

Language-Assisted 3D Feature Learning for Semantic Scene Understanding

Junbo Zhang¹, Guofan Fan², Guanghan Wang¹, Zhengyuan Su¹, Kaisheng Ma^{1*}, Li Yi^{1,3,4*}

¹Tsinghua University

²Xi'an Jiaotong University

³Shanghai Artificial Intelligence Laboratory

⁴Shanghai Qi Zhi Institute

zhangjb21@mails.tsinghua.edu.cn, kaisheng@mail.tsinghua.edu.cn, ericyi@mail.tsinghua.edu.cn

Abstract

Learning descriptive 3D features is crucial for understanding 3D scenes with diverse objects and complex structures. However, it is usually unknown whether important geometric attributes and scene context obtain enough emphasis in an end-to-end trained 3D scene understanding network. To guide 3D feature learning toward important geometric attributes and scene context, we explore the help of textual scene descriptions. Given some free-form descriptions paired with 3D scenes, we extract the knowledge regarding the object relationships and object attributes. We then inject the knowledge to 3D feature learning through three classification-based auxiliary tasks. This language-assisted training can be combined with modern object detection and instance segmentation methods to promote 3D semantic scene understanding, especially in a label-deficient regime. Moreover, the 3D feature learned with language assistance is better aligned with the language features, which can benefit various 3D-language multimodal tasks. Experiments on several benchmarks of 3D-only and 3D-language tasks demonstrate the effectiveness of our language-assisted 3D feature learning. Code is available at <https://github.com/AsterisCi/Language-Assisted-3D>.

Introduction

Computer vision community has studied 3D semantic scene understanding for a long time. Remarkable progress has been made in detecting 3D objects (Qi et al. 2019; Xie et al. 2020) or segmenting 3D instances (Jiang et al. 2020; Vu et al. 2022; Chen et al. 2021c). Most of these works struggle for learning effective 3D features to depict geometric attributes of objects as well as scene context. Though these features are commonly believed to be crucial for semantic scene understanding, it is not clear whether a network is truly emphasizing these features while being end-to-end trained for a specific 3D task.

Our goal is to improve the learning process so that these important 3D features can be explicitly emphasized. And our idea comes from the following observation: humans could concisely summarize the key geometric attributes of an object as well as its surrounding context with a short natural language description (see Figure 1) and these descriptions




Figure 1: Text-scene pairs used in language-assisted 3D feature learning. “Weak alignment” means that text-object alignment annotations are not needed during training.

are very informative for identifying the object of interest. We then ask the question: can free-form languages guide 3D feature learning? We aim to imbue 3D feature learning with weak supervision knowledge from language descriptions.

3D scene understanding should benefit from text descriptions in both context modeling and object attribute learning. First, object relationships and context information in the language help understand the structure of the scene and localize specific objects. For example, “*the towel hanging on a rack*” and “*the chair facing the door*” provide spatial position and pose information of the objects, which could guide the scene understanding process when the *towel* and *chair* are hard to detect. Second, the text descriptions contain rich attribute information about the objects in a scene. For example, “*a brown wooden door*” and “*a small circular table*” provide color, shape and size information of the objects, which help to understand their semantic and geometric nature.

To emphasize the two types of features mentioned above, we design a language-assisted 3D feature learning paradigm to introduce the language prior into 3D features. Previous 2D-language joint representation learning methods, such as CLIP (Radford et al. 2021), achieve similar objectives by aligning the global features from the 2D encoder and language encoder. However, unlike 2D images, 3D scenes are much more complex and contain a richer variety of geometry and objects. One text description corresponds to only a few objects and a small part of the scene. Moreover, the amount of 3D-language data pairs is very small compared to image-text pairs mined from the web. Therefore, we propose to parse the text descriptions into scene graphs (Armeni et al. 2019; Wald et al. 2020; Zhang et al. 2021a) with object cat-

*Corresponding Author.

egories, relationships and attributes, and integrate the parsed information into 3D features at object-level. We demonstrate that this is a more precise way of introducing linguistic information into 3D models with limited paired data. Notice that text data is easy to collect since text-object alignment annotations are not needed in our method.

We find the supplemented linguistic information in our feature learning paradigm especially beneficial for data-efficient 3D scene understanding (Hou et al. 2021), where limited annotations of the 3D data are available during training. Given that the annotation for 3D scenes is laborious while collecting text descriptions is relatively easy, exploring the guidance from language is a promising direction for data-efficient 3D scene understanding.

Moreover, language guidance promotes the compatibility of our learned 3D features with 3D-language tasks, such as 3D visual grounding (Zhao et al. 2021; Yuan et al. 2021; Huang et al. 2021; Roh et al. 2021), 3D captioning (Chen et al. 2021a; Wang et al. 2022) and 3D visual question answering (Azuma et al. 2022; Yan et al. 2021). Previous methods for these 3D-language tasks usually follow a two-stage process: first, use an off-the-shelf 3D detector to extract region proposals, and then, predict the required outcome with a multimodal interaction module. They mainly focus on multimodal reasoning in the second stage. Different with previous 3D-language methods, we focus on the bottleneck of 3D feature learning: the object detectors are typically trained for general object detection of close-set categories, while a text description may refer to visual entities with various attributes and contexts. Therefore, 3D visual features learned for classical 3D scene understanding tasks might be incompatible with multimodal reasoning. We demonstrate our language-assisted 3D feature learning could promote 3D-language multimodal tasks.

