

VL-SAT: Visual-Linguistic Semantics Assisted Training for 3D Semantic Scene Graph Prediction in Point Cloud

Ziqin Wang¹, Bowen Cheng¹, Lichen Zhao¹, Dong Xu², Yang Tang^{3*}, Lu Sheng^{1*}

¹School of Software, Beihang University

²The University of Hong Kong, ³East China University of Science and Technology

{wzqin, chengbowen052, zlc1114, lsheng}@buaa.edu.cn dongxu@cs.hku.hk yangtang@ecust.edu.cn

Abstract

The task of 3D semantic scene graph (3DSSG) prediction in the point cloud is challenging since (1) the 3D point cloud only captures geometric structures with limited semantics compared to 2D images, and (2) long-tailed relation distribution inherently hinders the learning of unbiased prediction. Since 2D images provide rich semantics and scene graphs are in nature coped with languages, in this study, we propose Visual-Linguistic Semantics Assisted Training (VL-SAT) scheme that can significantly empower 3DSSG prediction models with discrimination about long-tailed and ambiguous semantic relations. The key idea is to train a powerful multi-modal oracle model to assist the 3D model. This oracle learns reliable structural representations based on semantics from vision, language, and 3D geometry, and its benefits can be heterogeneously passed to the 3D model during the training stage. By effectively utilizing visual-linguistic semantics in training, our VL-SAT can significantly boost common 3DSSG prediction models, such as SGFN and SGG_{point}, only with 3D inputs in the inference stage, especially when dealing with tail relation triplets. Comprehensive evaluations and ablation studies on the 3DSSG dataset have validated the effectiveness of the proposed scheme. Code is available at <https://github.com/wz7in/CVPR2023-VLSAT>.

1. Introduction

Structurally understanding 3D geometric scenes is particularly important for tasks that require interaction with real-world environments, such as AR/VR [7, 21, 26–28, 49, 51] and navigation [4, 5, 10]. As one vital topic in this field, predicting 3D semantic scene graph (3DSSG) in point cloud [36] has received emerging attention in recent years. Specifically, given the point cloud of a 3D scene that is associated with class-agnostic 3D instance masks, the 3DSSG

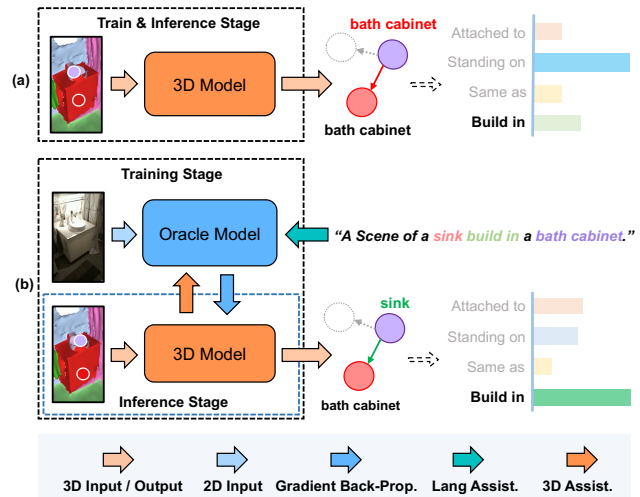


Figure 1. **Comparison between previous method and our VL-SAT.** (a) SGPN [36], as the 3D model, fails to find capture predicates such as *build in*. (b) VL-SAT creates an oracle model by heterogeneously fusing 2D semantics, and language knowledge along with the geometrical features, and the 3D model receives benefits from the oracle model during training. During inference, the enhanced 3D model can correctly detect the tail predicates.

prediction task would like to construct a directed graph whose nodes are semantic labels of the 3D instances and the edges recognize the directional semantic or geometrical relations between connected 3D instances.

However, in addition to common difficulties faced by scene graph prediction, there are several challenges specified to the 3DSSG prediction task. (1) 3D data such as point clouds only capture the geometric structures of each instance and may superficially define the relations by relative orientations or distances. (2) Recent 3DSSG prediction datasets [36, 47] are quite small and suffer from long-tailed predicate distributions, where semantic predicates are often rarer than geometrical predicates. For example, as shown in Fig. 1, the pioneering work SGPN [36] usually

*Lu Sheng and Yang Tang are the corresponding authors.

prefers a simple and common geometric predicate *standing on* between *sink* and *bath cabinet*, while the ground-truth relation triplet $\langle sink, build\ in, bath\ cabinet \rangle$ cares more about the semantics, and the frequency of *build in* in the training dataset is quite low, as shown in Fig. 3(a). Even though some attempts [38, 47, 48] have been proposed thereafter, the inherent limitations of the point cloud data to some extent hinder the effectiveness of these methods.

Since 2D images provide rich and meaningful semantics, and the scene graph prediction task is in nature aligned with natural languages, we explore using visual-linguistic semantics to assist the training, as another pathway to essentially enhance the capability of common 3DSSG prediction models with the aforementioned challenges.

How to assist 3D structural understanding with visual-linguistic semantics remains an open problem. Previous studies mainly focus on employing 2D semantics to enhance instance-level tasks, such as object detection [3, 20, 22, 29], visual grounding and dense captioning [4, 5, 44, 50]. Most of them require visual data both in training and inference, but a few of them, such as SAT [42] and \mathcal{X} -Trans2Cap [44] treat 2D semantics as auxiliary training signals and thus offer more practical inference only with 3D data. But these methods are specified to instance-level tasks and require delicately designed networks for effective assistance, thus they are less desirable to our structural prediction problem. Thanks to the recent success of large-scale cross-modal pretraining like CLIP [24], 2D semantics in images can be well aligned with linguistic semantics in natural languages, and the visual-linguistic semantics have been applied for alleviating long-tailed issue in tasks related to 2D scene graphs [1, 31, 32, 45] and human-object interaction [15]. But how to adapt similar assistance of visual-linguistic semantics to the 3D scenario remains unclear.

