

Free-form Description Guided 3D Visual Graph Network for Object Grounding in Point Cloud

Mingtao Feng^{1*} Zhen Li^{1*} Qi Li^{1*} Liang Zhang^{1†} XiangDong Zhang¹ Guangming Zhu¹
Hui Zhang^{2†} Yaonan Wang² Ajmal Mian³
¹Xidian University, ²Hunan University, ³The University of Western Australia

Abstract

3D object grounding aims to locate the most relevant target object in a raw point cloud scene based on a free-form language description. Understanding complex and diverse descriptions, and lifting them directly to a point cloud is a new and challenging topic due to the irregular and sparse nature of point clouds. There are three main challenges in 3D object grounding: to find the main focus in the complex and diverse description; to understand the point cloud scene; and to locate the target object. In this paper, we address all three challenges. Firstly, we propose a language scene graph module to capture the rich structure and long-distance phrase correlations. Secondly, we introduce a multi-level 3D proposal relation graph module to extract the object-object and object-scene co-occurrence relationships, and strengthen the visual features of the initial proposals. Lastly, we develop a description guided 3D visual graph module to encode global contexts of phrases and proposals by a nodes matching strategy. Extensive experiments on challenging benchmark datasets (ScanRefer [3] and Nr3D [42]) show that our algorithm outperforms existing state-of-the-art. Our code is available at <https://github.com/PNXD/FFL-3DOG>.

1. Introduction

Imagine a scenario where an old person with limited mobility wakes up in the morning, feeling unwell, and instructs a robot to fetch medicine from a brown table. He/she could say “A brown table is located in the corner of the room, it is to the right of a white cabinet and to the left of black shoes. The front of it is a light blue curtain.” For a human, finding the table based on the free-form language is an easy task. However, for assistive robots, parsing the large 3D visual scene, finding the target object and understanding the global context based on natural language descriptions is a challenging task. These sentences describe the appearance of the target object (*table*), its spatial location relative to other objects (*cabinet*, *curtain* and *shoes*) and the global

*Equal contribution

†Corresponding author

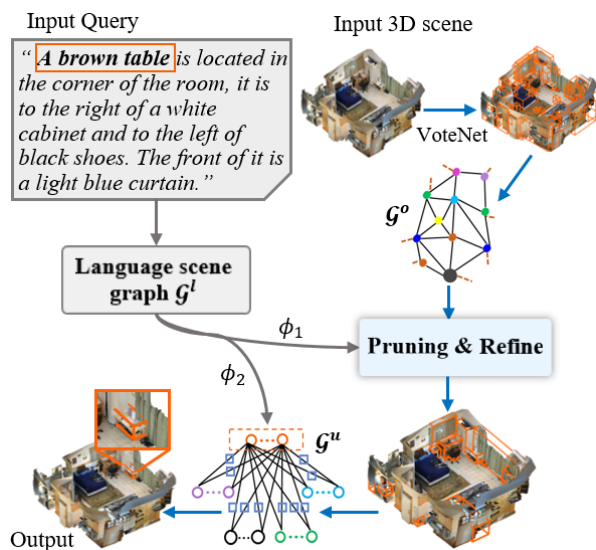


Figure 1. Proposed model for object grounding in 3D scenes. A multi-level proposal relation graph \mathcal{G}^o is formed to strengthen the visual features of the initial proposals, then the 3D visual graph \mathcal{G}^u is constructed under the guidance of the language scene graph \mathcal{G}^l which refines the initial coarse proposals. The language scene graph \mathcal{G}^l predicts the nodes matching with the 3D visual graph \mathcal{G}^u and the matching scores ϕ_1 and ϕ_2 are fused to make the final grounding predictions.

scene (*room*), which offer a rich source of information to localize the target object and guide the robot.

With the widespread availability of LiDARs, depth cameras and light field cameras [13, 12], 3D scene representation in the form of point clouds is becoming increasingly available and affordable in many application domains such as robotics, autonomous driving etc. Understanding the free-form descriptions and lifting them to the 3D point cloud scene is a new topic and challenging in the field of vision-and-language.

Research on the use of complex descriptions for 3D object grounding in point clouds is still in its infancy and only a few methods exist in the literature. To advance this line of research, ScanRefer [3] introduced the first large-scale

dataset that couples free-form descriptions to objects in 3D scenes. The ScanRefer [3] method consists of two stages: the first stage aims to use a 3D object detector VoteNet [24] to generate a set of 3D object proposals based on the input point cloud, the second stage correlates a global language expression with 3D visual features of each proposal, computes a confidence score for each fused feature, and then takes the proposal with the highest score as the target object. Similar to ScanRefer [3], Yuan et al. [42] replaced the 3D object detector with 3D panoptic segmentation backbone to obtain instance-level candidates, and captured the context of each candidate for visual and language feature matching, where the instance with the highest score is regarded as the target object.

The above methods suffer from several issues due to the inherent difficulties of both complex language processing and irregular 3D point cloud recognition. Firstly, free-form descriptions are complex and diverse in nature, and contain several sentences where strong context and long-term relation dependencies exist among them. Both ScanRefer [3] and InstanceRefer [42] do not consider that and only learn holistic representations of the description. Secondly, these two-stage solutions heavily rely on 3D object detectors or 3D panoptic segmentation backbones, and the quality of the object proposals obtained in the first stage are coarse which severely affect the overall performance. Conventional two-stage methods use the coarse object proposals directly in the following steps and fail to consider the relationships between the surrounding object proposals and global information in cluttered 3D scenes to refine the visual features of coarse proposals. Thirdly, the relationships between object proposals and phrases have not been fully explored. All existing methods [3, 42] neglect linguistic and visual structures, and only fuse the global input description embedding with 3D visual features for grounding.