Our feature learning method can be easily combined with the state-of-the-art bottom-up 3D object detection and segmentation methods, provided with natural language descriptions. Notice that the text-object alignment annotations are not needed in our method, the additional language data paired with each 3D scene can be collected with minimal effort. We test our method using various backbones and models in both 3D scene understanding benchmarks and 3D-language tasks, achieving significant progress. Our contributions can be summarized as follows:

- To the best of our knowledge, this is the first work to leverage weak language supervision to promote 3D semantic scene understanding.
- We propose a language-assisted training method for 3D scene understanding, introducing linguistic object relationship and attribute information into 3D features. Our method can be easily combined with state-of-the-art 3D detection and instance segmentation methods.
- With several benchmark experiments of 3D-only and 3D-language tasks, we demonstrate the benefits of introducing language prior into 3D scene understanding model, especially in a label-deficient regime.

Related Work

3D Object Detection and Instance Segmentation

3D object detection models take 3D data, such as point cloud and voxels, as inputs and generate 3D bounding boxes to localize the objects in a scene. Most existing methods can be considered a bottom-up manner, which requires a point grouping step to obtain object features and a detection head to generate the bounding box (Pan et al. 2021; Shi et al. 2020; Zhang et al. 2020). VoteNet (Qi et al. 2019) applies Hough Voting to group the points and vote to the similar center region. MLCVNet (Xie et al. 2020) further allows VoteNet to aggregate global contextual information via self-attention. In 3D instance segmentation, this bottom-up manner and grouping mechanisms are also widely used to generate instances. OccuSeg (Han et al. 2020a) adopts learnt occupancy signals to guide clustering. PointGroup (Jiang et al. 2020) proposes to cluster points with dual coordinate sets, which consider both the boundary and the centroid of an instance. We propose that these 3D detection and instance segmentation methods can be further promoted with additional language supervision to remedy their inherent flaws. Based on the object-level cluster features and the language describing the objects, the proposed language-assisted training method introduces object relationship and attribute information into 3D features, which helps to better recognize and localize the objects in the scene.

Multi-Modality Learning

2D & Language. For image-language reasoning, a series of works (Lu et al. 2019; Zhang et al. 2021b; Wang et al. 2021; Desai and Johnson 2021) perform BERT-like pretraining to obtain a powerful modality interaction module. Another line of work (Jia et al. 2021; Zhong et al. 2022; Li et al. 2022a) trains the image and text encoder with contrastive learning using web-scale image-text pairs. CLIP (Radford et al. 2021) first proposes to apply language supervision to the learning of images and achieves significant results in zero-shot learning. Different with 2D-language pre-training methods, we focus on promoting the 3D feature encoding module with weak language supervision, instead of the modality interaction module. We demonstrate that text-scene pairs help to remedy the inherent flaws in current 3D feature learning methods. Besides, compared with 2D data, 3D data contains a wider variety of geometry but a much smaller amount of data. Thus we propose to parse the texts and scenes into object-level, and introduce more accurate information (e.g. attribute & relation) rather than just object concept information from the language.

2D & 3D. Many works leverage 2D information fused with 3D data for 3D scene understanding (Ku et al. 2018; Qi et al. 2020; Dai and Nießner 2018; Chiang et al. 2019; Hou, Dai, and Nießner 2019). The effectiveness of 2D data has been explored in 3D representation learning (Liu et al. 2021; Li et al. 2022b; Afham et al. 2022; Liu, Qi, and Fu 2021). These works utilize contrastive learning to introduce the image prior to the 3D model.

3D & Language. In 3D-language multimodal learning, shape-text joint embedding learning (Tang et al. 2021; Koo

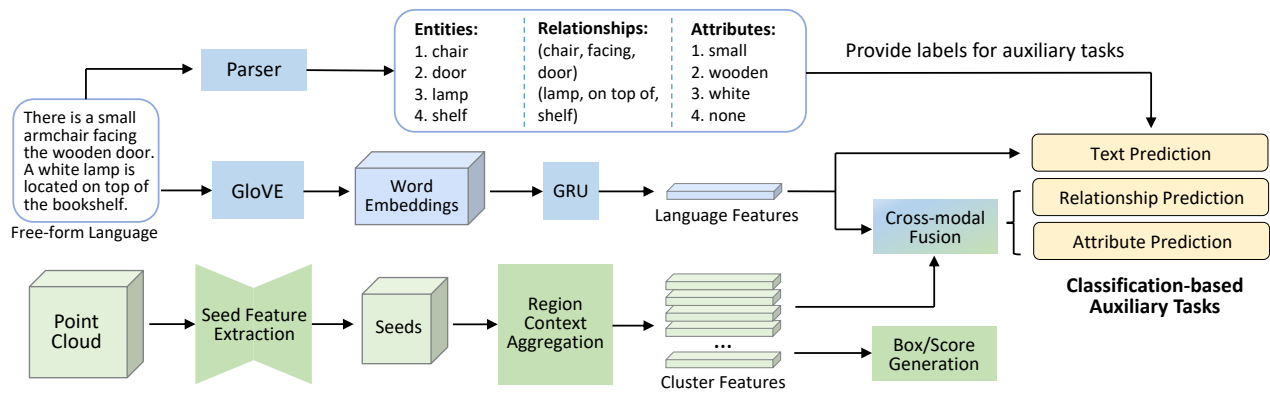


Figure 2: Overview of language-assisted training. The bottom-up 3D scene understanding branch at the bottom includes three modules: seed feature extraction, region aggregation, and box/score generation. The language processing branch at the top includes two modules: A language parser extracts entities, object relationships and attribute information from the text description. And a language model encodes the text tokens as a language feature. We fuse the language feature and object-level 3D features with a lightweight cross-attention module to perform three classification auxiliary tasks, which introduce the language prior into 3D features. The labels of these classification tasks are provided by the language parsing results.

et al. 2022) has been explored to facilitate the shape-text generation (Liu et al. 2022) and retrieval tasks (Chen et al. 2018). In contrast, we consider the multimodal learning of language and 3D objects in a scene for the first time, which has more complex forms of data correspondence and broader applications in 3D perception tasks.