In this study, we propose the Visual-Linguistic Semantics Assisted Training (VL-SAT) scheme to empower the point cloud-based 3DSSG prediction model (termed as the 3D model) with sufficient discrimination about long-tailed and ambiguous semantic relation triplets. In this scheme, we simultaneously train a powerful multi-modal prediction model as the oracle (termed as oracle model) that is heterogeneously aligned with the 3D model, which captures reliable structural semantics by extra data from vision, extra training signals from language, as well as the geometrical features from the 3D model. These introduced visual-linguistic semantics have been aligned by CLIP. Consequently, the benefits of the oracle model, especially the multi-modal structural semantics, can be efficiently embedded into the 3D model through the back-propagated gradient flows. In the inference stage, the 3D model can perform superior 3DSSG prediction performance with only 3D inputs. For example, in Fig. 1(b), the predicate *build in* can be reliably detected. To our best knowledge, VL-SAT is

the first visual-linguistic knowledge transfer work that is applied to 3DSSG prediction in the point cloud. Moreover, VL-SAT can successfully enhance SGFN [38] and SGG_{point} [47], validating that this scheme is generalizable to common 3DSSG prediction models.

We benchmark VL-SAT on the 3DSSG dataset [36]. Quantitative and qualitative evaluations, as well as comprehensive ablation studies, validate that the proposed training scheme leads to significant performance gains, especially for tail relations, as discussed in Sec. 4.

2. Related Work

Scene Graph Prediction in Point Cloud. Image-based semantic scene graph prediction has been extensively studied [6, 30–32, 37, 40, 41, 46] in recent years, but only a few works try to predict 3D semantic scene graph in the point cloud. Armeni *et al.* [2] presented the first 3D scene graph dataset, which maps 3D buildings into hierarchical structures. Wald *et al.* [36] constructed a point cloud-based semantic scene graph dataset, namely 3DSSG, with a GNN-based baseline model named SGPN. The follow-up work SGFN [38] predicted 3DSSG incrementally from RGB-D sequences. In recent years, a few methods were proposed to improve the GNN-based baseline. SGG_{point} [47] used an edge-oriented graph convolution network to exploit multi-dimensional edge features for relation modeling. Zhang *et al.* [48] proposed a graph auto-encoder network to automatically learn a group of embeddings for each class in advance, and then perform the 3DSSG prediction to recognize credible relation triplets from pre-learned knowledge.

3D Scene Understanding with 2D Semantics. A list of methods have employed 2D semantics to help 3D instance-level tasks, such as 3D object detection, segmentation, visual grounding and dense captioning [3, 20, 22, 29, 34, 43, 52]. They can be coarsely divided into two categories, *i.e.*, concatenating image features with each 3D point [3–5, 20, 34, 43, 50], and projecting object detection results into 3D space [12, 19, 22, 39, 42, 44]. Most methods require 2D semantics both in the training and inference stages. Recently, SAT [42] and \mathcal{X} -Trans2Cap [44] explore using 2D semantics only in training to assist 3D visual grounding and dense captioning. Both of them can learn enhanced models that only use 3D inputs in inference. But these methods are restrained to instance-level tasks and the networks have to be carefully designed. We follow similar ideas as [42, 44] and use 2D semantics only in training, but we would like to enhance the 3DSSG prediction that requires structural understanding rather than instance-level perception.

Knowledge-inserted Methods in Scene Graph Prediction. Zellers *et al.* [46] and Chen *et al.* [6] indicated that the statistical co-occurrences between object pairs and relationships are useful for relation prediction. Besides, [25, 48] is

generated class-level prototypical representations from all previous perceptual outputs as the prior knowledge. These methods explicitly encoded the data priors into the model. [1, 13, 14, 18] attempted to combine language priors with scene graph prediction. Zareian *et al.* [45] proposed a Graph Bridging Network to propagate messages between scene graphs and knowledge graphs. Our VL-SAT scheme uses CLIP to encode the linguistic semantics, which is thus better aligned with 2D semantics, and even the required 3D structural semantics during the training stage.

3. Method

We first overview the formulation of 3D semantic scene graph (3DSSG) prediction in point cloud (Sec. 3.1) and then elaborate on a GNN-based network that we experiment on as our 3D prediction model (Sec. 3.2). Since then we highlight how our Visual-Linguistic Semantics Assisted Training (VL-SAT) scheme comprehensively transfers the benefits of an oracle multi-modal prediction model to the 3D prediction model in discriminating challenging relations (Sec. 3.3). Finally, we depict the training objective in Sec. 3.4.

3.1. Problem Formulation

Suppose we have a point cloud $\mathbf{P} \in \mathbb{R}^{N \times 3}$ with N 3D points, and a set of *class-agnostic* instance masks $\mathcal{M} = \{\mathbf{M}_1, \dots, \mathbf{M}_K\}$ that associate the point cloud \mathbf{P} with K semantic instances, as indicated by SGPN [36], we aim at predicting a 3D semantic scene graph as a directed graph $\mathcal{G} = \{\mathcal{O}, \mathcal{R}\}$. The set of objects $\mathcal{O} = \{o_i\}_{i=1}^K$ are all *named* object instances that are specified by instance masks \mathcal{M} . Each edge r_{ij} in \mathcal{R} depicts the *predicate* in a relation triplet $\langle \text{subject}, \text{predicate}, \text{object} \rangle$, where the head node o_i of this edge is the *subject* and the tail node o_j is the *object*. To be specific, o_i indicates an object label from N_{obj} semantic classes. r_{ij} is a predicate label from N_{rel} predicate classes.

3.2. 3D Prediction Model

As depicted in Fig. 2, our employed 3D prediction model shares a similar network structure as those GNN-based scene graph prediction methods, such as SGFN [38] and SGG_{point} [47], which mainly consists of node encoder, edge encoder, and scene graph reasoning modules.

Node Encoder. Based on one class-agnostic instance mask \mathbf{M}_i along with the input point cloud \mathbf{P} , we can extract the set of points \mathbf{P}_i that correspond to one semantic instance. We employ a simple PointNet [23] to extract instance-level features. The node features $\mathbf{o}_i^{3d} \in \mathbb{R}^D$ before the GNN-based scene graph reasoning are thus given by these instance-level features.

Edge Encoder. We follow the same practice as in SGFN [38] to encode the edge features for the GNN-based

scene graph reasoning. It requires calculating the differences between several attributes between the linked instances. For each instance, these attributes include the mean $\boldsymbol{\mu}$ and standard deviation $\boldsymbol{\sigma}$ of the 3D points, the size $\mathbf{b} = (b_x, b_y, b_z)$, the volume $v = b_x b_y b_z$, and the maximum side length $l = \max(b_x, b_y, b_z)$ of the bounding box. Thus the edge features $\mathbf{r}_{ij}^{3d} \in \mathbb{R}^D$ are encoded by projecting the concatenated differences of these attributes between two instances, via multi-layer perceptron (MLP) layers, *i.e.*,

$$\mathbf{r}_{ij}^{3d} = \text{MLP}(\text{cat}(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j, \boldsymbol{\sigma}_i - \boldsymbol{\sigma}_j, \mathbf{b}_i - \mathbf{b}_j, \ln \frac{l_i}{l_j}, \ln \frac{v_i}{v_j})), \quad (1)$$

where the subscript i indicates the instance \mathbf{P}_i in the head node, and the j means the instance \mathbf{P}_j in the tail node.