To address the aforementioned limitations, we propose a free-form description guided 3D visual graph network for 3D object grounding in point clouds, as shown in Figure 1. Especially, to incorporate the rich structure and language context, we parse the complex free-form description into three components (noun phrases, pronouns and relation phrases) and construct a language scene graph to compute context-aware phrase presentation through message propagation, in which the nodes and edges correspond to noun phrase plus pronouns and relation phrases respectively. Moreover, through the 3D object detector of VoteNet [24], a set of initial 3D object proposals are extracted from the input raw point cloud. We introduce a multi-level 3D relation graph to leverage two co-occurrence relationships (object-object and object-scene), which strengthen the visual features of the initial proposals to boost the performance of the subsequent operations. Furthermore, we use the language scene graph to guide the pruning of redundant proposals

and then refine the selected ones. Built on top of the refined proposal set, we introduce a 3D visual graph to generate a context-aware object representation via message propagation. Concretely, nodes are the selected proposals relevant to the noun phrase, and the edges encode the relationships between object proposals. Finally, the nodes of 3D visual graph are adaptively matched with the nodes of language scene graph, then fused with the matching score in proposals pruning for the final 3D object grounding.

To sum up, our key contributions include: (1) We propose a free-form description guided 3D visual graph for object grounding that directly exploits the raw point cloud and is end-to-end trainable. (2) We propose a language scene graph module that captures the rich structure and long-distance phrase correlations; (3) We propose a multi-level 3D proposal relation graph module that extracts the object-object and object-scene co-occurrence relationships to strengthen the visual features of the initial proposals; (4) We propose a description 3D visual graph module that encodes global contexts of phrases and proposals through nodes matching. Experiments were performed on the benchmark ScanRefer [3] and Nr3D [3] datasets and state-of-the-art results [3, 2, 42] were achieved.

2. Related work

2D object grounding in images: 2D visual grounding aims to localize objects in an image corresponding to noun phrases from a given language description [26, 8]. A significant number of 2D visual grounding works can be found in the literature that focus on bounding-box-level [35, 31, 7, 41, 11, 17] and pixel-level [40, 18, 39] comprehensions from the input language description. Most 2D object grounding approaches follow a two stage approach where a pre-trained object detector like Faster RCNN [28] is first used to generate a set of 2D bounding box candidates based on the input image, and then the referred object is predicted in the second stage depending on the ranked matching score between each 2D bounding box candidate and the query sentence. However, most 2D object grounding methods mainly focus language description comprising a single sentence [14, 23, 19]. Although impressive progress has been made in the field of 2D object grounding, there is little work on lifting natural language processing tasks to 3D point clouds. Effective modeling of 3D vision-language tasks require establishing carefully designed connections between language and the 3D point cloud data. Whereas 2D visual grounding methods can provide some guidance for 3D relationships learning, they can not be directly used for 3D object grounding in point clouds.

3D Object Detection in Point Clouds: Point clouds are usually converted to canonical forms such as 2D images [6, 34, 27] or regular grids [43, 33, 16] for 3D object detection with Convolutional Neural Networks. Recently,

Qi et al. [24] proposed a framework to process raw point clouds directly and then predict 3D bounding boxes in cluttered scenes via a combination of PointNet++ [25] backbone and Hough voting. However, VoteNet [24] focuses on regressing each 3D object independently and the local geometric information is not well accounted for. Chen et al. [4] proposed a hierarchical graph network to aggregate features and capture shape information of objects in point clouds. Point-GNN [29] differs from previous works by introducing a graph neural network to detect the category and shape of 3D objects that each nodes in the graph belongs to. Since the 3D real world scene contains partially scanned objects with physical connections, dense placement, changing sizes, and a wide variety of challenging relationships, there is still much room for improvement in the accuracy of current 3D object detection methods.

3D Vision and Language: Compared to the significant progress made in joint inference of language and images, connecting 3D vision to natural language is a relatively new research topic. A recent dataset [20] combines RGB-D images with language expression to explore the potential gains from adding depth channel beyond a single RGB image. Kong et al. [15] exploit language description of single-view RGB-D images of scenes to align the nouns/pronouns with the referred visual objects. Chen et al. [5] presented conditional generation of 3D models from text, which could be useful in augmented reality applications. Achlioptas et al. [2] introduced a new large-scale dataset and task that identifies the specific object instances from the same category with known 3D bounding boxes. Unlike these methods, we focus on 3D object grounding task extended from ScanRefer [3] by capturing the rich structure in the free-form description, extracting the object-object and object-scene co-occurrence relationships in the 3D scene and encoding global contexts of phrases and object proposals by a nodes matching strategy.

3. Proposed Approach

Figure 2 shows an overview of our free-form description guided 3D object grounding method. It includes a language scene graph module, a multi-context proposal relation graph module and a description guided 3D visual graph module.

3.1. Language Scene Graph Module

Unlike single sentences used in most image-language tasks [36, 44, 32, 37], we adopt free-form language descriptions composed of several sentences in our work, where additional challenges such as long-distance phrase correlations emerge. Specifically, we first build a language scene graph \mathcal{G}^l from the free-form 3D scene description L to capture the rich structure and relationships between the phrases.

(1) Parsing free-form descriptions: As the language descriptions tend to describe not only the properties of the referred 3D object, but also its relationships with nearby 3D

objects or the 3D scene, we parse the complex free-form language description into three sub-components by adopting an off-the-shelf Language Scene Graph Parser [1]. We refine the Language Scene Graph Parser [1] for our free-form language description guided 3D visual grounding task by the rule-based post-processing.