3D-language Tasks

3D vision and language tasks have recently attracted much attention. Chen et al. (Chen et al. 2018) introduce descriptions for ShapeNet (Chang et al. 2015) objects, which enable text-to-shape generation (Sanghi et al. 2022; Liu et al. 2022) and shape captioning (Han et al. 2020b). On the scene level, Chen et al. (Chen, Chang, and Nießner 2020) release the ScanRefer dataset, which annotates natural language descriptions on ScanNet (Dai et al. 2017). ReferIt3D (Achlioptas et al. 2020) also proposes Nr3D and Sr3D datasets for distinguishing fine-grained objects in ScanNet. These datasets facilitate the study of 3D visual grounding (Chen, Chang, and Nießner 2020; Achlioptas et al. 2020; Zhao et al. 2021; Yuan et al. 2021; Chen et al. 2021b; Abdelreheem et al. 2022), 3D captioning (Chen et al. 2021a; Wang et al. 2022) and 3D question answering (Azuma et al. 2022; Yan et al. 2021) on scene data. Current methods for 3D-language tasks mainly focus on the interaction of the two modalities for better reasoning, instead of 3D feature learning that we focus on. However, the 3D feature learning part is a bottleneck for these tasks, due to the difficulty of detecting objects in 3D point clouds and the inconsistency between visual backbones and multimodal reasoning. We focus on leveraging attribute and context information to 3D features, which can bring general improvements to 3D-language tasks.

Method

Overview

Our goal is to introduce language information from paired text descriptions into 3D features. There are three critical

problems in our designs. The first is what kind of knowledge from the language should be extracted. The second is how to leverage language information into 3D features without text-object alignment annotations. The third is which level of features in the model needs to be introduced with language information. For question 1, we propose to extract object relationship and attribute information using an off-the-shelf language parser. For question 2, we introduce the cross-modal knowledge by training the fused language and 3D features to perform a series of auxiliary tasks. For question 3, in the 3D models, object-level features are used to perform the auxiliary tasks, since the core subjects of a language description are usually some objects in the scene.

Figure 2 shows an overview of our language-assisted training architecture for 3D scene understanding. There are mainly two branches in our method: the language processing branch at the top (in blue), and the 3D scene understanding branch at the bottom (in green).

For the 3D scene understanding branch, we ground our method on popular bottom-up object detection and instance segmentation methods, which generally include three modules as shown in Figure 2: seed feature extraction in which a backbone processes point clouds into a dense set of features, region context aggregation that summarizes spatial regions (such as voting and grouping (Qi et al. 2019), RoI pooling (Shi, Wang, and Li 2019) and clustering (Jiang et al. 2020)) to produce a set of cluster features, and box/score generation module that outputs scene understanding results from each cluster feature. We take the cluster features after the region context aggregation for cross-modal fusion.

For the language processing branch, we take the free-form language paired with each scene as input, which is tokenized into word embeddings using a pretrained GloVe (Pennington, Socher, and Manning 2014) model. The sequences of word embedding vectors are fed into a language encoder to aggregate the textual information and extracted as a language feature. Meanwhile, the free-form text descriptions

are parsed and post-processed to extract specific information, such as the entities contained in the language and their relationships and attributes.

After feature extracting and language processing, the object-level 3D features and language features are fused with a lightweight cross-attention module and used to perform three classification-based auxiliary tasks. The first is *text classification* that supervises the language encoder. The remaining tasks are *relationship classification* and *attribute classification* which help to introduce linguistic information into 3D features. We choose the most frequent relationships and attributes from the language parsing results as the classification ground truth for each task. By training these auxiliary objectives, the 3D features become “language-aware”, and the linguistic information overlooked by traditional can be supplemented.

Different with 2D-language training methods for multi-modal downstream tasks (Kamath et al. 2021), we only train the text/3D feature encoder with weak text-scene alignments, instead of multi-modal reasoning module. After language-assisted training, 3D model can be used to perform 3D-only tasks, while the language encoder and 3D model can be used in 3D-language tasks. Next, we describe the language processing and the design of auxiliary tasks in detail.

Language Parsing and Post-processing

Given a free-form text, we aim to extract all the entities mentioned in the language and parse their relationships and attributes. First, we deal with the problem of pronouns’ reference. Natural languages often face the long-distance coreference problem, such as “*This is a brown chair. It is located under the whiteboard.*”, where “*chair*” and “*it*” refer to the same 3D object in the scene. Following Feng et al. (Feng et al. 2021), we replace pronouns with the noun phrases of the referred 3D object in the descriptions. Because the statistical analysis of the languages in the dataset shows that the pronouns refer to the referred 3D object in most cases.

Then, we utilize a rule-based language parser to extract the entities and their relationships in the description. We add additional rules to the language parser to adapt the model to the 3D scene data and corresponding language descriptions. After parsing the language, we obtain several noun phrases, e.g., “*a small armchair*”, “*the brown circle table*”, each representing an entity in the scene. Notice that the language parser combines the attribute and the noun words for each entity to form a noun phrase. In addition, we can obtain the relationships of the entities, which are expressed in the triple form: (*subject, relation, object*). For example, for the description “*The small chair is facing the door*”, the resulting relationship of the two entities is (*chair, facing, door*).

Given the noun phrases and their relationships extracted from the language, we select the most frequent relation phrases and attribute words in the dataset. In total, we select 8 relationships, of which the frequency is greater than 500, and 26 attributes, of which the frequency is greater than 100. The selected attributes can be further divided into color, shape and size. Then we match entity names to the object classes in the 3D dataset. To sum up, for each text description, we find all the entities and save their relationships and

attributes. These will be regarded as the language prior and used as the ground truth labels for the auxiliary tasks.