Scene Graph Reasoning. In our experiment, we apply a similar GNN structure as in SGFN [38], which utilizes a Feature-wise Attention (FAT) module [38] to pass messages between nodes and edges, and then gets the updated node and edge features. Each GNN module is paired with a multi-head self-attention (MHSA) module, and they are repeated for T times to extract the final node and edge features $\{\tilde{\mathbf{o}}_i^{3d}\}_{i=1, \dots, K}$ and $\{\tilde{\mathbf{r}}_{ij}^{3d}\}_{i \neq j, i, j=1, \dots, K}$. Since then, an object classifier and a predicate classifier are to predict the elements $\{o_i, r_{ij}, o_j\}$ of each possible relation triplet from the triplet features $\{\tilde{\mathbf{o}}_i^{3d}, \tilde{\mathbf{r}}_{ij}^{3d}, \tilde{\mathbf{o}}_j^{3d}\}$. These relation triplets finally construct the semantic scene graph $\mathcal{G} = \{\mathcal{Q}, \mathcal{R}\}$.

Note that the 3D prediction model does not have to strictly follow the network of SGFN [38]. More recent GNN-based models, such as SGG_{point} [47] can also be applied. In Sec. 4.4, we show that the proposed VL-SAT scheme can also enhance both baselines with significant gains, validating that our method is generalizable to common 3DSSG prediction models.

3.3. Visual-Linguistic Semantics Assisted Training

In this subsection, we elaborate on how the visual-linguistic semantics assisted training (VL-SAT) scheme can empower the 3D prediction model with sufficient discrimination about long-tailed and ambiguous semantic relation triplets. The key idea is that this discriminative power comes from auxiliarily learning a powerful multi-modal prediction model that receives structural semantics from vision, and language, as well as the 3D geometry from the 3D prediction model. The multi-modal semantics are expected to be heterogeneously aligned with the 3D semantics at the node and edge levels, and the benefits from the oracle model can be efficiently absorbed by the 3D prediction model during the training process. To be specific, we first introduce the applied multi-modal prediction model that has heterogeneous collaboration with the 3D prediction model, and then the auxiliary training strategies that boost the performance of the oracle model and eventually enhance the 3D predic-

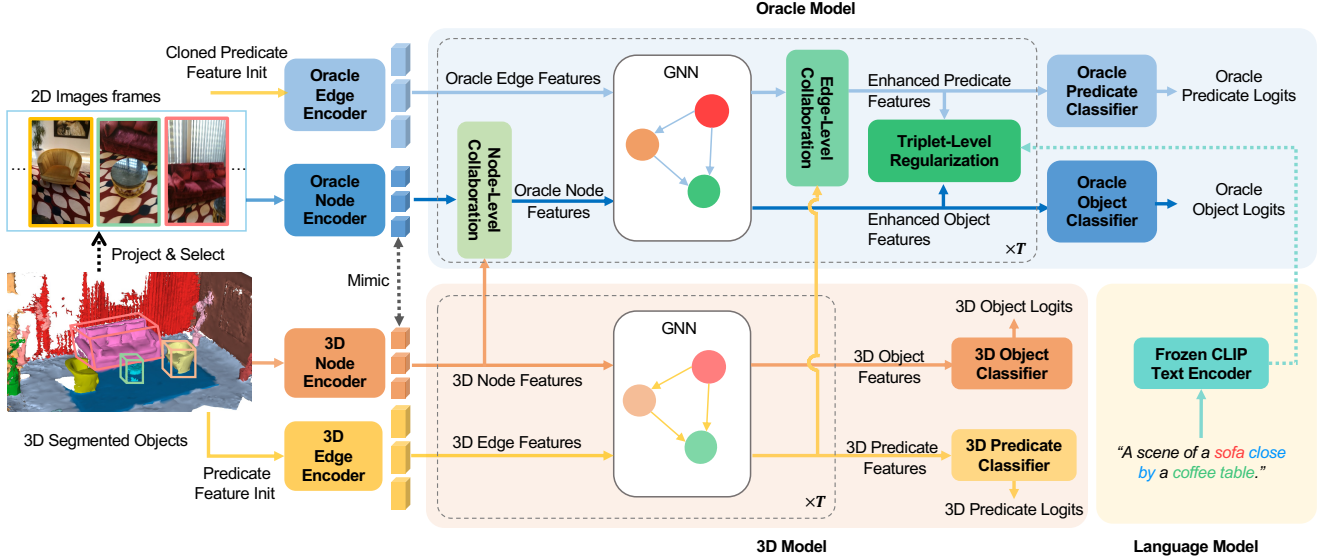


Figure 2. **The proposed Visual-Linguistic Semantics Assisted Training (VL-SAT) for 3D scene graph prediction.** In training, VL-SAT takes 2D and language semantics as extra inputs and helps 3D scene graph prediction with node and edge-level collaboration and triplet-level regularization. In inference, VL-SAT only takes the 3D point cloud to predict reliable 3D scene graphs.

tion model. We present the pipeline of VL-SAT in Fig. 2.

Multi-modal Prediction Model as the Oracle. This multi-modal prediction model acts as the oracle to our 3D prediction model. It copies the 3D prediction network in Sec. 3.2 and also learns to predict 3D semantic scene graphs, but its node features are represented by visual features. These visual features are extracted by a fixed 2D instance encoder, describing RGB image patches that are associated with each point cloud instance \mathbf{P}_i ¹. The edge features of the multi-modal prediction model are encoded in the same way in Sec. 3.2, thus still capturing 3D spatial structures in understanding relations.