Noun phrases: We first build a noun library using ground truth 3D object classes of interest in the training dataset, and then add synonyms and plural to the noun library. We extract nouns from the free-form description, and match them to the 3D object classes. We connect nouns not only to the 3D object class, but also to its synonyms and plural forms. To get 3D object attributes for a noun, such as color, shape and size, we search all attribute words that modify a noun of interest in the Scene Graph Parser [1] dependency information [15]. We generate noun phrases $\{v_i^l\}$ by combing the attribute words and the noun, e.g. for “*The table is light brown*” we have a noun phrase “*light brown table*”.

Pronouns: We consider free-form descriptions composed of several sentences which are not independent but typically refer to the same 3D object multiple times. Therefore, we are faced with the long-distance coreference problem. For example, in “*A light brown table is located in the corner of the room, it is to the right of a white cabinet*”, both “*table*” and “*it*” refer to the same 3D object and thus form a coreference. According to the statistical analysis of the free-form descriptions in dataset, we found that the pronouns refer to the referred 3D object in most cases. To address this, we extract part of speech of all words in a description L using the Language Scene Graph Parser [1], and replace pronouns with noun phrases of the referred 3D object in the descriptions, so that free-form descriptions which contain multiple sentences can deal with the long-distance correlations.

Relation phrases: To accurately parse the relation phrases $\{r_{ij}^l\}$ connecting noun phrases v_i^l and v_j^l in a free-form description, we build a relation phrase library from the training dataset. When a relation word can not be directly parsed, we look up all related words in the library and match it with the most similar word, and then expand it according to the standard form.

(2) Language Scene Graph Construction: Given a free-form description L , we first use the parsed sub-components to construct an initial language scene graph, where each node and edge of the graph correspond to an object and a relationship between this object and another object mentioned in L respectively. We define the language scene graph as a directed graph $\mathcal{G}^l = \{\mathcal{V}^l, \mathcal{R}^l\}$, where $\mathcal{V}^l = \{v_i^l\}_{i=1}^I$ and $\mathcal{R}^l = \{r_{ij}^l\}$ are nodes and edges set. As each object is represented as a noun phrase with a set of attributes, we first encode each word in description L into a sequence of word embedding $\{h_t\}_{t=1}^T$ using GLoVe [22], where T is the number of words in description L . We then encode the complete noun phrase by taking average pooling using the last

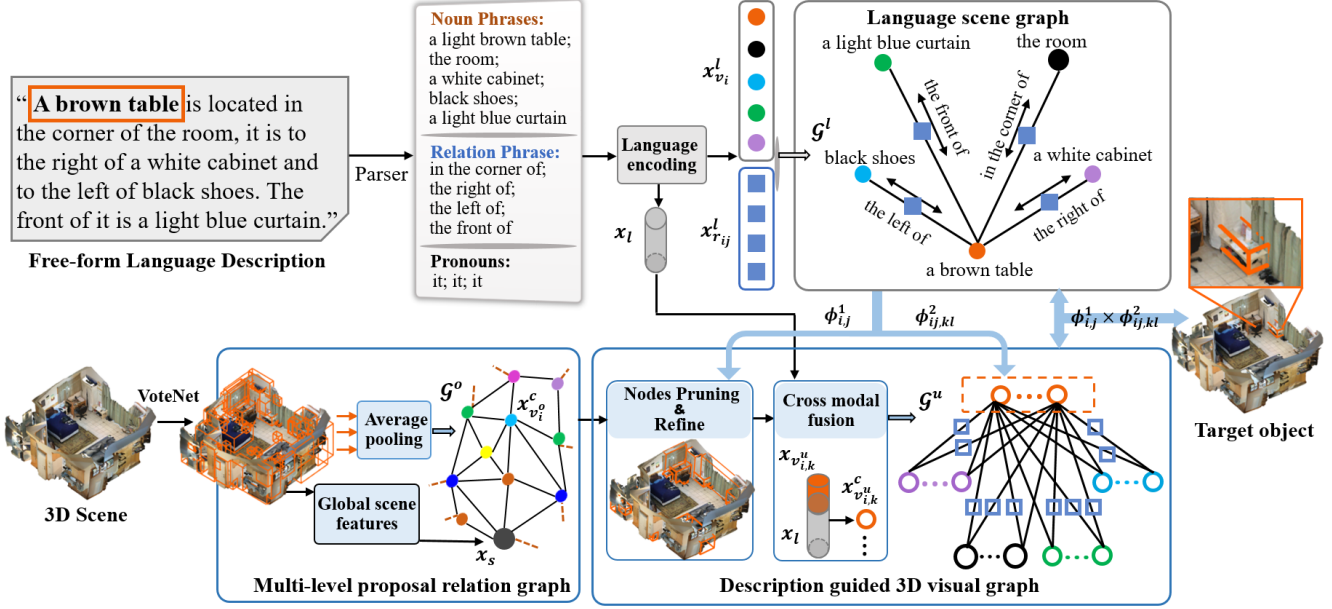


Figure 2. An overview of our proposed network. There are three modules in our method, the language scene graph \mathcal{G}^l incorporates the rich structure and language context; the multi-context proposal relation graph leverages two occurrence relationships (object-object and object-scene) to strengthen the visual features of the initial proposals set; the description guided 3D visual graph \mathcal{G}^u is defined on the pruned and refined proposals which is under the guidance of \mathcal{G}^l , then the nodes of \mathcal{G}^u are adaptively matched with the nodes of \mathcal{G}^l , and then we fuse this with the matching score in proposals pruning for the final 3D object grounding.

hidden state of a Bi-directional GRU [9]. Finally, we obtain noun phrase embedding $x_{v_i}^l$ for the corresponding node v_i^l in graph \mathcal{G}^l . Similar to the noun phrase, we also compute a vector representation $x_{r_{ij}}^l$ for edge r_{ij}^l in graph \mathcal{G}^l .

