

Visual Programming for Zero-shot Open-Vocabulary 3D Visual Grounding

Zhihao Yuan^{1,2}, Jinke Ren^{1,2}, Chun-Mei Feng³, Hengshuang Zhao⁴, Shuguang Cui^{2,1}, Zhen Li^{2,1*}

¹ FNii, CUHKSZ ² SSE, CUHKSZ ³ IHPC, A*STAR, Singapore ⁴ HKU

Abstract

3D Visual Grounding (3DVG) aims at localizing 3D object based on textual descriptions. Conventional supervised methods for 3DVG often necessitate extensive annotations and a predefined vocabulary, which can be restrictive. To address this issue, we propose a novel visual programming approach for zero-shot open-vocabulary 3DVG, leveraging the capabilities of large language models (LLMs). Our approach begins with a unique dialog-based method, engaging with LLMs to establish a foundational understanding of zero-shot 3DVG. Building on this, we design a visual program that consists of three types of modules, i.e., view-independent, view-dependent, and functional modules. These modules, specifically tailored for 3D scenarios, work collaboratively to perform complex reasoning and inference. Furthermore, we develop an innovative language-object correlation module to extend the scope of existing 3D object detectors into open-vocabulary scenarios. Extensive experiments demonstrate that our zero-shot approach can outperform some supervised baselines, marking a significant stride towards effective 3DVG. Code is available at <https://curryyuan.github.io/ZSVG3D>.

1. Introduction

3D Visual Grounding (3DVG) aims to localize specific objects within 3D scenes by using a series of textual descriptions. This has become a crucial component in a variety of burgeoning applications, such as autonomous robotics [12, 53, 58], virtual reality [40, 55], and metaverse [10, 33]. For illustration, given a 3D scan in Figure 1(a) along with its description — It is the keyboard closest to the door, the goal of 3DVG is to accurately pinpoint the keyboard in the green box, while eliminating potential distractions such as tables and desks. Despite the apparent simplicity of this task for humans, it poses a significant challenge for machines due to their inherently limited perceptual capabilities.

*Corresponding author.

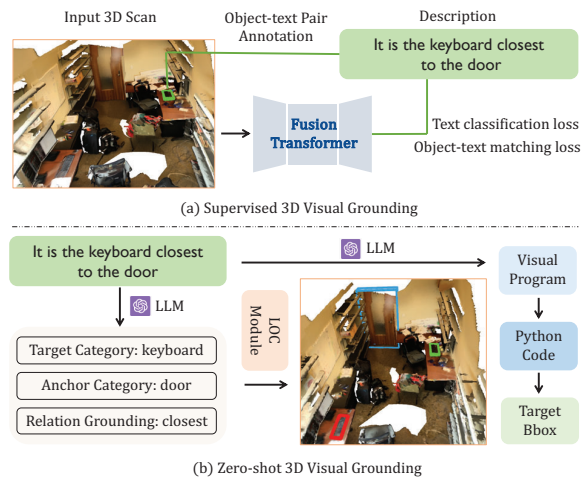


Figure 1. **Comparative overview of two 3DVG approaches**, where (a) *Supervised 3DVG* involves input from 3D scans combined with text queries, guided by object-text pair annotations, (b) *Zero-shot 3DVG* identifies the location of target objects using programmatic representation generated by LLMs, i.e., *target category*, *anchor category*, and *relation grounding*, thereby highlighting its superiority in decoding spatial relations and object identifiers within a given space, e.g., the location of the keyboard (outlined in green) can be retrieved based on the distance between the keyboard and the door (outlined in blue).

Traditional supervised 3DVG approaches [19, 60, 62] achieve this objective by leveraging the rich annotations in public datasets, such as ScanRefer [4] and ReferIt3D [1]. These approaches typically define 3DVG as a matching problem, generating possible objects via 3D detectors [22, 38], and identifying the best match by fusing the visual and textual features. While these approaches can yield precise results, the acquisition of sufficient annotations is prohibitively resource-intensive for real-world applications. Furthermore, these approaches are often constrained by the pre-defined vocabulary during training, making them sub-optimal in open-vocabulary scenarios.

To address these issues, we propose a novel visual programming approach for 3DVG that integrates zero-shot learning and large language models (LLMs). Zero-shot

learning [65, 67, 70] can generalize across new categories by leveraging the pre-trained capabilities of CLIP [41] in the 3D domain. LLMs [34, 47, 49] can facilitate 3DVG due to their strong planning and reasoning capabilities. Regarding this, we first propose a vanilla version *dialog with LLM*. It describes the location and size of all objects in the scene and instructs the LLM to distinguish the object of interest through an interactive dialog. Despite the simplicity of the base approach, the inherent stochasticity and control limitations of LLMs make it hard to capture the **view-dependent queries** and decipher **spatial relations in 3D space**, which are the main challenges of 3DVG. To overcome this limitation, we further develop a new visual programming approach, as shown in Figure 1(b). It mainly consists of three steps: (1) generating a 3D visual program using LLMs, (2) interpreting the program into Python code, and (3) identifying the target bounding box by executing the code. To enhance the localization accuracy, we further introduce a novel language-object correlation (LOC) module capable of merging the geometric discernment of 3D point clouds with the fine-grained appearance acumen of 2D images.

In summary, contributions are summarized as follows:

- We propose an innovative 3D visual programming approach. It eliminates the need for extensive object-text pair annotations required in supervised approaches.
- We transform the visual program to Python code by designing two types of modules, *i.e.*, relation modules and LOC modules. The former explicitly defines the view-dependent and view-independent relations in 3D space, while the latter captures both the geometric and appearance information for open-vocabulary localization.
- We conduct extensive experiments on two popular datasets, *i.e.*, ScanRefer [4] and Nr3d [1]. We for the first time evaluate the whole validation set rather than a few samples. The results demonstrate the superior performance of our approach, even comparable with existing supervised approaches.