The features of this oracle model heterogeneously collaborate with those in the 3D prediction model at the node and edge levels, which are conducted before and after each GNN layer in the scene graph reasoning module. The former is a node-level collaboration, and the latter is an edge-level collaboration. To be specific, these collaboration operations are implemented by multi-head cross-attention (MHCA) modules [33], where the keys and values are node/edge features from the 3D model, and the queries are their counterparts from the multi-modal model. The node-level collaboration has a distance-aware masking strategy to remove unnecessary attention between instances that are far apart without valid relations. The mask value between two instances \mathbf{P}_i and \mathbf{P}_j are learned by

$$D_{ij}^{\text{node}} = \text{MLP}(\text{cat}(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j, \|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|_2)), \quad (2)$$

¹Please refer to the supplementary for the details about how to gather associated image patches with each point cloud instance.

with respect to the mean coordinates $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_j$ of the point cloud instances \mathbf{P}_i and \mathbf{P}_j . The edge-level collaboration does not use a distance-aware masking strategy since the distance between edges is hard to define, thus it is safer to incorporate all edges into attention calculation.

Note that the heterogeneous collaboration is *unidirectional* from the 3D model to the oracle model, while the benefits of the oracle model are passed to the 3D model through the back-propagated gradient flows. It favors that in the inference stage, predicting 3D semantic scene graphs will not need extra data from other modalities.

Auxiliary Training Strategies. Since the oracle multi-modal model would like to perceive scene graphs from both the visual and linguistic perspective, it is natural to enhance the oracle model using visual-linguistic knowledge captured by CLIP [9]. Specifically, we can generate CLIP text embedding $\mathbf{e}_{ij}^{\text{text}}$ for each groundtruth relation triplet, and regularize the corresponding triplet features $\{\tilde{\mathbf{o}}_i^{\text{oracle}}, \tilde{\mathbf{r}}_{ij}^{\text{oracle}}, \tilde{\mathbf{o}}_j^{\text{oracle}}\}$ at the end of each GNN layer of the scene graph reasoning module. The CLIP text embeddings are offline extracted by the template ‘‘A scene of a/an [subject][predicate] a/an [object]’’ for each GT relation. Thus, the regularization becomes to minimize the embedding distances between the text embeddings $\mathbf{e}_{ij}^{\text{text}}$ and the fused triplet features $\mathbf{t}_{ij}^{\text{oracle}}$, *i.e.*,

$$L_{\text{tri-emb}} = \sum_{i=1}^K \sum_{j=1, j \neq i}^K \rho(\mathbf{t}_{ij}^{\text{oracle}}, \mathbf{e}_{ij}^{\text{text}}) \cdot \mathbb{I}_{[\mathbf{e}_{ij}^{\text{text}} \text{ is from GT triplet}]} \quad (3)$$

where $\mathbf{t}_{ij}^{\text{oracle}} = \text{MLP}(\text{cat}(\tilde{\mathbf{o}}_i^{\text{oracle}}, \tilde{\mathbf{r}}_{ij}^{\text{oracle}}, \tilde{\mathbf{o}}_j^{\text{oracle}}))$ is the fused

embedding of the concatenated features $\tilde{\mathbf{o}}_i^{\text{oracle}}, \tilde{\mathbf{r}}_{ij}^{\text{oracle}}$, and $\tilde{\mathbf{o}}_j^{\text{oracle}}$. $\rho(\cdot, \cdot)$ is a distance metric, we can apply ℓ_1 norm or negative cosine distance. $\mathbb{I}_{[\cdot]}$ is an indicator function that equals to 1 when the argument is true, and 0 otherwise. Thus Eq. (3) only regularizes the node and edge features whose triplets have ground-truth relations.

In addition, before being put into the scene graph reasoning modules, the 3D node features \mathbf{o}_i^{3d} from the 3D model and the 2D node features \mathbf{o}_i^{2d} from the oracle model can be aligned. We apply a same distance measurement as Eq. (3),

$$L_{\text{node-init}} = \sum_{i=1}^K \rho(\mathbf{o}_i^{3d}, \mathbf{o}_i^{2d}). \quad (4)$$

To enhance the representation ability of the initialized 2D node features, the 2D instance encoder is a fixed CLIP-pretrained vision encoder. Moreover, to enhance the object classifiers of both models, we use the CLIP object embeddings to initialize the weights of the object classifiers both at the 3D prediction model and the oracle multi-modal prediction model, as in [15, 24].

3.4. The Training Objective

The training objective of the entire network is defined as:

$$L = \lambda_{\text{obj}}(L_{\text{obj}}^{3d} + L_{\text{obj}}^{\text{oracle}}) + \lambda_{\text{pred}}(L_{\text{pred}}^{3d} + L_{\text{pred}}^{\text{oracle}}) + \lambda_{\text{aux}}(L_{\text{tri-emb}} + L_{\text{node-init}}) \quad (5)$$

L_{obj} indicates object classification loss and is implemented with cross-entropy loss. $L_{\text{obj}}^{3d/\text{oracle}}$ is applied on 3D/oracle object classifier. L_{pred} indicates predicate classification loss and is formulated as per-class binary cross-entropy loss as in [36]. $L_{\text{pred}}^{3d/\text{oracle}}$ is applied on 3D/oracle predicate classifier. $\lambda_{\text{node}}, \lambda_{\text{edge}}, \lambda_{\text{aux}}$ are hyper-parameters to balance each loss in the same scale.

4. Experiments and Discussions

4.1. Setups and Implementation Details

Datasets. We conduct experiments on 3DSSG [36]. It is a 3D semantic scene graph dataset drawn from the 3RScan dataset [35], with rich annotations about instance segmentation masks and relation triplets. It has 1553 3D reconstructed indoor scenes, 160 classes of objects, and 26 types of predicates. In the experiments, we use the same data preparation and training/validation split as in 3DSSG [36].

Metrics and Tasks. We follow the experiment settings in 3DSSG [36]. In both training and testing stages, 3D scenes are placed in the same 3D coordinate. The view-dependent spatial relation predicates are not ambiguous. To evaluate the prediction of the object and predicate, we use the top-k accuracy ($A@k$) metric. As for the triplets, we first multiply the subject, predicate, and object scores to get triplet

scores, and then compute the top-k accuracy ($A@k$) as the evaluation metric. The triplet is considered correct only if the subject, predicate, and object are all correct². To fairly evaluate the performance of long-tailed predicate distribution, we also compute the average top-k accuracy of the predicate across all predicate classes, denoted as the mean top-k accuracy ($\text{mA}@k$).