Furthermore, we use another Bi-directional GRU [9] to encode the complete description L and obtain the description presentation x_l , which provides the global context information missing in the encoding of individual noun phrases. In addition, fused with 3D point cloud features, this description presentation x_l is then used to construct the 3D visual graph.

(3) Phrase feature refinement Inspired by the message propagation operators proposed in [17], we first learn the context-aware features for all edges with its connected nodes in graph \mathcal{G}^l , where all nodes are associated with embedded features of noun phrases. We aggregate messages from nodes to update the relation features of their corresponding edge:

$$x_{r_{ij}}^c = x_{r_{ij}}^l + F_r^l([x_{v_i}^l; x_{v_j}^l; x_{r_{ij}}^l]), \quad (1)$$

where $x_{r_{ij}}^c \in \mathbb{R}^{D_1}$ is the refined relation feature, and F_r^l is a multilayer network with fully connected layers.

Then, we update each node v_i^l in graph \mathcal{G}^l by aggregating messages from connected all nodes $\mathcal{N}(i)$ and edges to it via self-attention mechanism [30]:

$$x_{v_i}^c = x_{v_i}^l + \sum_{j \in \mathcal{N}(i)} w_{v_{ij}} F_v^l([x_{v_j}^l; x_{r_{ij}}^c]), \quad (2)$$

where $x_{v_i}^c$ is the refined noun phrase feature, F_v^l is a multilayer network with fully connected layers, and $w_{v_{ij}}^l$ is an attention weight between node v_i^l and v_j^l , which is defined as follows:

$$w_{v_{ij}}^l = \text{softmax}_{j \in \mathcal{N}(i)}(F_v^l([x_{v_j}^l; x_{r_{ij}}^c])^T F_v^l([x_{v_i}^l; x_{r_{ij}}^c])), \quad (3)$$

where softmax computes the normalized attention values.

3.2. Multi-level 3D proposal relation graph module

As the first visual processing step of our model, we need to use a 3D object detection framework to predict an initial set of 3D object proposals. The quality of the 3D object proposals obtained in this step can severely affect the subsequent 3D object grounding performance. In fact, since the state-of-the-art 3D object detection methods still yield limited performance in real world 3D scenes, generally large misalignment exists between the generated 3D bounding box candidates and the ground-truth 3D objects. This could impede the learning procedure of the following stage. Leveraging two co-occurrence relationships (object-object and object-scene), we introduce the multi-level 3D proposal relation graph to encode the context dependency among the global information, which strengthens the visual features of the initial proposals by incorporating co-occurrence relation cues in the 3D scene.

We adapt the VoteNet [24] backbone to process the input point cloud and output a set of point clusters with enriched appearance features $x_a^o \in \mathbb{R}^{K_o \times D_2}$, where K_o is the

number of proposals. Next, the 3D object detection module takes in the point clusters and processes them to predict 3D bounding boxes candidates $\mathcal{B} = \{b_i\}_{i=1}^{K_o}$ and its object class for all K_o proposals, where each $b_i \in \mathbb{R}^{D_3}$ denotes 3D object location. Additionally, we use average pooling to compute a global scene feature $x_s \in \mathbb{R}^{D_2}$ depending on the set of point clusters. To encode the spatial features for each proposal, we combine two types of features, geometric structure feature S_b represented by the relative parameters of each bounding box and relative spatial location feature L_b represented by the relative center coordinates of each bounding box. To obtain strong representational power that captures co-occurrence relationships between different proposals, both appearance and spatial features are considered:

$$x_{pv}^o = F_{vf}^o([x_a^o; F_p^o([S_b; L_b])]), \quad (4)$$

where $x_{pv}^o \in \mathbb{R}^{D_4}$ is the 3D proposal visual feature, F_{vf}^o and F_p^o denote MLP layers.

We define the multi-level 3D proposal relation graph as $\mathcal{G}^o = \{\mathcal{V}^o, \mathcal{R}^o\}$, where $\mathcal{V}^o = \{v_i^o\}_{i=1}^{K_o+1}$ is the set of nodes including the initial 3D object proposals and the global scene; $\mathcal{R}^o = \{r_{ij}^o\}_{i,j=1}^{K_r}$ is the set of edges, and r_{ij}^o is the edge from v_i^o to v_j^o . Considering that an object typically only interacts with the global scene and objects nearby, the graph \mathcal{G}^o aggregates features from all its neighbour nodes $\mathcal{N}(i)$ and global scene via self-attention mechanism [30] to augment the input visual features,

$$x_{v_i^o}^c = x_{pv}^o + \sum_{j \in \mathcal{N}(i)} r_{ij}^o F_f(x_{pv}^o), \quad (5)$$

where $x_{v_i^o}^c$ is the augmented context-aware 3D object visual feature, F_f denote MLP layers, and r_{ij}^o denotes the self-attention weight [30] between nodes v_i^o and v_j^o .