Language-assisted Training for 3D Semantic Scene Understanding

In order to leverage language prior to 3D models, we propose to learn three auxiliary classification tasks along with the 3D perception task. The auxiliary tasks are based on the outputs of 3D encoder, language encoder and the language parsing result. We first introduce the data preparations and upstream models for the auxiliary tasks and then the detailed design of each task.

Preparation for the Auxiliary Tasks. Given a 3D scene data and a paired text description, we extract global language features and object-level 3D features with a language encoder and a 3D detection/segmentation model. The word tokens from the text description are embedded using a pre-trained GloVe (Pennington, Socher, and Manning 2014) model and then fed into a GRU cell (Chung et al. 2014) to aggregate the textual information. We take the final hidden state of the GRU cell as the final language embedding $f_{lang} \in \mathbb{R}^{c_l}$. The point cloud data is fed into the seed feature extraction and region context aggregation modules in the 3D scene understanding model to obtain the object-level clusters $P = \{x_i, f_i\}_{i=1}^M$, where M is the number of clusters. $x_i \in \mathbb{R}^3$ and $f_i \in \mathbb{R}^{c_v}$ are cluster center and feature, respectively. c_l and c_v are the channel numbers of the language and cluster features. Next, the cluster features are fed to the following module to output the 3D scene understanding results, such as box regression in VoteNet (Qi et al. 2019) and score prediction in PointGroup (Jiang et al. 2020). Meanwhile, a semantic classification score $s_i \in \mathbb{R}^{N_c}$ is predicted for each cluster, where N_c is the number of object classes in the 3D dataset.

The extracted language feature and object-level 3D features are then fed into a lightweight cross-attention module to integrate multimodal information. A QKV attention module is used, where language embedding $f_{lang} \in \mathbb{R}^{c_l}$ serve as query vectors and 3D features $f_i \in \mathbb{R}^{c_v}$ in the point cluster P as key/value vectors to compute the final multi-modal features. The fused multimodal features, the semantic classification scores $S = \{s_i\}_{i=1}^M$ and the language parsing results are used to build the auxiliary tasks as followings.

Text Classification. To supervise the language encoder, we include an object classification loss based on the input description as in ScanRefer (Chen, Chang, and Nießner 2020). We consider the benchmark classes in the corresponding 3D dataset and apply a linear classifier to the language embedding. The language-to-object classification loss L_{text} is a multi-class cross-entropy loss.

Relationship Classification. If any relationship triplet $t = (s, r, o)$ is extracted from the text description, we perform a relation classification task, aiming to enhance the contextual information and the awareness of object relationships for each object feature. s, r, o represent subject, relationship and object class, respectively. For each triplet, we first select all the fused multimodal features that the predicted semantic class c_i matches the subject (s) and object (o) class. The indexes of the selected features for subjects

Method	mAP@0.5	Method	mAP@0.5
VoteNet	31.67	MLCVNet	30.18
+Ours	33.41	+Ours	33.69

Table 1: 3D object detection results on ScanNetV2.

No. of Boxes	1	4	All
VoteNet	13.74	23.08	51.92
+Ours	18.55	30.61	52.60

Table 2: 3D object detection results on ScanNetV2 using limited bounding box annotations. Each scene contains an average of about 13 bounding boxes. The metric is mAP@0.25.

and objects are denoted as $I = \{i : c_i = s\}$, $J = \{j : c_j = o\}$, $1 \leq i, j \leq M$, respectively. If using detection backbones, we only choose the cluster features that the predicted objectness mask $m_i \in \{0, 1\}$ is equal to 1, indicating that the box regressed from this feature encloses an object. Assuming N_{sub} subject features and N_{obj} object features are selected, we then concatenate each subject and object feature and obtain $N_{sub} \times N_{obj}$ pairs of features for relationship classification.

Notice that the relationship classes we considered are not mutually exclusive. For instance, a chair “next to” a door could also be “facing” the door. Therefore, we formulate the relationship prediction loss L_{rel} as per-class binary cross-entropy. This way, whether a subject-object pair should be assigned to a specific relationship label is judged independently. In practice, for each triplet, we have only one ground truth relationship type. Thus, we use the dependencies between relationships to build labels for the other relationships. For instance, if the triplet (*chair*, *left of*, *table*) is given, the relationships between the two mentioned objects can be determined to be “left” and “next to”, and not to be “above” or “under”. However, we can not determine whether the relationship between two objects is “facing”. So, in this case, the labels used in the per-class binary cross-entropy are 1 for “left” and “next to”, 0 for “above” and “under”, and 0.5 for “facing”. The label of 0.5 indicates that the ground truth of this category can not be determined based on the given relationship triplet.

In this way, for each relationship triplet $t = (s, r, o)$, we have $N_{sub} \times N_{obj}$ loss values: $l_{ij} = L_{BCE}(p_{ij}, y_r)$, $i \in I, j \in J$, where p_{ij} is the relation prediction and y_r is the generated label for per-class binary cross-entropy loss. Since the object features selected according to the predicted semantic class do not necessarily match the exact objects described by language, we only consider the most confident prediction and take the minimum of these loss values as the final relationship prediction loss: $L_t = \min(l_{ij})$. The overall relationship prediction loss is $L_{rel} = \sum_{t=1}^T L_t$, where T is the number of relationship triplets extracted from the text description.

Attribute Classification. For each entity in the text description, we select the corresponding multimodal features according to the predicted semantic class in the same way as

Method	mAP@0.25	mAP@0.5
PointGroup	71.3	55.7
+Ours	72.3	56.9

Table 3: 3D instance segmentation results.

the relationship prediction task. A linear classifier is used to predict the attribute category for each selected feature. The classification labels are provided by the language parsing results. We group the attribute types that are not in the chosen lists into “Others”. Similarly, we take the minimum of the loss values as the final attribute prediction loss L_{attr} , which is a multi-class cross-entropy loss.