We also conduct two 2D scene graph tasks proposed in [40] in the 3D scenario, as what Zhang *et al.* [48] did, *i.e.*, (1) Scene Graph Classification (SGCls) that evaluates the triplet together. (2) Predicate Classification (PredCls) that only evaluates the predicate with the ground-truth labels of object entities. Following Zhang *et al.* [48], we compute the recall at the top-k ($R@k$) triplets. The triplet is considered correct when the subject, predicate, and object are all valid. Additionally, we also adopt mean recall ($\text{mR}@k$) to evaluate the performance on the unevenly sampled relations using a similar strategy as $\text{mA}@k$.

Implementation Details. Our network is end-to-end optimized using AdamW optimizer [11, 17] with the batch size as 8. We train the network for 100 epochs, and the base learning rate is set as 0.001 with a cosine annealing learning rate decay strategy [16]. $N_{\text{obj}} = 160$ and $N_{\text{rel}} = 26$ in our experiments. GNN modules are repeated for $T = 2$ times in both 3D and oracle multi-modal models. $\lambda_{\text{obj}} = \lambda_{\text{aux}} = 0.1$, $\lambda_{\text{pred}} = 1$ in Eq. (5). All experiments are carried out on the PyTorch platform equipped with one NVIDIA GeForce RTX 2080 Ti GPU card, and each experiment takes about 48 hours until model convergence. Note that 2D inputs are only used during the training stage. During the inference stage, we follow the same strategy in [40], which selects the top@1 class of both object and predicate while giving an object instance index tuple. Please refer to the supplementary for the details of the network structures.

4.2. Comparison with the State-of-the-art Methods

We compare our method with a list of reference methods, *i.e.* SGPN [36], SGG_{point} [47], SGFN [38], Co-Occurrence [48], KERN [6], Schemata [25], Zhang *et al.* [48]. In addition, to gain a deeper understanding of our approach, we also report the performances of the oracle multi-modal prediction model (termed as VL-SAT (oracle)), as well as the baseline performance of the 3D prediction model that is trained purely by 3D data (term as non-VL-SAT). The proposed method is term as VL-SAT.

Quantitative Results. The comparison results are summarized in Tab. 1. The baseline “non-VL-SAT” has a similar performance as SGFN. The only difference between them is that “non-VL-SAT” adds a multi-head self-attention (MHSA) module [33] before each GNN module in SGFN.

²However, the metric top-k accuracy is written as the top-k recall or $R@k$ in 3DSSG [36] and SGFN [38].

Table 1. Quantitative Results of 3D semantic scene graph prediction on the 3DSSG validation set [36]. Evaluations are conducted in terms of object, predicate, and triplet. The results of SGPN, SGG_{point}, and SGFN are based on our reproduced model with point cloud-only inputs, since they don’t compute the mA@k metric in their papers.

Model	Object			Predicate						Triplet			
	A@1	A@5	A@10	A@1	A@3	A@5	mA@1	mA@3	mA@5	A@50	A@100	mA@50	mA@100
SGPN [36]	48.28	72.94	82.74	91.32	98.09	99.15	32.01	55.22	69.44	87.55	90.66	41.52	51.92
SGG _{point} [47]	51.42	74.56	84.15	92.4	97.78	98.92	27.95	49.98	63.15	87.89	90.16	45.02	56.03
SGFN [38]	53.67	77.18	85.14	90.19	98.17	99.33	41.89	70.82	81.44	89.02	91.71	58.37	67.61
non-VL-SAT	54.79	77.62	85.84	89.59	97.63	99.08	41.99	70.88	81.67	88.96	91.37	59.58	67.75
VL-SAT (ours)	55.66	78.66	85.91	89.81	98.45	99.53	54.03	77.67	87.65	90.35	92.89	65.09	73.59
VL-SAT (oracle)	66.39	86.53	91.46	90.66	98.37	99.40	55.66	76.28	86.45	92.67	95.02	74.10	81.38




Figure 3. The line chart shows the predicate frequency in the train set of 3DSSG [36]. The bar chart shows the results on mA@1 of the predicate prediction of SGPN [36] and our VL-SAT.

Thanks to the delicate visual-linguistic assisted training scheme, our “VL-SAT” tremendously improves the baseline, according to the evaluation with respect to the predicate, and the triplet. Moreover, according to the less biased mA@k metrics with respect to long-tailed distribution, when evaluating the predicate, the proposed “VL-SAT” outperforms the baseline “non-VL-SAT” with around 12.0%, 6.8% and 6.0% gains at mA@1, mA@3, and mA@5 respectively. Our method reaches new state-of-the-art results on triplet prediction, with 6.8% gain on mA@50 and 5.9% gain on mA@100 over SGFN [38]. Note that the results of object classification just have a marginal improvement, which means that a simple PointNet-based 3D encoder may not be able to convey similar instance-level representative power as the 2D vision encoder.

As illustrated in Tab. 2 and Tab. 3, we also compare our “VL-SAT” with the reference methods, with respect to two tasks named SGClS and PredClS, according to the settings introduced by Zhang *et al.* [48]. Our method outperforms Zhang *et al.* [48] by a large margin. For example, with graph constraint [40] (as a more rigorous testing scenario [45]), “VL-SAT” has 2.5% gain on R@20 in SGClS, 8.5% gain on R@20 in PredClS. Moreover, with respect to the less biased metrics in Tab. 3, “VL-SAT” even achieves

Table 2. Quantitative results of the compared methods with respect to the SGClS and PredClS tasks, with and without graph constraint. The evaluation metric is top-k recall.

Model	SGClS	PredClS
	R@20/50/100	R@20/50/100
with Graph Constraints		
Co-Occurrence [48]	14.8/19.7/19.9	34.7/47.4/47.9
KERN [6]	20.3/22.4/22.7	46.8/55.7/56.5
SGPN [36]	27.0/28.8/29.0	51.9/58.0/58.5
Schemata [25]	27.4/29.2/29.4	48.7/58.2/59.1
Zhang <i>et al.</i> [48]	28.5/30.0/30.1	59.3/65.0/65.3
SGFN [38]	29.5/31.2/31.2	65.9/78.8/79.6
VL-SAT (ours)	32.0/33.5/33.7	67.8/79.9/80.8
without Graph Constraints		
Co-Occurrence [48]	14.1/20.2/25.8	35.1/55.6/70.6
KERN [6]	20.8/24.7/27.6	48.3/64.8/77.2
SGPN [36]	28.2/32.6/35.3	54.5/70.1/82.4
Schemata [25]	28.8/33.5/36.3	49.6/67.1/80.2
Zhang <i>et al.</i> [48]	29.8/34.3/37.0	62.2/78.4/88.3
SGFN [38]	31.9/39.3/45.0	68.9/82.8/91.2
VL-SAT (ours)	33.8/41.3/47.0	70.5/85.0/92.5

Table 3. Quantitative results of the compared methods with respect to the SGClS and PredClS tasks, with graph constraint. The evaluation metric is top-k mean recall.