3.3. Description guided 3D visual graph module

Although the visual features of the initial candidates in the previous step have been further refined, there are still hundreds of proposals and a lot of noise, which makes it infeasible to identify the target object by exploring the knowledge of language scene graph \mathcal{G}^l . We demonstrate this by introducing a description guided 3D visual graph to capture the global scene context via message propagation for few selected proposals, and to reduce the gap between 3D visual information to language structure.

(1) Nodes pruning and refinement: Most 3D object proposals are unlikely to have relationships with the language scene graph \mathcal{G}^l because the object proposals in large 3D scenes are usually redundant. To model these regularities, we introduce a description guided nodes pruning module which exploit the knowledge of our language scene graph \mathcal{G}^l to efficiently estimate the relatedness of the noun phrase nodes and the 3D object proposal nodes, and refine the parameters of the selected 3D bounding boxes. To achieve

this, we compute a matching score Φ_{ij}^1 for each noun phrase and 3D object proposal pair, and a noun phrase guided bounding box offset $\psi_{i,j} \in \mathbb{R}^{D_3}$ depending on each noun phrase feature $x_{v_i^c}^c$ and 3D object visual feature $x_{v_j^o}^c$,

$$\Phi_{i,j}^1 = H_p^n(x_{v_i^c}^c, x_{v_j^o}^c), \quad (6)$$

$$\psi_{i,j} = H_{reg}^b(x_{v_i^c}^c, x_{v_j^o}^c), \quad (7)$$

where H_p^n and H_{reg}^b are learning steps which transform the input features as in [21].

After obtaining the matching scores Φ_{ij}^1 for all noun phrase and 3D object proposal pairs, we sort the scores to choose the top K proposals for each noun phrase and obtain their new 3D bounding boxes $\mathcal{B}' = \{b_i\}_{i=1}^K$ according to the regression offset $\psi_{i,j}$. We represent the refined proposal set of each noun phrase v_i^l as $\mathcal{V}_i^u = \{v_{i,k}^u\}_{k=1}^K$. We then apply 3D Non-Maximal Suppression (NMS) to filter out proposals that have significant overlap with others. We obtain the selected 3D object proposal nodes which are much sparser and more accurate than the initial set of proposals.

(2) 3D visual graph: Instead of using noisy dense 3D object proposals of ScanRefer [3], we introduce a 3D visual graph $\mathcal{G}^u = \{\mathcal{V}^u, \mathcal{R}^u\}$ on top of the refined set of 3D object proposals. Similar to Eq. (4), we extract an new visual feature $x_{v_i^u}$ for each proposal by fusing its appearance feature and spatial feature. We then concatenate the new visual feature $x_{v_{i,k}^u}$ with the holistic feature x_l of the description as the final node feature $x_{v_{i,k}^u}^c$ of 3D visual graph. Moreover, we compute the edge feature $x_{v_{i,j,kl}^u}$ of two proposal $v_{i,k}^u$ and $v_{j,l}^u$ using the minimum box region that covers them. We adopt a set of message propagation operators to generate context aware representations for all the nodes and edges in the graph \mathcal{G}^u . Similar to Eq. (1), Eq. (2), and Eq. (3), we obtain the updated relation features $x_{v_{i,j,kl}^u}^c$, context-aware object feature $x_{v_{i,k}^u}^c$ using the respective types of message propagation operators.

4. Prediction model

Given the language scene graph and 3D visual graph, we formulate the 3D grounding as a node matching problem between the 3D visual graph \mathcal{G}^u and the language scene graph \mathcal{G}^l . We introduce a matching score $\Phi_{i,j}^2$ for each noun phrase v_i^l and proposal v_j^u pair,

$$\Phi_{i,j}^2 = H_p^u(x_{v_i^l}^c, v_j^o), \quad (8)$$

where H_p^u is a two layer MLP network. We then fuse this with the score $\Phi_{i,j}^1$ used in object pruning to generate the node matching score $\Phi_{i,j} = \Phi_{i,j}^1 \times \Phi_{i,j}^2$. Finally, the box with the highest score is selected from the K proposal boxes as the target 3D bounding box of the noun phrase. Here, only the target 3D bounding box corresponding to the subject is selected.

Methods	Unique Acc@0.5	Multiple Acc@0.5	Overall Acc@0.5
One-stage [38]	22.82	6.49	9.04
VoteNet [24] + rand.	19.35	2.81	5.28
Ours (GT boxes)	75.40	30.20	43.16
Ours (xyz)	64.04	24.13	32.47
Ours (xyz+rgb)	66.87	25.00	33.55
Ours (xyz+rgb+nor.)	67.94	25.70	34.01
Ours (w/o LSG)	66.92	23.15	32.87
Ours (w/o MLPG)	67.10	24.98	33.14
Ours (w/o LGVG)	65.34	23.75	31.90

Table 1. Ablation studies on ScanRefer [3] validation set. We measure the percentage of predictions whose IoU with the ground truth boxes is greater than 0.5. We also report scores on “unique” (single object of class in scene) and “multiple” subsets.

Loss function. The final loss is a linear combination of the vote loss \mathcal{L}_{vt} , abjectness of loss \mathcal{L}_{obj} , bounding box loss \mathcal{L}_b , semantic classification loss \mathcal{L}_{sm} , description classification loss \mathcal{L}_{cls} and reference loss \mathcal{L}_{rf} :

$$\mathcal{L} = \lambda_1 \mathcal{L}_{vt} + \lambda_2 \mathcal{L}_{obj} + \lambda_3 \mathcal{L}_b + \lambda_4 \mathcal{L}_{sm} + \lambda_5 \mathcal{L}_{cls} + \lambda_6 \mathcal{L}_{rf}, \quad (9)$$

where $\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5$ and λ_6 are the weights for the individual loss terms. Specially, \mathcal{L}_{vt} supervises the vote regression step defined in [24], \mathcal{L}_{obj} represents whether the point clusters obtained by voting and aggregation belong to a certain instance object, \mathcal{L}_b supervises the box center regression, classifying the box size classification and the box size regression processes respectively, \mathcal{L}_{sm} supervises the semantic classification process for the N ScanNet dataset classes, \mathcal{L}_{cls} is applied for an object classification based on the input description, \mathcal{L}_{rf} supervises the similarity scores and offset respectively (details in supplementary material).