Overall Loss Function. The final loss is a combination of the perception loss and auxiliary losses:

$$L_{overall} = L_{perception} + \alpha L_{text} + \beta L_{rel} + \gamma L_{attr}, \quad (1)$$

where α , β and γ are the weights for the individual loss terms. The perception loss $L_{perception}$ is the same as the used 3D detection/segmentation model.

Experiments

Experimental Setup

Dataset. We adopt the ScanNetV2 (Dai et al. 2017) dataset for 3D detection and instance segmentation. ScanNetV2 provides 1, 513 indoor scans with semantic and instance segmentation annotations for 18 object categories. Since ScanNetV2 does not provide oriented bounding box annotation, we train the detection model to predict axis-aligned bounding boxes as in VoteNet (Qi et al. 2019).

We perform language-assisted training based on a widely used point cloud based visual grounding dataset, ScanRefer (Chen, Chang, and Nießner 2020). Text-object alignment annotations are not used in our 3D feature learning process. The ScanRefer dataset provides 51, 583 free-form text descriptions based on the 800 ScanNet scenes. Each scene has an average of 13.81 objects and 64.48 descriptions. The 3D visual grounding and 3D captioning tasks are performed on ScanRefer. We follow the ScanRefer benchmark to split the train/val/test set with 36, 655, 9, 508, and 5, 410 samples, respectively. Results and analysis are conducted on the val split, as the hidden test set is not officially available.

Implementation Details. For the detection and instance segmentation backbone, we follow the same implementation details as in VoteNet (Qi et al. 2019), MLCVNet (Xie et al. 2020) and PointGroup (Jiang et al. 2020). In VoteNet and MLCVNet, we obtain the object-level features from the cluster features after voting and grouping, which will be used for the following auxiliary tasks. In PointGroup, we take the cluster features after ScoreNet as the encoded object-level features. For language-assisted training, we jointly train the 3D perception task and the auxiliary tasks for 70 epochs with an Adam optimizer and an initial learning rate of 0.001. The learning rate is multiplied by 0.3 after 50 and 60 epochs. We set α , β and γ to 0.1, 0.05 and 0.05. VoteNet and MLCVNet are trained with the batch size of 12, in which each scene is paired with one description. PointGroup is trained with the batch size of 4, in which each

Method	Input	Unique	Multiple	Overall
ScanRefer	3D	47.86	20.50	25.81
+Ours	3D	49.80	22.75	28.00
ScanRefer	2D+3D	56.01	23.52	29.83
+Ours	2D+3D	59.31	25.07	31.71

Table 4: Visual grounding results on ScanRefer. *Unique* means there is only one object of its class. *Multiple* means there are distractors of the same class as the referenced one.

Method	CiDEr@0.25IoU	CiDEr@0.5IoU
Scan2Cap	50.95	37.29
+Ours	52.97	39.06

Table 5: Captioning results on ScanRefer. We average the captioning metric with the percentage of the predicted bounding boxes whose IoU are greater than 0.25 and 0.5.

scene is paired with 8 descriptions to balance the training time. Note that each scene is trained for multiple times in an epoch since an average of 64.48 descriptions are paired with each scene. For 3D-language tasks, we train the detection backbone and the multimodal interaction module end-to-end following ScanRefer (Chen, Chang, and Nießner 2020) and Scan2Cap (Chen et al. 2021a). All the evaluation results are the average of 5 runs. Details are contained in the appendix.

Language Parser. After parsing the language, we select 8 relationships and 26 attributes to form the training set of auxiliary tasks. In total, we extract relationships from 12,325 descriptions and extract attributes from 11,402 descriptions among the 36,665 descriptions in ScanRefer’s training set. Our improved language parser successfully parses 83.28% of the training texts. More language parsing results and analysis can be found in the appendix.

Object Detection and Instance Segmentation

We first show the 3D detection and instance segmentation results on ScanNetV2 in Table 1 and Table 3. VoteNet (Qi et al. 2019) and MLCVNet (Xie et al. 2020) are used as the detection backbone, and PointGroup (Jiang et al. 2020) is used as the instance segmentation backbone. We reproduce all the baseline methods and apply the language-assisted training in our method. For language-assisted training, we only train the 3D model on the ScanNet scenes that are contained in the ScanRefer dataset and evaluate the model on the whole ScanNet test set. Specifically, we train the 3D model on 562 scenes from ScanRefer’s training set. The results show that the language prior from text descriptions is successfully introduced to 3D model and is beneficial for detection and instance segmentation tasks with different backbones.

Data-Efficient 3D Scene Understanding

We study the data-efficient 3D scene understanding to further explore the effect of language supervision. A widely used benchmark is to use limited bounding box annotations to perform 3D object detection. Following (Hou et al. 2021), we random sample $\{1, 4\}$ bounding boxes per scene and

train the detector while the average number of bounding boxes per scene is about 13. The results are shown in Table 2 with the metric of mAP@0.25. It is observed that our method outperforms the vanilla method significantly when limited bounding box annotations are provided, indicating that the knowledge from language helps 3D models deduce other object and geometry information in a scene from limited annotations. Since collecting free-form descriptions paired with 3D scenes is relatively convenient, exploring the effect of language supervision is a promising direction to promote data-efficient 3D scene understanding.