Model	SGClS	PredClS
	mR@20/50/100	mR@20/50/100
Co-Occurrence [48]	8.8/12.7/12.9	33.8/47.4/47.9
KERN [6]	9.5/11.5/11.9	18.8/25.6/26.5
SGPN [36]	19.7/22.6/23.1	32.1/38.4/38.9
Schemata [25]	23.8/27.0/27.2	35.2/42.6/43.3
Zhang <i>et al.</i> [48]	24.4/28.6/28.8	56.6/63.5/63.8
SGFN [38]	20.5/23.1/23.1	46.1/54.8/55.1
VL-SAT(ours)	31.0/32.6/32.7	57.8/64.2/64.3

6.6% gains on mR@20 in SGClS than Zhang *et al.* [48].

Qualitative Results. We provide some qualitative results between SGFN and our “VL-SAT” in Fig. 4. These results demonstrate that our method can predict more reliable scene graphs with more accurate edges and nodes. For example, our method successfully distinguishes some similar predi-

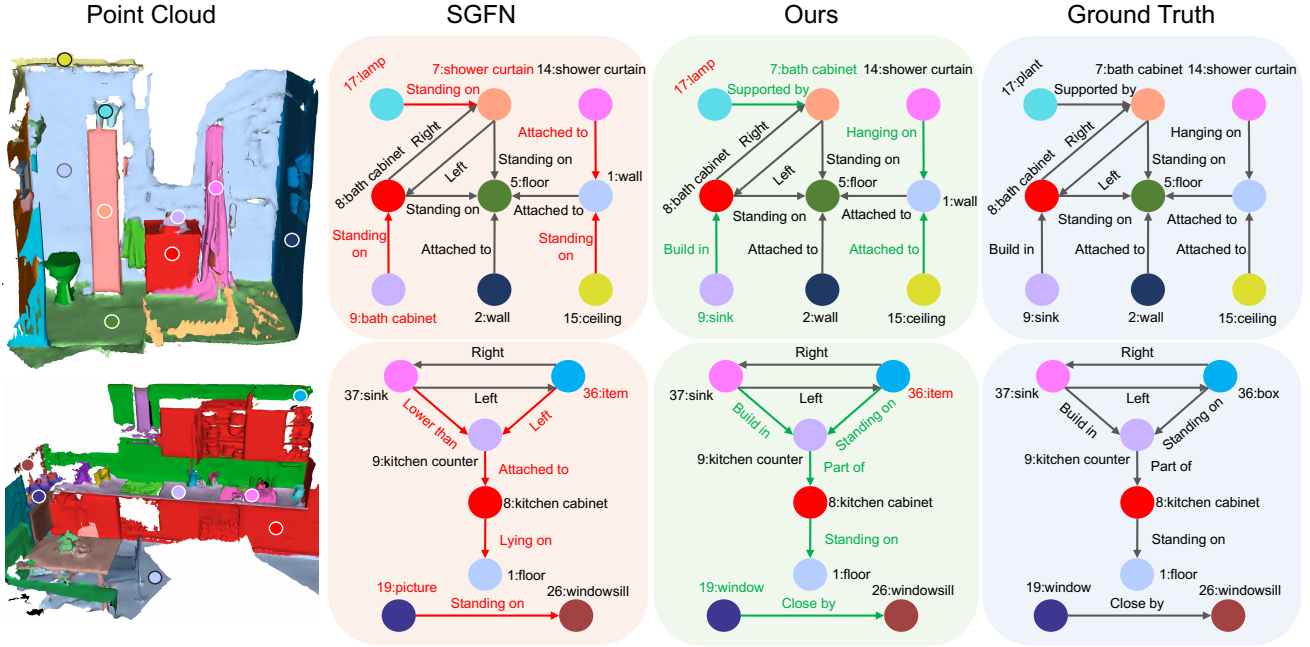


Figure 4. Qualitative results from SGFN [38] and our method on the 3DSSG [36] dataset. *Red edge*: miss-classified edges from SGFN, *green edge*: edges corrected by our method, *red node*: miss-classified node.




Figure 5. Qualitative results from SGFN [38] and our method on ScanNet [8] dataset. Note that there are no annotations of scene graphs on ScanNet [8], so we utilize the descriptions from ScanRefer [5] to parse the relationships between different objects manually. *Red edge*: miss-classified edges from SGFN, *green edge*: edges corrected by our method, *red node*: miss-classified node.

cate, like *standing on* versus *supported by* and further disambiguates related instances, such as *shower curtain* versus *bath cabinet*. The results conducted on ScanNet [8] in Fig. 5 validate that VL-SAT is generalizable to more datasets.

4.3. More Evaluations about Predicate and Triplet

Tail Predicates. In Fig. 3, we visualize the frequency of the predicates in the train set as the line chart, which shows the long-tail distribution. We also show the per-class predicate prediction performances of “VL-SAT” and SGFN [36] in the bar chart. Compared with SGFN, our method gets a significant improvement in the tail categories. To further explore the improvements brought by “VL-SAT”, we split

the 26 predicate classes into three parts: *head*, *body*, and *tail* according to their frequencies in the train set, and calculate $ma@k$ metric. In Tab. 4, we obtain 13.71% improvement on $ma@3$ when predicting the predicate on tail categories. Compared with SGFN, our method slightly drops on some head classes but significantly increases on tail classes. Since our VL-SAT boosts the overall performance by a large margin, such a slight performance degradation in head predicates is acceptable. We also provide some examples of tail predicates in Fig. 4. In the top row, our method correctly predicts tail predicates like *supported by*. In the bottom row, our method corrects the relation between the kitchen counter and the kitchen cabinet from *attached to* to *part of*.

Table 4. Based on the distribution of the predicates in the train set of the 3DSSG dataset [36], we split the 26 predicate classes into the head, body, and tail classes, and then compute mA@3 and mA@5 metrics on each split. Moreover, we test several methods on unseen and seen triplets in the validation set to evaluate the generalization ability of these methods.