5. Experiments

We introduce two experimental 3D object grounding datasets and the implementation details of our method. We perform detailed analysis of our method to demonstrate the efficacy of the proposed modules and compare the performance with state-of-the-art.

5.1. Datasets

Experiments were performed on ScanRefer [3] and Nr3D [2] datasets, which are built on top of ScanNet [10].

ScanRefer contains 51, 583 descriptions of 11, 046 3D objects for 800 ScanNet [10] real world scenes. It is the first large-scale dataset to induct 3D object grounding in point clouds via complex and diverse natural language descriptions. There are an average of 13.81 objects, 64.48 descriptions per scene, and 4.67 descriptions per object. Following the official ScanRefer [3] splits, we use the training set of 36, 665 descriptions and validation set of 9, 508 descriptions. Each description is annotated by a ground truth 3D

bounding box presented by center coordinates, orientations and dimensions in an indoor scene.

Nr3D is a real-world 3D scene dataset with extensive free-form natural language descriptions. It contains 41, 503 human utterances collected by deploying an online reference game in AMT. Following the official Nr3D [2] split, we use the training set of 29, 500 descriptions and validation set of 7, 650 descriptions. Each description is annotated by a ground truth 3D bounding box presented by center coordinates, orientations and dimensions in an indoor scene.

To augment the training data, we follow the operations in [24], randomly flip each scene in both horizontal direction, randomly rotate the scene points around the Z-axis by an angle selected uniformly between $[-30^\circ; 30^\circ]$ and globally scale the scene points between $[0.9; 1.1]$. Results and analysis are reported for the validation split.

5.2. Implementation details

We generate an initial set of $K_o = 256$ 3D object proposals using VoteNet [24]. In graph \mathcal{G}^l module, the refined relation feature dimension D_1 is set to 128. In graph \mathcal{G}^o module, the dimensions of the appearance feature D_2 is set to 256; since each 3D bounding box is parameterized by the box center, the box size and semantic classes, the parameter dimension D_3 is set to 24; the 3D proposal visual feature dimension is set to 128. In graph \mathcal{G}^u module, the selected proposals number K is set to 20. For model training, we use ADAM optimizer with initial learning rate $1e-3$. We train 30 epoches with batch size 32 and decay the learning rate 10 times after 5, 15 and 25 epoches. The loss weights of regression terms λ_1 and λ_3 are set to 1 while λ_2 is set to 0.5 and $\lambda_4 = \lambda_5 = \lambda_6$ are set to 0.1.

5.3. Ablation studies

We select the ScanRefer [3] dataset, as it contains complex and multi-sentences descriptions making it more challenging, to conduct three groups of ablation studies

Task complexity: The top part of Table 1 shows the effect of using a 2D object grounding method [38] where the 2D proposal with the highest confidence score is projected to 3D using the recorded camera parameters for that view. Grounding in 2D images results in inaccurate 3D bounding boxes as it suffers from the limited view of the 3D scene. Hence, it is necessary to design a method for direct object grounding in a 3D scene. Next, we randomly selected one of the proposals from the VoteNet [24] backbone that matches the ground truth semantic class label and found that it is insufficient to identify the referred 3D object from only the semantic label. Furthermore, we investigate the case where our model uses ground truth 3D bounding boxes and observe a huge improvement of 24.97% on overall Acc@0.5. This confirms that current 3D object detection backbones have a large room for improvement and that our description guided 3D visual graph module is necessary to refine the initial 3D bounding box candidates.

Methods	Input	Unique		Multiple		Overall	
		Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5
One-stage [38]	2D image	29.32	22.82	18.72	6.49	20.38	9.04
IntanceRefer [42]	xyz+rgb+normals	<u>77.13</u>	<u>66.40</u>	28.83	<u>22.92</u>	38.20	<u>31.35</u>
ScanRefer [3]	xyz+rgb+normals	67.64	46.19	<u>32.06</u>	21.26	38.97	26.10
Ours	xyz+rgb+normals	78.80	67.94	35.19	25.70	41.33	34.01

Table 2. Comparison with state-of-the-art on ScanRefer [3] dataset. Best performance is marked in bold and the second best is underlined.

Method	Overall	Easy	Hard	VD	VI
ScanRefer [3]	34.2	41.0	23.5	29.9	35.4
IntanceRefer [42]	<u>38.8</u>	<u>46.0</u>	<u>31.8</u>	<u>34.5</u>	<u>41.9</u>
ReferIt3D [42]	35.6	43.6	27.9	32.5	37.1
Ours	41.7	48.2	35.0	37.1	44.7

Table 3. Comparison with state-of-the-art on Nr3D [2] dataset. Here “easy” and “hard” are determined by whether there are more than two instances of the same object class in the scene. VD/VI mean view-dependent/view-independent i.e. whether the language description depends on the camera view or not.