2. Related Work

Supervised 3DVG. 3DVG has received much attention in many emerging applications ranging from automatic robotics [12, 53, 58] to metaverse [10, 33]. On the one hand, densely-annotated datasets like ScanRefer [4] and Referit3D [1] can provide well-aligned object-text pairs for ScanNet [9]. On the other hand, most existing methods [6, 14, 42, 60, 62] treat 3DVG as a matching problem, where object identifiers [22, 38] are utilized to generate candidate objects and find the best-matching one by fusing visual and textual features. Building on this, [57, 63] attempt to explore the object attributes and relations between different proposals. Moreover, some works [7, 23] have also investigated 3D language pretraining using advanced techniques, such as mask modeling and contrastive learning on paired

object-caption data, followed by finetuning on downstream tasks. Additionally, NS3D [17] has employed CodeX [5] to generate hierarchical programs. However, it still needs annotations to train the neuro-symbolic networks, thus lacking open-vocabulary and zero-shot capabilities.

Indoor 3D Scene Understanding. 3D scene understanding of indoor environments has been widely studied. In specific, the emergence of RGB-D scans datasets [9, 44, 52] greatly push the boundary of several tasks, including 3D object classification [36, 37], 3D object detection [31, 38], 3D semantic segmentation [8, 66], 3D instance segmentation [22, 43, 51], and so on [16, 59, 64]. However, these methods are often constrained to a closed set of semantic class labels, limiting their applicability in real-world scenarios. Recent progress in open-vocabulary image segmentation [13, 29] has inspired research into 3D scene understanding under the open-vocabulary setting. For instance, LERF [24] learns a language field inside NeRF [32] by volume rendering CLIP [41] features along training rays, enabling it to generate 3D relevancy maps for arbitrary language queries. OpenScene [35] extracts image features using 2D open-vocabulary segmentation models [13, 27], then trained a 3D network to produce point features aligned with multi-view fused pixel features. OpenMask3D [46] utilizes the closed-vocabulary network to generate instance masks while discarded the classification head. Despite these advancements, these methods still lack spatial and common-sense reasoning abilities.

LLMs for Vision-Language Tasks. Recent progress on LLMs has provided impressive zero-shot planning and reasoning abilities [34, 47, 49]. Advanced prompting technologies such as Least-to-Most [69], Think-Step-by-Step [26], and Chain-of-Thought [54] are proposed to elicit the capabilities of LLMs. These methods can understand human instructions, break complex goals into sub-goals, and control robot agents to execute tasks without additional training [2, 20, 30]. Moreover, when combined with specialized vision models, LLMs can significantly enhance the performance of vision-language tasks. For instance, Visual ChatGPT [56] uses ChatGPT as a central orchestrator, interfacing with a variety of visual foundation models to solve more challenging problems. VISPROG [15] leverages the in-context learning ability to generate high-level modular programs for solving complex and compositional natural language reasoning and image editing tasks. ViperGPT [45] directly feeds the API of available modules to LLM and then generated executable Python code for image grounding. However, leveraging these capabilities for zero-shot 3D language grounding remains an unexplored area.

3. Methodology

In Section 3.1, we introduce the vanilla approach, *i.e.*, dialog with LLM to overcome the annotation issue in 3DVG.

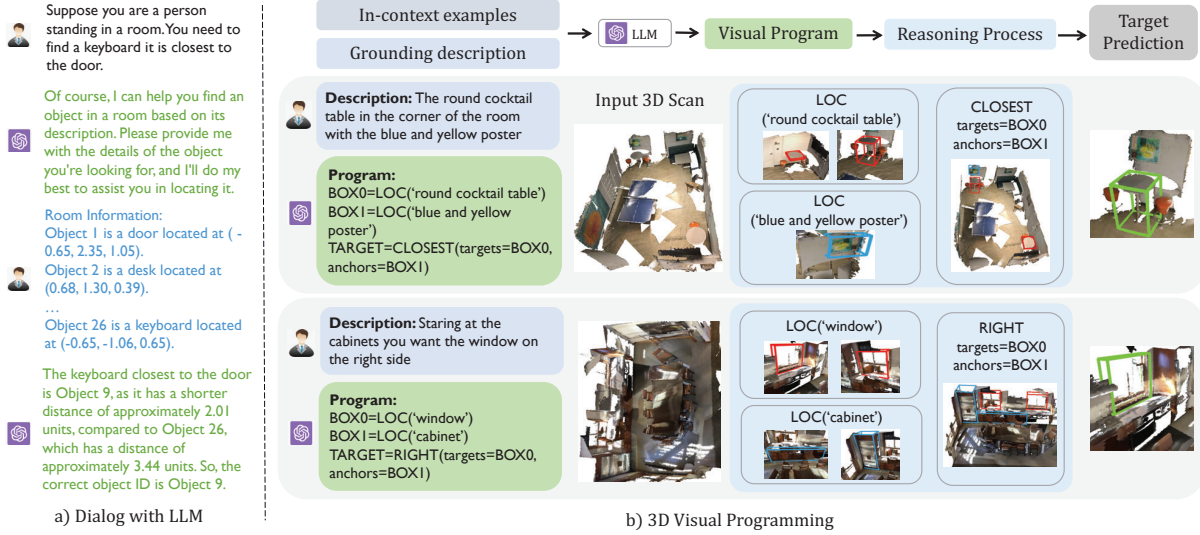


Figure 2. **Overview of two zero-shot approaches for 3DVG.** (a) shows the working mechanism of the vanilla dialog with LLM approach. First, we describe the 3DVG task and provide the text descriptions of the room. Then, LLMs identify the objects relevant to the query sentence and perform human-like reasoning. (b) presents the 3D visual programming approach. We first input in-context examples into LLMs. Then, LLMs generate 3D visual programs through the grounding descriptions and perform human-like reasoning. Next, these programs are transformed into executable Python codes via the LOC module for predicting the location of the object. For example, the upper example uses the view-independent module, i.e., CLOSEST to determine the proximity in 3D space, while the lower example applies the view-dependent module, i.e., RIGHT to establish the relative positioning.