Model	Predicate						Triplet			
	Head		Body		Tail		Unseen		Seen	
	mA@3	mA@5	mA@3	mA@5	mA@3	mA@5	A@50	A@100	A@50	A@100
SGPN [36]	96.66	99.17	66.19	85.73	10.18	28.41	15.78	29.60	66.60	77.03
SGFN [38]	95.08	99.38	70.02	87.81	38.67	58.21	22.59	35.68	71.44	80.11
non-VL-SAT	95.32	99.01	71.88	88.64	40.01	58.33	21.99	35.44	71.52	80.34
VL-SAT (ours)	96.31	99.21	80.03	93.64	52.38	66.13	31.28	47.26	75.09	82.25

Table 5. Results of our method when different modules are ablated. CI means CLIP-initialized object classifier. NC means node-level collaboration. EC means edge-level collaboration. TR means triplet-level CLIP-based regularization.

CI	NC	EC	TR	Object		Predicate		Triplet	
				A@5	A@10	mA@3	mA@5	mA@50	mA@100
✓				77.62	85.84	70.88	81.67	59.58	67.75
✓	✓			79.03	86.81	72.50	83.59	60.65	69.71
✓	✓	✓		79.28	86.82	73.92	84.78	62.88	71.84
✓	✓	✓	✓	78.71	86.17	76.92	87.08	64.00	72.42
✓	✓	✓	✓	78.66	85.91	77.67	87.65	65.09	73.59

Table 6. Our method with different cross-modal collaboration operations. NC means node-level collaboration. EC means edge-level collaboration. CT means concatenation. CA means cross-attention in our method.

NC	EC	Object		Predicate		Triplet	
		A@1	A@5	mA@1	mA@3	mA@50	mA@100
CT	CT	55.78	77.58	51.64	74.13	60.37	72.66
CT	CA	56.14	78.38	52.28	75.04	61.50	73.80
CA	CT	56.00	77.68	52.14	73.54	63.92	73.10
CA	CA	55.66	78.66	54.03	77.67	65.09	73.59

Table 7. Performance gains brought by our VL-SAT scheme with two reference 3DSSG prediction models.

	Object		Predicate		Triplet	
	A@1	A@5	mA@1	mA@3	mA@50	mA@100
SGG _{point} [47]	51.42	74.56	27.95	49.98	45.02	56.03
+VL-SAT	52.08	75.76	38.04	60.36	52.51	64.31
SGFN [38]	53.67	77.18	41.89	70.82	58.37	67.61
+VL-SAT	55.43	78.88	52.91	72.37	63.57	72.02

Unseen Triplets. We consider relation triplets that do not appear in the train set as unseen triplets. In Tab. 4, our method gains about 8.69% on A@50 on unseen triplets compared with SGFN [36]. The results validate that thanks to the VL-SAT scheme, our model can convey more robust feature representations based on the 3D point cloud, which leads to a better generalization ability on unseen triplets.

4.4. Ablation Study and Analysis

Ablation Study. In Tab. 5, we conduct a comprehensive ablation study. The first row denotes the baseline method “no-VL-SAT”. From Tab. 5, we could observe that the CLIP-initialized object classifier brings about 1.41% gains

on object A@5. Node and edge-level collaboration and triplet-level CLIP-based regularization steadily bring gains on triplet prediction, with 2.23%, 1.12%, and 1.09% boost on mA@50 metric. It is worth noting that regularizing the training of predicates and triplets may bring bias to the representation of objects, which leads to a slight drop in object prediction when EC/TR is employed. Thanks to the NC/CI modules, this degradation is not severe, which validates the effectiveness of our VL-SAT training scheme.

Different Cross-modal Collaboration Strategies. We investigate the effects of different cross-modal collaboration operations in Tab. 6. We compare a simple operation named CT, in which we just concatenate corresponding features between two models. When CT is applied, the mA@50 metric of the triplet prediction drops significantly. By employing our multi-head cross-attention (termed as CA) in both collaborations, significant gains can be observed.

Generalization Ability. In Tab. 7, we also show the performance gains brought by our VL-SAT scheme with two reference 3DSSG prediction models, namely SGG_{point} [47] and SGFN [38]. VL-SAT shows consistent performance gains with different 3D prediction models, especially with respect to the evaluation of predicate and triplet.

5. Conclusions

We have introduced a visual-linguistic semantics assisted training (VL-SAT) scheme to boost 3D semantic scene graph prediction in the point cloud. We build a strong oracle multi-modal model, which captures structural semantics using extra input data from vision, auxiliary training signals from language, and geometric features from the 3D model. The oracle multi-modal model enhances the 3D prediction model via back-propagated gradient flows. Consequently, the 3D prediction model can predict reliable scene graphs with only a 3D point cloud as input. Qualitative and quantitative results demonstrate that our method remarkably outperforms the existing methods.

Acknowledgements. This work was supported by National Key Research and Development Program of China (2021YFB1714300), and National Natural Science Foundation of China (62132001, 62233005).