Different inputs: We conduct an ablation study on our model to examine what components and point cloud features contribute to the performance. Results are reported in the middle part of Table 1. The performance of our model improves when *rgb* information is added to the location *xyz* features as input. The performance further improves when normals (*nor.*) are also included from the ScanNet [10] meshes. Color information and additional geometric information can enhance the performance of our network, because they are usually included in noun phrases as attributes in the language description and as part of the features of the language scene graph nodes. If the input point cloud contains both color information and additional geometric information, the expressive power of the 3D visual scene graph becomes stronger, and higher similarity matching scores will be obtained during the two matching processes, thereby improving the final prediction accuracy.

Units effectiveness: We conducted an ablation study on the effect of the proposed modules in our model. Results are reported in the bottom part of Table 1. We set the baseline when our complete framework uses spatial coordinates, color and normal information as input (i.e. *xyz + rgb + nor.*). When the language scene graph module (LGS) or the multi-level proposal graph module (MLPG) is not used, there is a drop in performance over the baseline. This is mainly because the context-aware phrases presentation in language description and the mutual occurrence relationships has the ability to better handle complex scenes. There is a higher drop in performance over the baseline when the description guided 3D visual graph module (LGVG) is not used. The reason is that, with complex and diverse descriptions, the relationships between 3D object proposals and the phrase can be ambiguous without LGVG.

5.4. Quantitative comparisons

We first compare our method with the state-of-the-art approaches including One-stage [38], IntanceRefer [42] and

ScanRefer [3] on ScanRefer [3] dataset in Table 2. Particularly, One-stage [38] is a 2D image based method, which predicts the referred 2D object in each frame of the scan video. Then they select the 2D bounding box with the highest probability value and project it to 3D space using the camera parameters of that frame. Note that IntanceRefer [42], ScanRefer [3] and our method all use coordinates, RGB information and normals of point clouds as input. We report the precision using the percentage of predictions whose IoU with the ground truth boxes is greater than 0.25 and 0.5 as the evaluation metric. As summarized in Table 2, our approach outperforms the prior methods by large margins. 2D based method One-stage [38] cannot achieve satisfactory results since it is limited by the view of a single frame. It is worth noting that our method achieves a $\sim 1.5\%$ improvement of Acc@0.5 for scenes with a single object of its class, and achieves a remarkable $\sim 2.8\%$ improvement of Acc@0.5 for more complex scenes containing multiple classes of objects. Furthermore, we observe that our proposed method performs better by a larger margin in “Multiple” than “Unique” cases. This supports our claim that our proposed pipeline comprising three novel modules can better handle complex interactions between 3D objects and the cross-modal communications from free-form descriptions to point clouds.

Table 3 compares our method with state-of-the-art methods ScanRefer [3], IntanceRefer [42] and ReferIt3D [2] on the Nr3D [2] dataset. ReferIt3D [2] assumes that ground truth 3D bounding boxes for each 3D scene are already given in the input, and the task is just to select which 3D bounding box is the referred 3D object. As shown in Table 3, our method achieves the best performance, which are 2.9% higher than IntanceRefer [42], 6.1% higher than ReferIt3D [2], and 7.5% higher than ScanRefer [3]. Also note that our method has the highest gain in accuracy for ‘Hard’ cases which again supports our claim that our proposed framework comprising the three novel modules can effectively handle complex scenes and free-form descriptions. Furthermore, our graph formulation of 3D visual information is more effective than ReferIt3D [2] in modeling the free-form description and geometric context.

5.5. Visualization

Qualitative Visualization results: Figure 3 shows 5 qualitative visual grounding results produced by the ScanRefer [3] method and our method on the ScanRefer dataset (first four columns) and the Nr3D dataset (last column).




Figure 3. Results of ScanRefer [3] method and our method on ScanRefer [3] dataset (columns 1-4) and Nr3D [2] dataset (last column).




Figure 4. Results of the most relative 3D bounding boxes for each noun phrase in description guided 3D visual graph module.

The successful detection of objects in the first two columns show that our multi-level 3D proposal graph module and description guided 3D visual module can handle the spatial relations to regress more accurate 3D bounding boxes than ScanRefer. The failure cases in the last three columns show that ScanRefer is unable to distinguish ambiguous objects in 3D scenes with complex and diverse descriptions. The performance of ScanRefer is limited since it fuses holistic language representations with visual features and ignores the relationships between proposals and phrase.

Model visualization: We visualize the results of the top $K = 20$ 3D proposals selected by the relevant noun phrase

in our description guided 3D visual graph module, as shown in Figure. 4 second row. It is evident that each noun phrase node in language scene graph matches all instances of the same category and all objects that have a strong relationship with it in the scene, which prunes and refines the initial redundant proposals set to boost the performance of the subsequent steps.

6. Conclusion

We proposed a free-form description guided 3D visual graph network for 3D object grounding in point clouds. Our method achieves accurate detection via capturing the intra-modal and cross modal relationships between the natural language descriptions and the 3D scenes. The complex free-form description is first parsed and then a language scene graph is constructed to compute a context-aware phrases presentation through message propagation. A multi-level 3D relation graph was introduced to leverage two co-occurrence relationships (object-object and object-scene) and strengthen the visual features of the initial proposals. A 3D visual graph was constructed over the refined proposals to encode global contexts of phrases and proposals under the guidance of language scene graph. Experiments on two challenging benchmark datasets show that our method quantitatively and qualitatively outperforms existing state-of-the-art in 3D object grounding.