Vision-Language Task

Table 4 shows the 3D visual grounding results with ScanRefer (Chen, Chang, and Nießner 2020) baseline on ScanRefer dataset. In the *Input* column, *3D* stands for $xyz + RGB$, and *2D + 3D* means an extra 128-dimensional multiview feature for each point is added, as indicated by ScanRefer (Chen, Chang, and Nießner 2020). We use Acc@0.5IoU as the visual grounding metrics, i.e., the percentage of the correctly predicted bounding boxes whose IoUs with the ground-truth bounding boxes are larger than 0.5. The overall accuracy and the accuracies on both “unique” and “multiple” subsets are reported. *Unique* means there is only one object of its class in the scene. *Multiple* means there are distractors of the same class as the referenced object. Results on ScanRefer validation set are reported. The VoteNet backbone in ScanRefer architecture is trained with language assistance in our method, while the multimodal interaction module is kept the same.

Table 5 shows the 3D captioning results with Scan2Cap (Chen et al. 2021a) baseline on ScanRefer dataset. The VoteNet backbone in Scan2Cap is trained with language assistance in our method. We use the same metrics as in Scan2Cap (Chen et al. 2021a) to jointly measure the quality of the generated descriptions and the detected bounding boxes. We use the standard captioning metric of CiDEr (Vedantam, Zitnick, and Parikh 2015), and employ CiDEr@0.25IoU and CiDEr@0.5IoU where only the prediction whose IoU is larger than 0.25 and 0.5 will be considered. Specifically, we evaluate the descriptions by combining the captioning metric with IoU scores between predicted bounding boxes and the target bounding boxes. The results show that our method outperforms the baselines significantly, indicating that introducing language prior to 3D features is beneficial for general 3D-language tasks.

Discussion and Ablation Study

Language-assisted Training for a Specific Category

To further explore the effect of language-assisted training, we train the VoteNet with text descriptions only referring to a specific category. Table 6 shows the changes (Δ) in object detection results of the referred category and the rest categories after applying the language assistance. We show the results of those object categories that have more than 500 text descriptions and do not belong to “Others” class in ScanNetV2. The results show a significant boost in the detection results for the referred category, which is larger than the general improvement in Table 1, possibly because

	Metric	<i>Cabinet</i>	<i>Door</i>	<i>Chair</i>	<i>Desk</i>	<i>Sink</i>	<i>Shelf</i>
Refer Categories	Δ mAP@0.5	+2.68	+3.45	+2.06	+2.11	+2.38	+1.81
Rest Categories	Δ mAP@0.5	+1.21	+1.26	+1.24	+0.96	+0.72	+0.39

Table 6: mAP changes of specific categories (paired with more than 500 text descriptions) after applying the language assistance.

Method	Obj. Det.	Inst. Seg.
Baseline	31.67	55.7
+relation	32.47	56.5
+attribute	32.80	56.1
+relation+attribute	33.41	56.9

Table 7: Ablation study on auxiliary tasks in object detection and instance segmentation (metric is mAP@0.5).

Text Percent	<i>Cabinet</i>	<i>Rest</i>	<i>Door</i>	<i>Rest</i>
30%	+0.31	+0.49	+0.21	+0.95
60%	+1.43	+0.46	+1.42	+0.71
100%	+2.68	+1.21	+3.45	+1.26

Table 8: Ablation study on the number of texts (mAP@0.5). We show results for the two categories with the most textual descriptions. This result shows the potential to further improve performance by collecting more language data.

the chosen languages contain less noise for the referred category. Meanwhile, changes of the detection results for the rest categories are relatively small as shown in the last row in Table 6. Since the chosen descriptions describing a specific category can also involve other category objects, the detection of the rest categories can be improved to some degree. This experiment provides the insight that, in practical applications, we can use the descriptions referring to a specific category to improve the accuracy of those categories that are important for application or difficult to detect.

Ablation Study

Table 7 shows the ablation study on the relationship and attribute classification auxiliary tasks. The baseline method is VoteNet and PointGroup for object detection and instance segmentation. The metric is mAP@0.5. *relation* and *attribute* represent the relationship and attribute auxiliary task, respectively. It is observed that each auxiliary task is effective, and using both of them can further improve the performance. Table 8 shows the ablation study on the number of text descriptions used in language-assisted training for a specific category. We show the changes (Δ) of mAP@0.5 in the object detection model trained with different percentages of texts. It is observed that the detection results improve with the increase of text numbers. Given that collecting free-form languages describing a 3D scene is convenient, larger amount of text descriptions can be used to further improve the results in the future work.

Qualitative Comparison

Figure 3 shows the visualization of visual grounding results. The ground-truth boxes are displayed in blue. The predicted boxes of our method and the baseline are displayed in green



Figure 3: Qualitative comparison results on visual grounding. The ground-truth boxes are displayed in blue. The predicted boxes of our method and the baseline are displayed in green and red, respectively.

and red, respectively. In the first picture, there are two bookshelves in the scene, and the target of the language input is the left one. Two relation triplets can be extracted from the language description, namely, (*bookshelf*, *right of*, *cabinet*) and (*bookshelf*, *left of*, *door*). Our model better preserves this relation information and finds the correct bookshelf. In contrast, the baseline model fails to infer from relational information and predicts the wrong results. The second picture shows that our method utilizes the object relationship to correctly localize the *cabinet* among distractors, indicating that the 3D features in our method better preserve relationship information. The failure case of ScanRefer illustrates that it can not distinguish ambiguous objects according to spatial relationships. In the third example, there are two chairs and the target one is on the left. In the language descriptions, the attribute word “*long*” used to describe the table is an essential basis for finding the target object. Our model better utilizes attribute clue and predicts the right location, while the baseline method fails to distinguish between the two tables. The last picture shows that when relationship information is insufficient to make a decision, the attribute information can be used to localize the target in our method. We provide more qualitative comparison results for both visual grounding and 3D object detection in the appendix.