From Section 3.2 to Section 3.4, we present the visual programming approach, address the issue of view-dependent relations, and design the LOC module, respectively.

3.1. Dialog with LLM

To accomplish the goal of 3DVG, we propose to initiate a dialogue with LLMs. The input for the dialogue consists of a real-world RGB-D scan and a free-form text description \mathcal{T} . The text description provides specific information about the target object within a point cloud representation $\mathbf{P} \in \mathbb{R}^{N \times 6}$, where \mathbf{P} is a collection of color-enriched 3D points and N is the total number of such points. The LLM acts as an agent located in the scanned room, which aims to identify the specified object based on the given text description. To bridge the gap between the model’s proficiency in understanding text and the spatial nature of the 3DVG task, we first transform the scene into a textual narrative. This narrative can provide a comprehensive account of the objects \mathcal{O} presented in the scene, including their positions and dimensions, which can be expressed as:

Object <id> is a <category>
located at (x, y, z) with sizes
(width, length, height).

Given such textual layout, we dialog with the LLM by providing the scene’s description and query. Our objec-

tive is to guide the LLM to identify the object mentioned in the query, while also understand and explain its reasoning process in the identification duration. Particularly, LLM is capable of mimicking the reasoning steps undertaken by humans. As illustrated in Figure 2(a), if the LLM get the object information, it can extract the objects relevant to the query sentence, i.e., targets *keyboard* and anchors *door*, and successfully identify the correct target *keyboard* by calculating its distance with *door*.

While LLMs show powerful human-like reasoning capabilities, they still have some limitations. First, it cannot handle the view-dependent issue such as *the right window*. This is because the 3D scene can freely rotate to different views while it keeps static in 2D images. LLMs usually make decisions by comparing their x-y values of 3D coordinates despite hinting it in the conversation. Second, mathematical calculation is a common weakness of LLMs but is necessary for 3DVG [11]. For example, in Figure 2(a), distance computing is crucial to solve the *closest* relations, whereas the LLMs cannot always provide accurate results. These two issues stem from LLM’s training limitations, which affect the reliability of the dialog with LLM approach.

3.2. 3D Visual Programming

To address the above two issues, we now introduce a new approach that generates visual programs through LLMs. As

View-independent	<i>near, close, next to, far, above, below, under, top, on, opposite, middle</i>
View-dependent	<i>front, behind, back, right, left, facing, leftmost, rightmost, looking, across, between</i>
Functional	<i>min, max, size, length, width</i>

Table 1. Common relations in 3DVG.

shown in Figure 2(b), we first construct a set of sample programs to encapsulate human-like problem-solving tactics in 3DVG. Each program includes a sequence of operations, where each operation contains a module name, several input parameters, and an assigned output variable. The output of each step can be reused in the subsequent step, thus creating an interlinked sequence that reflects logical reasoning within a visual context.

We transform the reasoning process of 3DVG into a scripted visual program. Specifically, we collect a set of in-context examples and the corresponding grounding descriptions, and then use LLMs to extrapolate new visual programs tailored to the task. For example, in Figure 2(b), we consider the task prompted by the following description:

The round cocktail table in
the corner of the room with the
blue and yellow poster.

In this case, the objective is to identify the round cocktail table, which can be transformed into an operation: $\text{BOX0} = \text{LOC}(\text{'round cocktail table'})$, where the LOC operator processes the textual query and outputs the bounding boxes for the target objects. We will elaborate the design of LOC module in Section 3.4. Nevertheless, since there may exist multiple similar objects in 3D scenarios, the identified results may not be unique. To overcome this issue, we further pinpoint the blue and yellow poster as an auxiliary reference point by an operation: $\text{BOX1} = \text{LOC}(\text{'blue and yellow poster'})$. Then, the CLOSEST module computes the proximity between BOX0 (potential tables) and BOX1 (poster), and selects the table closest to the poster as the result.

Table 1 summarizes the common relations in 3DVG. Based on this, we present the detailed visual program by developing three types of modules tailored for 3D contexts:

- **View-independent modules:** They operate on the 3D spatial relations between objects. For example, the CLOSEST module can discern proximity independent of the viewer’s position.
- **View-dependent modules:** They depend on the observer’s vantage point. For instance, the RIGHT module determines the right window (TARGET) when looking at cabinets (BOX1) from all windows (BOX0).
- **Functional modules:** They include multiple operations such as MIN and MAX, which select objects based on the extremal criteria.




Figure 3. Addressing view-dependent relations: A shift to 2D egocentric view.

These three types of modules allow the output of one operation to be fed into another operation, thus providing flexible composability. They not only facilitate structured and accurate inference sequences, but also integrate 3D and 2D data to yield a robust and interpretable result for 3DVG.

3.3. Addressing View-Dependent Relations

In this section, we discuss the intricacies of the view-dependent relations, which are essential for interpreting spatial relations within 3D space. Particularly, the main challenge is the dynamic nature of these relations that will change with the observer’s viewpoint. Although traditional supervised approaches can learn these relations implicitly, they cannot provide a definitive resolution.