References

- [1] <https://github.com/vacancy/scenegraphparser>. 3
- [2] Panos Achlioptas, Ahmed Abdelreheem, Fei Xia, Mohamed Elhoseiny, and Leonidas Guibas. Referit3d: Neural listeners for fine-grained 3d object identification in real-world scenes. In *European Conference on Computer Vision*, pages 422–440, 2020. 2, 3, 6, 7, 8
- [3] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. *16th European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 3, 5, 6, 7, 8
- [4] Jintai Chen, Biwen Lei, Qingyu Song, Haochao Ying, Danny Z Chen, and Jian Wu. A hierarchical graph network for 3d object detection on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 392–401, 2020. 3
- [5] Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. Text2shape: Generating shapes from natural language by learning joint embeddings. In *Asian Conference on Computer Vision*, pages 100–116. Springer, 2018. 3
- [6] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. 2
- [7] Xinpeng Chen, Lin Ma, Jingyuan Chen, Zequn Jie, Wei Liu, and Jiebo Luo. Real-time referring expression comprehension by single-stage grounding network. *arXiv preprint arXiv:1812.03426*, 2018. 2
- [8] Zhenfang Chen, Peng Wang, Lin Ma, Kwan-Yee K Wong, and Qi Wu. Cops-ref: A new dataset and task on compositional referring expression comprehension. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10086–10095, 2020. 2
- [9] Junyoung Chung, Caglar Gulcehre, Kyung Hyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling, december 2014. *arXiv preprint arXiv:1412.3555*, 2019. 4
- [10] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017. 6, 7
- [11] Pelin Dogan, Leonid Sigal, and Markus Gross. Neural sequential phrase grounding (seqground). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4175–4184, 2019. 2
- [12] Mingtao Feng, Syed Zulqarnain Gilani, Yaonan Wang, Liang Zhang, and Ajmal Mian. Relation graph network for 3d object detection in point clouds. *IEEE Transactions on Image Processing*, 30:92–107, 2020. 1
- [13] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennis. Deep learning for 3d point clouds: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 1
- [14] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014. 2
- [15] Chen Kong, Dahua Lin, Mohit Bansal, Raquel Urtasun, and Sanja Fidler. What are you talking about? text-to-image coreference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3558–3565, 2014. 3
- [16] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019. 2
- [17] Yongfei Liu, Bo Wan, Xiaodan Zhu, and Xuming He. Learning cross-modal context graph for visual grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11645–11652, 2020. 2, 4
- [18] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Liujuan Cao, Chenglin Wu, Cheng Deng, and Rongrong Ji. Multi-task collaborative network for joint referring expression comprehension and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10034–10043, 2020. 2
- [19] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016. 2
- [20] Cecilia Mauerer, Martha Palmer, and Christoffer Heckman. Sun-spot: an rgb-d dataset with spatial referring expressions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 3
- [21] Lili Mou, Rui Men, Ge Li, Yan Xu, Lu Zhang, Rui Yan, and Zhi Jin. Natural language inference by tree-based convolution and heuristic matching. In *ACL*, 2016. 5
- [22] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014. 3
- [23] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. 2
- [24] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9277–9286, 2019. 2, 3, 4, 6
- [25] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++ deep hierarchical feature learning on point sets in a metric space. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5105–5114, 2017. 3

- [26] Yanyuan Qiao, Chaorui Deng, and Qi Wu. Referring expression comprehension: A survey of methods and datasets. *IEEE Transactions on Multimedia*, 2020. 2
- [27] Mohammad Muntasir Rahman, Yanhao Tan, Jian Xue, and Ke Lu. Recent advances in 3d object detection in the era of deep neural networks: A survey. *IEEE Transactions on Image Processing*, 29:2947–2962, 2019. 2
- [28] Shaoqing Ren, Kaiming He, Ross B Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 2
- [29] Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1711–1719, 2020. 3
- [30] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations*, 2018. 4, 5
- [31] Peng Wang, Qi Wu, Jiwei Cao, Chunhua Shen, Lianli Gao, and Anton van den Hengel. Neighbourhood watch: Referring expression comprehension via language-guided graph attention networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1960–1968, 2019. 2
- [32] Chenyun Wu, Zhe Lin, Scott Cohen, Trung Bui, and Subhransu Maji. Phrasecut: Language-based image segmentation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10216–10225, 2020. 3
- [33] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018. 2
- [34] Guandao Yang, Yin Cui, Serge Belongie, and Bharath Hariharan. Learning single-view 3d reconstruction with limited pose supervision. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 86–101, 2018. 2
- [35] Sibe Yang, Guanbin Li, and Yizhou Yu. Dynamic graph attention for referring expression comprehension. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4644–4653, 2019. 2
- [36] Sibe Yang, Guanbin Li, and Yizhou Yu. Propagating over phrase relations for one-stage visual grounding. In *European Conference on Computer Vision*, pages 589–605. Springer, 2020. 3
- [37] Sibe Yang, Guanbin Li, and Yizhou Yu. Relationship-embedded representation learning for grounding referring expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 3
- [38] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4683–4693, 2019. 6, 7
- [39] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10502–10511, 2019. 2
- [40] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. Mtnet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1307–1315, 2018. 2
- [41] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer, 2016. 2
- [42] Zhihao Yuan, Xu Yan, Yinghong Liao, Ruimao Zhang, Zhen Li, and Shuguang Cui. Instancerefer: Cooperative holistic understanding for visual grounding on point clouds through instance multi-level contextual referring. *arXiv preprint arXiv:2103.01128*, 2021. 1, 2, 7
- [43] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4490–4499, 2018. 2
- [44] Haidong Zhu, Arka Sadhu, Zhaoheng Zheng, and Ram Nevatia. Utilizing every image object for semi-supervised phrase grounding. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2210–2219, 2021. 3