Conclusion

This paper proposes a language-assisted training method to promote 3D scene understanding. With free-form languages paired with 3D scenes, we extract object relationship and attribute information from languages and introduce this language prior to object-level 3D features. Three auxiliary tasks are designed to transmit cross-modal knowledge based on multimodal features. We apply the proposed method to several 3D object detection and instance segmentation models. Several benchmark experiments of 3D-only and 3D-language tasks (e.g. visual grounding, captioning) demonstrate the benefits of language supervision in 3D scene understanding. We also discuss our method’s application in data-efficient 3D scene understanding and improving the detection performance for specific categories.

References

- Abdelreheem, A.; Upadhyay, U.; Skorokhodov, I.; Yahya, R. A.; Chen, J.; and Elhoseiny, M. 2022. 3DRefTransformer: Fine-Grained Object Identification in Real-World Scenes Using Natural Language. *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 607–616.
- Achlioptas, P.; Abdelreheem, A.; Xia, F.; Elhoseiny, M.; and Guibas, L. J. 2020. ReferIt3D: Neural Listeners for Fine-Grained 3D Object Identification in Real-World Scenes. In *ECCV*.
- Afham, M.; Dissanayake, I.; Dissanayake, D.; Dharmasiri, A.; Thilakarathna, K.; and Rodrigo, R. 2022. CrossPoint: Self-Supervised Cross-Modal Contrastive Learning for 3D Point Cloud Understanding. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9892–9902.
- Armeni, I.; He, Z.-Y.; Gwak, J.; Zamir, A. R.; Fischer, M.; Malik, J.; and Savarese, S. 2019. 3D Scene Graph: A Structure for Unified Semantics, 3D Space, and Camera. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 5663–5672.
- Azuma, D.; Miyanishi, T.; Kurita, S.; and Kawanabe, M. 2022. ScanQA: 3D Question Answering for Spatial Scene Understanding. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19107–19117.
- Chang, A. X.; Funkhouser, T. A.; Guibas, L. J.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; Xiao, J.; Yi, L.; and Yu, F. 2015. ShapeNet: An Information-Rich 3D Model Repository. *ArXiv*, abs/1512.03012.
- Chen, D. Z.; Chang, A. X.; and Nießner, M. 2020. ScanRefer: 3D Object Localization in RGB-D Scans using Natural Language. In *ECCV*.
- Chen, D. Z.; Gholami, A.; Nießner, M.; and Chang, A. X. 2021a. Scan2Cap: Context-aware Dense Captioning in RGB-D Scans. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3192–3202.
- Chen, D. Z.; Wu, Q.; Nießner, M.; and Chang, A. X. 2021b. D3Net: A Speaker-Listener Architecture for Semi-supervised Dense Captioning and Visual Grounding in RGB-D Scans. *ArXiv*, abs/2112.01551.
- Chen, K.; Choy, C. B.; Savva, M.; Chang, A. X.; Funkhouser, T. A.; and Savarese, S. 2018. Text2Shape: Generating Shapes from Natural Language by Learning Joint Embeddings. *ArXiv*, abs/1803.08495.
- Chen, S.; Fang, J.; Zhang, Q.; Liu, W.; and Wang, X. 2021c. Hierarchical Aggregation for 3D Instance Segmentation. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 15447–15456.
- Chiang, H.; Lin, Y.-L.; Liu, Y.-C.; and Hsu, W. H. 2019. A Unified Point-Based Framework for 3D Segmentation. *2019 International Conference on 3D Vision (3DV)*, 155–163.
- Chung, J.; Çağlar Gülçehre; Cho, K.; and Bengio, Y. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *ArXiv*, abs/1412.3555.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T. A.; and Nießner, M. 2017. ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2432–2443.
- Dai, A.; and Nießner, M. 2018. 3DMV: Joint 3D-Multi-View Prediction for 3D Semantic Scene Segmentation. In *ECCV*.
- Desai, K.; and Johnson, J. 2021. VirTex: Learning Visual Representations from Textual Annotations. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11157–11168.
- Feng, M.; Li, Z.; Li, Q.; Zhang, L.; Zhang, X.; Zhu, G.; Zhang, H.; Wang, Y.; and Mian, A. S. 2021. Free-form Description Guided 3D Visual Graph Network for Object Grounding in Point Cloud. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3702–3711.
- Han, L.; Zheng, T.; Xu, L.; and Fang, L. 2020a. OccluSeg: Occupancy-Aware 3D Instance Segmentation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2937–2946.
- Han, Z.; Chen, C.; Liu, Y.-S.; and Zwicker, M. 2020b. ShapeCaptioner: Generative Caption Network for 3D Shapes by Learning a Mapping from Parts Detected in Multiple Views to Sentences. *Proceedings of the 28th ACM International Conference on Multimedia*.
- Hou, J.; Dai, A.; and Nießner, M. 2019. 3D-SIS: 3D Semantic Instance Segmentation of RGB-D Scans. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4416–4425.
- Hou, J.; Graham, B.; Nießner, M.; and Xie, S. 2021. Exploring Data-Efficient 3D Scene Understanding with Contrastive Scene Contexts. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15582–15592.
- Huang, P.-H.; Lee, H.-H.; Chen, H.-T.; and Liu, T.-L. 2021. Text-Guided Graph Neural Networks for Referring 3D Instance Segmentation. In *AAAI*.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q. V.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In *ICML*.
- Jiang, L.; Zhao, H.; Shi, S.; Liu, S.; Fu, C.-W.; and Jia, J. 2020. PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4866–4875.
- Kamath, A.; Singh, M.; LeCun, Y.; Misra, I.; Synnaeve, G.; and Carion, N. 2021. MDETR - Modulated Detection for End-to-End Multi-Modal Understanding. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 1760–1770.
- Koo, J.; Huang, I. Y.; Achlioptas, P.; Guibas, L. J.; and Sung, M. 2022. PartGlot: Learning Shape Part Segmentation from Language Reference Games. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16484–16493.
- Ku, J.; Mozifian, M.; Lee, J.; Harakeh, A.; and Waslander, S. L. 2018. Joint 3D Proposal Generation and Object Detection from View Aggregation. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1–8.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; Chang, K.-W.;