On 2D planes, the relations, especially the *left* and *right* are well defined. More specifically, *right* often corresponds to the positive direction of the x-axis while the *left* implies the negative direction. Motivated by this, we adopt a 2D egocentric view approach to ensure a consistent frame of reference for the spatial relations in Table 1.

Our view-dependent modules accept a *target* argument and an optional *anchors* parameter. They output the target objects that fulfill the spatial relation to the anchors. When grounding queries do not specify targets, we treat *targets* as *anchors* as well. This approach aligns with our intuition, such as identifying *the left window* by treating all windows themselves as the anchors.

As shown in Figure 3, we assume there is a virtual camera in the center of the room, i.e., P_{center} , which can rotate to align with the location of the anchor objects, i.e., P_{O_a} . The 3D objects are projected onto a 2D plane from this vantage point. Assume that the orthogonal camera has an intrinsic matrix I , then the 2D projections can be obtained by

$$R, T = \text{Lookat}(P_{\text{center}}, P_{\text{O}_a}, up), \quad (1)$$

$$(u, v, w)^T = I \cdot (R|t) \cdot P, \quad (2)$$

where $\text{Lookat}(\cdot)$ is a view transformation function that computes the rotation matrix R and translation matrix T

[50], $P = (x, y, z, 1)^T$ denotes the 3D coordinate vector, u and v respectively signify the x-axis and y-axis on the 2D plane, and w is the depth value. According to the value of u of an object’s center, we can determine its left or right position — a lower u value indicates *left*. Similarly, w allows us to distinguish the *front* from *behind*. By synthesizing these concepts, we can define the *between* relation.

The transition from 3D to 2D egocentric perspective provides a clear and consistent solution to interpret view-dependent relations in 3D space, thus enhancing our model’s spatial reasoning ability.

3.4. Language-Object Correlation Module

Although our zero-shot 3DVG approach does not need extensive grounding annotations, it still requires a basic vision model for object localization. To overcome this issue, previous works [4, 62] usually use pre-trained 3D detectors [22, 38] to generate object proposals and the corresponding labels within a fixed vocabulary. However, this approach is restricted to a predefined object class set, thus limiting the scope of class prediction. To enable open-vocabulary segmentation, we develop an LOC module, combining the advantages of 3D and 2D networks to extend the labeling capability beyond the closed set. For example, in Figure 4, considering the operation: $\text{BOX0} = \text{LOC}(\text{‘round cocktail table’})$, we first filter a subset of objects whose predicted label is *table* using a 3D instance segmentation network [43]. Then we only need to identify a round cocktail table from this subset using the corresponding 2D imagery. By mapping each 3D proposal to its 2D image, we can extract the color and texture details pertinent to our query. To further pinpoint the round cocktail table shown in the Figure 4, we consider three types of 2D multi-modal models:

- Image classification models: We construct a dynamic vocabulary, including both the query term “round cocktail table” and the class “table” using popular tools such as CLIP [41]. Then we evaluate the cosine similarity between these terms and the imagery to find the best correlation to our query.
- Visual question answering models: We raise the question: *Is there a [query]?* to the model such as ViLT [25]. Then the model sifts through its dictionary to suggest the most likely answer, i.e., yes or no.
- General large models: We submit the same inquiry and anticipate a response based on the generated text. This process is crucial for verifying the alignment between the detected table and the query.

We shall note that our approach is not limited to specific 3D or 2D models, allowing versatile incorporation of various models. In the experiments, we will demonstrate that the benefit of the LOC modules by comparing with the 3D-only and 2D-only counterparts. Our design indicates a leap forward in 3D open-vocabulary instance segmentation and

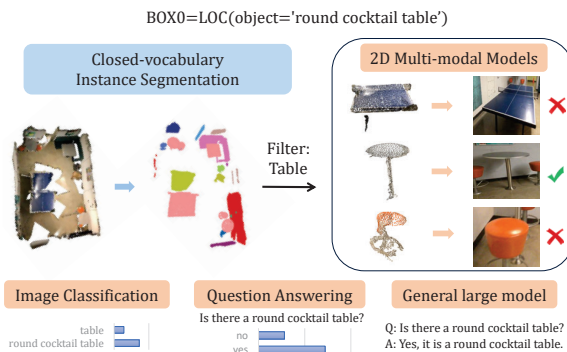


Figure 4. Illustration of the language-object correlation module.

can improve the object recognition accuracy in 3DVG.

4. Experiments

4.1. Experimental Settings

Datasets. We use two popular datasets, i.e., ScanRefer [4] and Nr3D [1] for experiments. ScanRefer is tailored for 3DVG that contains 51,500 sentence descriptions for 800 ScanNet scenes [9]. Nr3D is a human-written and free-form dataset for 3DVG, collected by 2-player reference game in 3D scenes. The sentences are divided into “easy” and “hard” subsets, where the target object only contains one same-class distractor in the “easy” subset but contains multiple ones in the “hard” subset. Depending on whether the sentence requires a specific viewpoint to ground the referred object, the dataset can also be partitioned into “view dependent” and “view independent” subsets. For both datasets, we evaluate the zero-shot approaches on the validation split.

Evaluation metrics. We consider two settings for performance evaluation. The first one mandates the generation of object proposals, aligning closely with real-world applications. The evaluation metrics are $\text{Acc}@0.25$ and $\text{Acc}@0.5$, representing the percentage of correctly predicted bounding boxes whose IoU exceeds 0.25 or 0.5 with the ground-truth, respectively. This is the default setting for ScanRefer dataset. The second one furnishes ground-truth object masks, necessitating only classification, with an objective to eradicate localization error and achieve high grounding accuracy. This is the default setting for Nr3D dataset.