References

- [1] Sherif Abdelkarim, Aniket Agarwal, Panos Achlioptas, Jun Chen, Jiayi Huang, Boyang Li, Kenneth Church, and Mohamed Elhoseiny. Exploring long tail visual relationship recognition with large vocabulary. In *ICCV*, pages 15921–15930, 2021. 2, 3
- [2] Iro Armeni, Zhi-Yang He, JunYoung Gwak, Amir R Zamir, Martin Fischer, Jitendra Malik, and Silvio Savarese. 3d scene graph: A structure for unified semantics, 3d space, and camera. In *ICCV*, pages 5664–5673, 2019. 2
- [3] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *CVPR*, pages 1090–1099, 2022. 2
- [4] Daigang Cai, Lichen Zhao, Jing Zhang, Lu Sheng, and Dong Xu. 3djcg: A unified framework for joint dense captioning and visual grounding on 3d point clouds. In *CVPR*, pages 16464–16473, 2022. 1, 2
- [5] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *ECCV*, pages 202–221. Springer, 2020. 1, 2, 7
- [6] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *CVPR*, pages 6163–6171, 2019. 2, 5, 6
- [7] Bowen Cheng, Lu Sheng, Shaoshuai Shi, Ming Yang, and Dong Xu. Back-tracing representative points for voting-based 3d object detection in point clouds. In *CVPR*, pages 8963–8972, 2021. 1
- [8] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 7
- [9] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. 4
- [10] Dailan He, Yusheng Zhao, Junyu Luo, Tianrui Hui, Shaofei Huang, Aixi Zhang, and Si Liu. Transrefer3d: Entity-and-relation aware transformer for fine-grained 3d visual grounding. In *ACM MM*, pages 2344–2352, 2021. 1
- [11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 5
- [12] Jean Lahoud and Bernard Ghanem. 2d-driven 3d object detection in rgb-d images. In *ICCV*, pages 4622–4630, 2017. 2
- [13] Yikang Li, Wanli Ouyang, Xiaogang Wang, and Xiao’ou Tang. ViP-CNN: Visual phrase guided convolutional neural network. In *CVPR*, pages 1347–1356, 2017. 3
- [14] Wentong Liao, Bodo Rosenhahn, Ling Shuai, and Michael Ying Yang. Natural language guided visual relationship detection. In *CVPR*, pages 0–0, 2019. 3
- [15] Yue Liao, Aixi Zhang, Miao Lu, Yongliang Wang, Xiaobo Li, and Si Liu. GEN-VLKT: Simplify association and enhance interaction understanding for hoi detection. In *CVPR*, pages 20123–20132, 2022. 2, 5
- [16] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 5
- [17] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101, 2017. 5
- [18] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*, pages 852–869. Springer, 2016. 3
- [19] Su Pang, Daniel Morris, and Hayder Radha. Clocs: Camera-lidar object candidates fusion for 3d object detection. In *IROS*, pages 10386–10393. IEEE, 2020. 2
- [20] Charles R Qi, Xinlei Chen, Or Litany, and Leonidas J Guibas. Invotenet: Boosting 3d object detection in point clouds with image votes. In *CVPR*, pages 4404–4413, 2020. 2
- [21] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *ICCV*, pages 9277–9286, 2019. 1
- [22] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *CVPR*, pages 918–927, 2018. 2
- [23] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017. 3
- [24] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 2, 5
- [25] Sahand Sharifzadeh, Sina Moayed Baharlou, and Volker Tresp. Classification by attention: Scene graph classification with prior knowledge. In *AAAI*, volume 35, pages 5025–5033, 2021. 2, 5, 6
- [26] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *CVPR*, pages 10529–10538, 2020. 1
- [27] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Point-rcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, pages 770–779, 2019. 1
- [28] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE TPAMI*, 43(8):2647–2664, 2020. 1
- [29] Vishwanath A Sindagi, Yin Zhou, and Oncel Tuzel. Mvx-net: Multimodal voxelnet for 3d object detection. In *ICRA*, pages 7276–7282. IEEE, 2019. 2
- [30] Mohammed Suhail, Abhay Mittal, Behjat Siddiquie, Chris Broaddus, Jayan Eledath, Gerard Medioni, and Leonid Sigal. Energy-based learning for scene graph generation. In *CVPR*, pages 13936–13945, 2021. 2
- [31] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from bi-ased training. In *CVPR*, pages 3716–3725, 2020. 2
- [32] Kaihua Tang, Hanwang Zhang, Baoyuan Wu, Wenhan Luo, and Wei Liu. Learning to compose dynamic tree structures for visual contexts. In *CVPR*, pages 6619–6628, 2019. 2

- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 4, 5
- [34] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *CVPR*, pages 4604–4612, 2020. 2
- [35] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. Rio: 3d object instance re-localization in changing indoor environments. In *ICCV*, pages 7658–7667, 2019. 5
- [36] Johanna Wald, Helisa Dharmo, Nassir Navab, and Federico Tombari. Learning 3d semantic scene graphs from 3d indoor reconstructions. In *CVPR*, pages 3961–3970, 2020. 1, 2, 3, 5, 6, 7, 8
- [37] Sanghyun Woo, Dahun Kim, Donghyeon Cho, and In So Kweon. Linknet: Relational embedding for scene graph. *NeurIPS*, 31, 2018. 2
- [38] Shun-Cheng Wu, Johanna Wald, Keisuke Tateno, Nassir Navab, and Federico Tombari. Scenegrphfusion: Incremental 3d scene graph prediction from rgb-d sequences. In *CVPR*, pages 7515–7525, 2021. 2, 3, 5, 6, 7, 8
- [39] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *CVPR*, pages 244–253, 2018. 2
- [40] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *CVPR*, pages 5410–5419, 2017. 2, 5, 6
- [41] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *ECCV*, pages 670–685, 2018. 2
- [42] Zhengyuan Yang, Songyang Zhang, Liwei Wang, and Jiebo Luo. Sat: 2d semantics assisted training for 3d visual grounding. In *ICCV*, pages 1856–1866, 2021. 2
- [43] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multi-modal virtual point 3d detection. *NeurIPS*, 34:16494–16507, 2021. 2
- [44] Zhihao Yuan, Xu Yan, Yinghong Liao, Yao Guo, Guanbin Li, Shuguang Cui, and Zhen Li. X-trans2cap: Cross-modal knowledge transfer using transformer for 3d dense captioning. In *CVPR*, pages 8563–8573, 2022. 2
- [45] Alireza Zareian, Svebor Karaman, and Shih-Fu Chang. Bridging knowledge graphs to generate scene graphs. In *ECCV*, pages 606–623. Springer, 2020. 2, 3, 6
- [46] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *CVPR*, pages 5831–5840, 2018. 2
- [47] Chaoyi Zhang, Jianhui Yu, Yang Song, and Weidong Cai. Exploiting edge-oriented reasoning for 3d point-based scene graph analysis. In *CVPR*, pages 9705–9715, 2021. 1, 2, 3, 5, 6, 8
- [48] Shoulong Zhang, Aimin Hao, Hong Qin, et al. Knowledge-inspired 3d scene graph prediction in point cloud. *NeurIPS*, 34:18620–18632, 2021. 2, 5, 6
- [49] Zaiwei Zhang, Bo Sun, Haitao Yang, and Qixing Huang. H3dnet: 3d object detection using hybrid geometric primitives. In *ECCV*, pages 311–329. Springer, 2020. 1
- [50] Lichen Zhao, Daigang Cai, Lu Sheng, and Dong Xu. 3dvg-transformer: Relation modeling for visual grounding on point clouds. In *ICCV*, pages 2928–2937, 2021. 2
- [51] Lichen Zhao, Jinyang Guo, Dong Xu, and Lu Sheng. Transformer3d-det: Improving 3d object detection by vote refinement. *IEEE TCSVT*, 31(12):4735–4746, 2021. 1
- [52] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, pages 15838–15847, 2021. 2