- and Gao, J. 2022a. Grounded Language-Image Pre-training. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10955–10965.
- Li, Z.; Chen, Z.; Li, A.; Fang, L.; Jiang, Q.; Liu, X.; Jiang, J.; Zhou, B.; and Zhao, H. 2022b. SimIPU: Simple 2D Image and 3D Point Cloud Unsupervised Pre-Training for Spatial-Aware Visual Representations. In *AAAI Conference on Artificial Intelligence*.
- Liu, Y.-C.; Huang, Y.-K.; Chiang, H.; Su, H.-T.; Liu, Z.-Y.; Chen, C.-T.; Tseng, C.-Y.; and Hsu, W. H. 2021. Learning from 2D: Pixel-to-Point Knowledge Transfer for 3D Pre-training. *ArXiv*, abs/2104.04687.
- Liu, Z.; Qi, X.; and Fu, C.-W. 2021. 3D-to-2D Distillation for Indoor Scene Parsing. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4462–4472.
- Liu, Z.; Wang, Y.; Qi, X.; and Fu, C.-W. 2022. Towards Implicit Text-Guided 3D Shape Generation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17875–17885.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *NeurIPS*.
- Pan, X.; Xia, Z.; Song, S.; Li, L. E.; and Huang, G. 2021. 3D Object Detection with Pointformer. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7459–7468.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. GloVe: Global Vectors for Word Representation. In *EMNLP*.
- Qi, C.; Chen, X.; Litany, O.; and Guibas, L. J. 2020. ImVoteNet: Boosting 3D Object Detection in Point Clouds With Image Votes. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4403–4412.
- Qi, C.; Litany, O.; He, K.; and Guibas, L. J. 2019. Deep Hough Voting for 3D Object Detection in Point Clouds. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9276–9285.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*.
- Roh, J.; Desingh, K.; Farhadi, A.; and Fox, D. 2021. LanguageRefer: Spatial-Language Model for 3D Visual Grounding. In *CoRL*.
- Sanghi, A.; Chu, H.; Lambourne, J.; Wang, Y.; Cheng, C.-Y.; and Fumero, M. 2022. CLIP-Forge: Towards Zero-Shot Text-to-Shape Generation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18582–18592.
- Shi, S.; Guo, C.; Jiang, L.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020. PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10526–10535.
- Shi, S.; Wang, X.; and Li, H. 2019. PointRCNN: 3D Object Proposal Generation and Detection From Point Cloud. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–779.
- Tang, C.; Yang, X.; Wu, B.; Han, Z.; and Chang, Y. 2021. Part2Word: Learning Joint Embedding of Point Clouds and Text by Matching Parts to Words. *ArXiv*, abs/2107.01872.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. CIDER: Consensus-based image description evaluation. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4566–4575.
- Vu, T.; Kim, K.; Luu, T. M.; Nguyen, X. T.; and Yoo, C.-D. 2022. SoftGroup for 3D Instance Segmentation on Point Clouds. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2698–2707.
- Wald, J.; Dhano, H.; Navab, N.; and Tombari, F. 2020. Learning 3D Semantic Scene Graphs From 3D Indoor Reconstructions. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3960–3969.
- Wang, H.; Zhang, C.; Yu, J.; and Cai, W. T. 2022. Spatiality-guided Transformer for 3D Dense Captioning on Point Clouds. In *International Joint Conference on Artificial Intelligence*.
- Wang, Z.; Yu, J.; Yu, A. W.; Dai, Z.; Tsvetkov, Y.; and Cao, Y. 2021. SimVLM: Simple Visual Language Model Pre-training with Weak Supervision. *ArXiv*, abs/2108.10904.
- Xie, Q.; Lai, Y.-K.; Wu, J.; Wang, Z.; Zhang, Y.; Xu, K.; and Wang, J. 2020. MLCVNet: Multi-Level Context VoteNet for 3D Object Detection. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10444–10453.
- Yan, X.; Yuan, Z.; Du, Y.; Liao, Y.; Guo, Y.; Li, Z.; and Cui, S. 2021. CLEVR3D: Compositional Language and Elementary Visual Reasoning for Question Answering in 3D Real-World Scenes. *ArXiv*, abs/2112.11691.
- Yuan, Z.; Yan, X.; Liao, Y.; Zhang, R.; Li, Z.; and Cui, S. 2021. InstanceRefer: Cooperative Holistic Understanding for Visual Grounding on Point Clouds through Instance Multi-level Contextual Referring. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 1771–1780.
- Zhang, C.; Yu, J.; Song, Y.; and Cai, W. T. 2021a. Exploiting Edge-Oriented Reasoning for 3D Point-based Scene Graph Analysis. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9700–9710.
- Zhang, P.; Li, X.; Hu, X.; Yang, J.; Zhang, L.; Wang, L.; Choi, Y.; and Gao, J. 2021b. VinVL: Revisiting Visual Representations in Vision-Language Models. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5575–5584.
- Zhang, Z.; Sun, B.; Yang, H.; and Huang, Q.-X. 2020. H3DNet: 3D Object Detection Using Hybrid Geometric Primitives. In *ECCV*.
- Zhao, L.; Cai, D.; Sheng, L.; and Xu, D. 2021. 3DVG-Transformer: Relation Modeling for Visual Grounding on Point Clouds. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2908–2917.
- Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N. C. F.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; and Gao, J. 2022. RegionCLIP: Region-based Language-Image Pre-training. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16772–16782.