Baselines. We use six supervised and two open-vocabulary 3D scene understanding approaches for performance comparison. For supervised approaches, ScanRefer [4] and ReferIt3DNet [1] encode the 3D point clouds and language separately, and then fuse them to rank the objects by predicted scores. TGNN [18] and InstanceRefer [62] make one further step by learning instance-wise features. 3DVG-Transformer [68] and BUTD-DETR [21] respectively utilize the Transformer [48] and DETR [3] architectures, representing the SoTA approaches. For open-vocabulary ap-

Methods	Supervision	Unique		Multiple		Overall	
		Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5
ScanRefer [4]	fully	65.0	43.3	30.6	19.8	37.3	24.3
TGNN [18]	fully	64.5	53.0	27.0	21.9	34.3	29.7
InstanceRefer [62]	fully	77.5	66.8	31.3	24.8	40.2	32.9
3DVG-Transformer [68]	fully	81.9	60.6	39.3	28.4	47.6	34.7
BUTD-DETR [21]	fully	84.2	66.3	46.6	35.1	52.2	39.8
LERF [24]	-	-	-	-	-	4.8	0.9
OpenScene [35]	-	20.1	13.1	11.1	4.4	13.2	6.5
Ours (2D only)	-	32.5	27.8	16.1	14.6	20.0	17.6
Ours (3D only)	-	57.1	49.4	25.9	23.3	33.1	29.3
Ours	-	63.8	58.4	27.7	24.6	36.4	32.7

Table 2. 3DVG results on ScanRefer validation set. The accuracy on the “unique” subset, “multiple” subset, and whole validation set are all provided. Following [4], we label the scene as “unique” if it only contains a single object of its class. Otherwise, we label it as “multiple”.

Method	Easy	Hard	Dep.	Indep.	Overall
ReferIt3DNet [1]	43.6	27.9	32.5	37.1	35.6
InstanceRefer [62]	46.0	31.8	34.5	41.9	38.8
3DVG-Transformer [68]	48.5	34.8	34.8	43.7	40.8
BUTD-DETR [21]	60.7	48.4	46.0	58.0	54.6
Ours (2D only)	29.4	18.4	23.0	23.9	23.6
Ours (3D only)	45.9	27.9	34.9	38.4	36.7
Ours	46.5	31.7	36.8	40.0	39.0

Table 3. Performance analysis of language grounding on Nr3D. We evaluate the top-1 accuracy using ground-truth boxes.

proaches, OpenScene [35] and LERF [24] aims to learn a 3D representation aligned with the 2D CLIP feature, thus enabling free-form language grounding. The query \mathcal{T} is processed by the CLIP text encoder, and its similarity is computed against the extracted point features. Finally, they cluster the points with the highest score to determine the target object.

4.2. Quantitative Results

ScanRefer. Table 2 provides a quantitative assessment of the proposed approach on the ScanRefer dataset. We can see that our zero-shot approach outperforms all baseline approaches. Specifically, our approach can achieve a 32.7 Acc@0.5 score, which surpasses part of supervised approaches such as ScanRefer and TGNN. On the other hand, the open-vocabulary approaches LERF and OpenScene can respectively achieve the overall accuracy of 4.8 and 13.2, even with the 0.25 IoU threshold. This is due to their limitations in reasoning and localization precision. Moreover, our zero-shot approach outperforms the approaches that only utilize the 3D or 2D information in the LOC module. This result demonstrates the effectiveness of incorporating visual programming and perception modules, highlighting our zero-shot approach in navigating 3DVG.

Nr3D. Table 3 shows the performance of different approaches on the Nr3D dataset, in which the ground-truth instance mask is also provided. We can see that our zero-

shot approach further excels the supervised approach InstanceRefer. Specifically, our zero-shot approach on the “view-dependent” split can achieve a 2% accuracy gain than the 3DVG-Transformer approach. This performance gain comes from the relation modules, strengthening the potential of our zero-shot approach for 3DVG tasks.

4.3. Qualitative Results

Figure 5 shows the visualizations of the selected samples from the ScanRefer validation set. The four columns present the ground-truth result, the supervised approach BUTD-DETR, the dialog with LLM, and the visual programming approaches, respectively. From Figure 5(a) and Figure 5(b), we can observe that the dialog with LLM and the visual programming approaches can achieve accurate prediction results for view-independent relations, i.e., (*above*, *under*) without much training. On the contrary, both the BUTD-DETR and the dialog with LLM approaches cannot address the view-dependent relations, i.e., (*left*, *front*), as shown in Figure 5(c) and Figure 5(d). The inherent uncertainty of these relations reflects the limitations of existing methods. However, our visual programming approach can leverage the 2D egocentric views, thus achieving accurate predictions in 3D scenarios.

Figure 5(e) presents a failure case, where the dialog with LLM approach cannot recognize *chair has wheels* since it lacks open-vocabulary detection ability. Besides, the visual programming approach makes wrong predictions because the LLM cannot correctly recognize the relation *pushed*. Fortunately, when we correct the program using the CLOEST module, the visual programming approach can make correct predictions.

