene graph generation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2021. 2, 5, 6, 7, 8
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. *arXiv*, 2016. 8
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 5
- [18] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv*, 2019. 2
- [19] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, 2017. 1
- [20] Xin Lin, Changxing Ding, Jinquan Zeng, and Dacheng Tao. Gps-net: Graph property sensing network for scene graph generation. In *CVPR*, 2020. 2
- [21] Xin Lin, Changxing Ding, Yibing Zhan, Zijian Li, and Dacheng Tao. Hl-net: Heterophily learning network for scene graph generation. In *CVPR*, 2022. 2
- [22] Cewu Lu, Ranjay Krishna, Michael S. Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *ECCV*, 2016. 2, 5, 6, 7, 8
- [23] Dimitrios Mallis, Enrique Sanchez, Matt Bell, and Georgios Tzimiropoulos. From keypoints to object landmarks via self-training correspondence: A novel approach to unsupervised landmark discovery. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2
- [24] Tomáš Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2013. 4, 5
- [25] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *ECCV*, 2016. 3
- [26] Maria Pirelli, Dimitrios Mallis, Markos Diomataris, and Vassilis Pitsikalis. Interpretable visual question answering via reasoning supervision. *ICIP*, 2023. 1
- [27] Deepak Pathak, Ross B. Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *CVPR*, 2017. 3
- [28] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In Marina Meila and Tong Zhang, editors, *ICML*, 2021. 3
- [29] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 2, 5, 6
- [30] Mohammad Amin Sadeghi and Ali Farhadi. Recognition using visual phrases. In *CVPR*, 2011. 2

- [31] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: pre-training of generic visual-linguistic representations. In *ICLR*, 2020. 2, 7, 8
- [32] Mohammed Suhail, Abhay Mittal, Behjat Siddiquie, Chris Broaddus, Jayan Eledath, Gérard G. Medioni, and Leonid Sigal. Energy-based learning for scene graph generation. In *CVPR*, 2021. 1
- [33] Mohammed Suhail, Abhay Mittal, Behjat Siddiquie, Chris Broaddus, Jayan Eledath, Gérard G. Medioni, and Leonid Sigal. Energy-based learning for scene graph generation. In *CVPR*, 2021. 3
- [34] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *CVPR*, 2019. 1
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, 2017. 4
- [36] Jesse Vig. A multiscale visualization of attention in the transformer model. In *ACL*, Florence, Italy, July 2019. Association for Computational Linguistics. 8
- [37] Sangmin Woo, Junhyug Noh, and Kangil Kim. Tackling the challenges in scene graph generation with local-to-global interactions. *IEEE transactions on neural networks and learning systems*, 2021. 2
- [38] Danfei Xu, Yuke Zhu, Christopher B. Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *CVPR*, 2017. 2, 3, 5, 7
- [39] Gengcong Yang, Jingyi Zhang, Yong Zhang, Baoyuan Wu, and Yujiu Yang. Probabilistic modeling of semantic ambiguity for scene graph generation. In *CVPR*, 2021. 2
- [40] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph R-CNN for scene graph generation. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *ECCV*, 2018. 2
- [41] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *CVPR*, 2018. 2, 5, 6, 7
- [42] Ao Zhang, Yuan Yao, Qianyu Chen, Wei Ji, Zhiyuan Liu, Maosong Sun, and Tat-Seng Chua. Fine-grained scene graph generation with data transfer. In *ECCV*, 2022. 2
- [43] Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, and Tat-Seng Chua. Visual translation embedding network for visual relation detection. In *CVPR*, 2017. 2, 5, 6, 7
- [44] Ji Zhang, Kevin J. Shih, Ahmed Elgammal, Andrew Tao, and Bryan Catanzaro. Graphical contrastive losses for scene graph parsing. In *CVPR*, 2019. 2