4.4. Ablation Studies

Dialog with LLM vs. visual programming. We compare the performance of the two proposed zero-shot 3DVG approaches on the ScanRefer validation set with 700 exam-

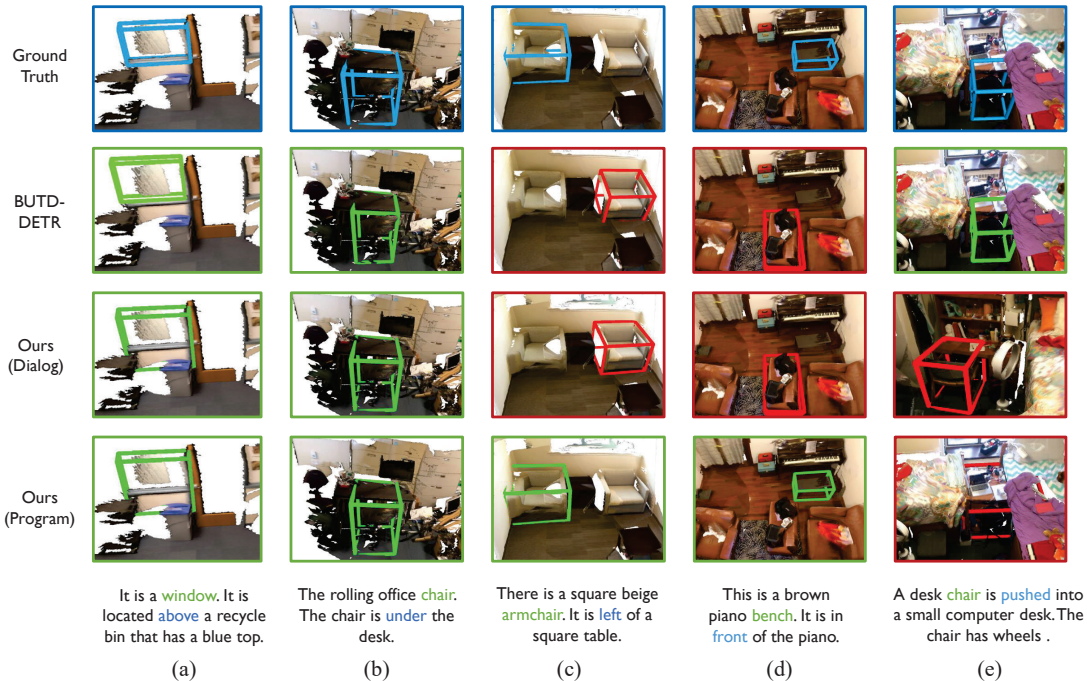


Figure 5. **Visualization results of 3D visual grounding.** Rendered images of 3D scans are presented, including the ground-truth (blue), incorrectly identified objects (red), and correctly identified objects (green).

ples. For both approaches, we use two GPT versions, i.e., GPT-3.5-turbo-0613 and GPT-4-0613. The cost of each GPT version depends on the number of input and output tokens. The experimental results are shown in Table 4. We can observe that for both zero-shot approaches, GPT4-based approach can achieve higher accuracy than GPT3.5-based approach, even it induces a larger economic cost. On the other hand, the visual programming approach always outperforms the dialog with LLM approach in terms of accuracy and cost, which demonstrates the effectiveness of the proposed visual programming approach. For other experiments, we use GPT3.5 to save cost.

Relation modules. We now ablate different relation modules in Section 3.2 to analyze their impact on the system performance. The most important view-dependent and view-independent modules are presented in Table 5 and 6, respectively. We can see that LEFT and RIGHT are the most important view-dependent relations, while CLOSEST is the most important view-independent relation. This result is coherent with our motivation and design.

LOC module. We juxtapose our approach by separately omitting the 3D component and 2D component. Both models utilize the instance mask prediction of Mask3D [43]. Particularly, the 2D-only model solely employs the paired 2D images for classification, while the 3D-only model just uses the 3D result. As can be seen from Tables 2 and 3, the 2D-only model performs worst when the images of in-

Method	LLM	Acc@0.5	Tokens	Cost
Dialog	GPT3.5	25.4	1959k	\$3.05
Dialog	GPT4	27.5	1916k	\$62.6
Program	GPT3.5	32.1	121k	\$0.19
Program	GPT4	35.4	115k	\$4.24

Table 4. Performance comparison of the dialog with LLM and the visual programming approaches.

LEFT	RIGHT	FRONT	BEHIND	BETWEEN	Accuracy
✓					26.5
✓	✓				32.4
✓	✓	✓			35.9
✓	✓	✓	✓		36.8
✓	✓	✓	✓	✓	38.4
✓	✓	✓	✓	✓	39.0

Table 5. Ablation study of different **view-dependent** modules.

door scenes are complicated and have domain gaps with the training samples. The 3D-only model performs better since it can utilize the geometric information and is trained on closed-set labels. Our full model can always achieve the best performance because it integrates the geometric distinctiveness of point clouds and the open-vocabulary ability of the image models.

Generalization. As discussed in Section 3.4, our framework has strong adaptability for a spectrum of 3D and 2D perception models. To validate this claim, we conduct ex-

CLOSEST	FARTHEST	LOWER	HIGHER	Accuracy
				18.8
✓				30.7
✓	✓			34.0
✓	✓	✓		36.8
✓	✓	✓	✓	39.0

Table 6. Ablation study of different **view-independent** modules.

periments using several representative models. For 3D perception, we utilize three backbones, i.e., PointNet++ [37], PointNeXt [39], and PointBERT [61]. For 2D perception, we use an image classification model proposed in [41], a visual question answering model in [25], and a general large model BLIP-2 [28] for testing. The results are shown in Tables 7 and 8. We can observe that our framework is compatible with other models. Also, it can leverage the advancements within both 2D and 3D foundational models to improve the performance. This cross-model effectiveness demonstrates the robustness and future-proof nature of our approach in the ever-evolving landscape of visual perception models.

2D Assistance	Unique	Multiple	Acc@0.25
CLIP	62.5	27.1	35.7
ViLT	60.3	27.1	35.1
BLIP-2	63.8	27.7	36.4

Table 7. Ablation study on different 2D models.

3D Backbone	View-dep.	View-indep.	Overall
PointNet++	35.8	39.4	38.2
PointBert	36.0	39.8	38.6
PointNeXt	36.8	40.0	39.0

Table 8. Ablation study on different 3D backbones.

Effect of prompt size. We use different numbers of in-context examples in the prompt for program generation. The result is shown in Figure 6. It can be seen that the performance on ScanRefer and Nr3d improve with the number of examples. This is because more examples can guide LLMs to handle more cases in the visual program generation process. Meanwhile, it still follows the law of diminishing marginal utility. Moreover, we test the voting technique [15] to aggregate the results from multiple runs, which brings some performance gains.

Error analysis. We select a representative subset for each dataset with around 100 samples and manually check the rationales offered by the visual program. As illustrated in Figure 7, the generation of accurate visual programs is the primary error source. Therefore, the performance can be improved by using more in-context examples and more powerful LLMs. The second error source is the object localization and classification, indicating that object detection and classification in 3D space remains a critical component.




Figure 6. Ablation study on the number of in-context examples. The performance on Nr3D and ScanRefer improves with the number of in-context examples.




Figure 7. Breakdown of error sources.

5. Conclusion

In this paper, we present a novel zero-shot approach for 3DVG to eliminate the need for extensive annotations and predefined vocabularies. A vanilla dialog with LLM approach is first proposed by taking interactive dialog with LLMs. A visual programming approach is further developed, which leverages three types of modules to navigate the intricate 3D relations. To adapt to open-vocabulary scenarios, we also develop a LOC module to seamlessly integrate both 3D and 2D features. Experimental results demonstrate the superiority of the proposed approach and highlight its potential to advance the field of 3DVG.

6. Acknowledgement

This work was supported by NSFC with Grant No. 62293482, by the Basic Research Project No. HZQB-KCZYZ-2021067 of Hetao Shenzhen HK S&T Cooperation Zone, by Shenzhen General Program No. JCYJ2022053014360001, by Shenzhen-Hong Kong Joint Funding No. SGDXX20211123112401002, by the National Key R&D Program of China with grant No. 2018YFB1800800, by the Shenzhen Outstanding Talents Training Fund 202002, by Guangdong Research Project No. 2017ZT07X152 and No. 2019CX01X104, by the Guangdong Provincial Key Laboratory of Future Networks of Intelligence (Grant No. 2022B1212010001), by the Guangdong Provincial Key Laboratory of Big Data Computing, CHUK-Shenzhen, by the NSFC 61931024&12326610, by the Shenzhen Key Laboratory of Big Data and Artificial Intelligence (Grant No. ZDSYS201707251409055), and the Key Area R&D Program of Guangdong Province with grant No. 2018B03033800, by Tencent & Huawei Open Fund.

References

- [1] 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. Springer, 2020. 1, 2, 5, 6
- [2] Anthony Brohan, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Alexander Herzog, Daniel Ho, Julian Ibarz, Alex Irpan, Eric Jang, Ryan Julian, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, pages 287–318. PMLR, 2023. 2
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 5
- [4] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *European Conference on Computer Vision*, pages 202–221. Springer, 2020. 1, 2, 5, 6
- [5] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021. 2
- [6] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Language conditioned spatial relation reasoning for 3d object grounding. *Advances in Neural Information Processing Systems*, 35:20522–20535, 2022. 2
- [7] Zhenyu Chen, Ronghang Hu, Xinlei Chen, Matthias Nießner, and Angel X Chang. Unit3d: A unified transformer for 3d dense captioning and visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18109–18119, 2023. 2
- [8] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, pages 3075–3084, 2019. 2
- [9] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 2, 5
- [10] John David N Dionisio, William G Burns Iii, and Richard Gilbert. 3d virtual worlds and the metaverse: Current status and future possibilities. *ACM Computing Surveys (CSUR)*, 45(3):1–38, 2013. 1, 2
- [11] Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jian, Bill Yuchen Lin, Peter West, Chandra Bhagavathula, Ronan Le Bras, Jena D Hwang, et al. Faith and fate: Limits of transformers on compositionality. *arXiv preprint arXiv:2305.18654*, 2023. 3
- [12] Qi Feng, Vitaly Ablavsky, and Stan Sclaroff. Cityflow-nl: Tracking and retrieval of vehicles at city scale by natural language descriptions. *arXiv preprint arXiv:2101.04741*, 2021. 1, 2
- [13] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *European Conference on Computer Vision*, pages 540–557. Springer, 2022. 2
- [14] Zoey Guo, Yiwen Tang, Ray Zhang, Dong Wang, Zhigang Wang, Bin Zhao, and Xuelong Li. Viewrefer: Grasp the multi-view knowledge for 3d visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15372–15383, 2023. 2
- [15] Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14953–14962, 2023. 2, 8
- [16] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36:20482–20494, 2023. 2
- [17] Joy Hsu, Jiayuan Mao, and Jiajun Wu. Ns3d: Neuro-symbolic grounding of 3d objects and relations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2614–2623, 2023. 2
- [18] Pin-Hao Huang, Han-Hung Lee, Hwann-Tzong Chen, and Tyng-Luh Liu. Text-guided graph neural networks for referring 3d instance segmentation. *35th AAAI Conference on Artificial Intelligence*, 2021. 5, 6
- [19] Shijia Huang, Yilun Chen, Jiaya Jia, and Liwei Wang. Multi-view transformer for 3d visual grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15524–15533, 2022. 1
- [20] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pages 9118–9147. PMLR, 2022. 2
- [21] Ayush Jain, Nikolaos Gkanatsios, Ishita Mediratta, and Katerina Fragkiadaki. Bottom up top down detection transformers for language grounding in images and point clouds. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*, pages 417–433. Springer, 2022. 5, 6
- [22] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4867–4876, 2020. 1, 2, 5
- [23] Zhao Jin, Munawar Hayat, Yuwei Yang, Yulan Guo, and Yinjie Lei. Context-aware alignment and mutual masking for 3d-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10984–10994, 2023. 2
- [24] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lrf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 2, 6
- [25] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region su-

- pervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021. 5, 8
- [26] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022. 2
- [27] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *International Conference on Learning Representations*, 2022. 2
- [28] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 8
- [29] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7061–7070, 2023. 2
- [30] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9493–9500. IEEE, 2023. 2
- [31] Ze Liu, Zheng Zhang, Yue Cao, Han Hu, and Xin Tong. Group-free 3d object detection via transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2949–2958, 2021. 2
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [33] Stylianos Mystakidis. Metaverse. *Encyclopedia*, 2(1):486–497, 2022. 1, 2
- [34] OpenAI OpenAI. Gpt-4 technical report. 2023. 2
- [35] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 815–824, 2023. 2, 6
- [36] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 2
- [37] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 2, 8
- [38] 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. 1, 2, 5
- [39] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *Advances in Neural Information Processing Systems*, 35:23192–23204, 2022. 8
- [40] Xing-yue Qiu, Chuang-Kai Chiu, Lu-Lu Zhao, Cai-Feng Sun, and Shu-jie Chen. Trends in vr/ar technology-supporting language learning from 2008 to 2019: A research perspective. *Interactive Learning Environments*, 31(4):2090–2113, 2023. 1
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 5, 8
- [42] Junha Roh, Karthik Desingh, Ali Farhadi, and Dieter Fox. Languagerefer: Spatial-language model for 3d visual grounding. In *5th Annual Conference on Robot Learning*, 2021. 2
- [43] Jonas Schult, Francis Engelmann, Alexander Hermans, Or Litany, Siyu Tang, and Bastian Leibe. Mask3d: Mask transformer for 3d semantic instance segmentation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8216–8223. IEEE, 2023. 2, 5, 7
- [44] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015. 2
- [45] Dídac Surís, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for reasoning. *arXiv preprint arXiv:2303.08128*, 2023. 2
- [46] Ayça Takmaz, Elisabetta Fedele, Robert W Sumner, Marc Pollefeys, Federico Tombari, and Francis Engelmann. Openmask3d: Open-vocabulary 3d instance segmentation. *arXiv preprint arXiv:2306.13631*, 2023. 2
- [47] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 2
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 5
- [49] Sai Vemprala, Rogerio Bonatti, Arthur Buckner, and Ashish Kapoor. Chatgpt for robotics: Design principles and model abilities. *Microsoft Auton. Syst. Robot. Res.*, 2:20, 2023. 2
- [50] John Vince and John A Vince. *Mathematics for computer graphics*. Springer, 2006. 5
- [51] Thang Vu, Kookhoi Kim, Tung M Luu, Thanh Nguyen, and Chang D Yoo. Softgroup for 3d instance segmentation on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2708–2717, 2022. 2
- [52] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. Rio: 3d object instance re-localization in changing indoor environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7658–7667, 2019. 2

- [53] Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6629–6638, 2019. 1, 2
- [54] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022. 2
- [55] Wei Wei. Research progress on virtual reality (vr) and augmented reality (ar) in tourism and hospitality: A critical review of publications from 2000 to 2018. *Journal of Hospitality and Tourism Technology*, 10(4):539–570, 2019. 1
- [56] Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*, 2023. 2
- [57] Yanmin Wu, Xinhua Cheng, Renrui Zhang, Zesen Cheng, and Jian Zhang. Eda: Explicit text-decoupling and dense alignment for 3d visual grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19231–19242, 2023. 2
- [58] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9068–9079, 2018. 1, 2
- [59] Xu Yan, Zhihao Yuan, Yuhao Du, Yinghong Liao, Yao Guo, Shuguang Cui, and Zhen Li. Comprehensive visual question answering on point clouds through compositional scene manipulation. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2
- [60] Zhengyuan Yang, Songyang Zhang, Liwei Wang, and Jiebo Luo. Sat: 2d semantics assisted training for 3d visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1856–1866, 2021. 1, 2
- [61] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19313–19322, 2022. 8
- [62] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1791–1800, 2021. 1, 2, 5, 6
- [63] Zhihao Yuan, Xu Yan, Zhuo Li, Xuhao Li, Yao Guo, Shuguang Cui, and Zhen Li. Toward explainable and fine-grained 3d grounding through referring textual phrases. *arXiv preprint arXiv:2207.01821*, 2022. 2
- [64] Zhihao Yuan, Xu Yan, Yinghong Liao, Yao Guo, Guanbin Li, Zhen Li, and Shuguang Cui. X-trans2cap: Cross-modal knowledge transfer using transformer for 3d dense captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [65] Yihan Zeng, Chenhan Jiang, Jiageng Mao, Jianhua Han, Chaoqiang Ye, Qingqiu Huang, Dit-Yan Yeung, Zhen Yang, Xiaodan Liang, and Hang Xu. Clip2: Contrastive language-image-point pretraining from real-world point cloud data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15244–15253, 2023. 2
- [66] Jiazhao Zhang, Chenyang Zhu, Lintao Zheng, and Kai Xu. Fusion-aware point convolution for online semantic 3d scene segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4534–4543, 2020. 2
- [67] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. 2
- [68] Lichen Zhao, Daigang Cai, Lu Sheng, and Dong Xu. 3dvg-transformer: Relation modeling for visual grounding on point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2928–2937, 2021. 5, 6
- [69] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*, 2022. 2
- [70] Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyu Guo, Ziyao Zeng, Zipeng Qin, Shanghang Zhang, and Peng Gao. Point-clip v2: Prompting clip and gpt for powerful 3d open-world learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2639–2650, 2023. 2