

Z3D: Zero-Shot 3D Visual Grounding from Images

Nikita Drozdov¹ Andrey Lemeshko² Nikita Gavrillov¹
Anton Konushin¹ Danila Rukhovich³ Maksim Kolodiazhnyi^{1†}

¹Lomonosov Moscow State University ²Higher School of Economics

³M3L Lab, Institute of Mechanics, Armenia

Abstract

3D visual grounding (3DVG) aims to localize objects in a 3D scene based on natural language queries. In this work, we explore zero-shot 3DVG from multi-view images alone, without requiring any geometric supervision or object priors. We introduce Z3D, a universal grounding pipeline that flexibly operates on multi-view images while optionally incorporating camera poses and depth maps. We identify key bottlenecks in prior zero-shot methods causing significant performance degradation and address them with (i) a state-of-the-art zero-shot 3D instance segmentation method to generate high-quality 3D bounding box proposals and (ii) advanced reasoning via prompt-based segmentation, which utilizes full capabilities of modern VLMs. Extensive experiments on the ScanRefer and Nr3D benchmarks demonstrate that our approach achieves state-of-the-art performance among zero-shot methods.

1 Introduction

3D visual grounding (3DVG) seeks to localize target objects in a scene based on natural language descriptions. It is a fundamental capability for embodied AI, robotics, and human–scene interaction, where agents must reason jointly over language, visual appearance, and spatial structure.

3DVG from point clouds is the most natural and well-studied scenario. The recent emergence of LLM allows reducing the 3D labeling burden, leveraging generalization capabilities of large models to eliminate the need for 3D supervision. Earlier methods, such as Vil3DRel (Chen et al., 2022) and MiKASA (Chang et al., 2024), and their recent follow-ups, SceneVerse (Jia et al., 2024), LIBA (Wang et al., 2025d), ROSS3D (Wang et al., 2025a), LLaVA-3D (Zhu et al., 2025), Video-3D LLM (Zheng et al., 2025b), GPT4Scene (Qi et al.,

2025), and MPEC (Wang et al., 2025c), rely on full supervision, with both 3D bounding boxes and texts being exposed to the model during the training phase. ZSVG3D (Yuan et al., 2024b), CSVG (Yuan et al., 2024a), SeeGround (Li et al., 2025), LaSP (Mi et al., 2025b), EaSe (Mi et al., 2025a), SPAZER (Jin et al., 2025) do not require texts for training but still exploit annotated 3D bounding boxes. The training-free, zero-shot approach is represented with LLM-Grounder (Yang et al., 2024) and VLM-Grounder (Xu et al., 2024). Both of them rely on proprietary VLMs but still use non-generative language models, such as BERT and CLIP, in critical parts of the pipeline, which severely limits their performance. Overcoming this weakness, we achieve up to +40% in accuracy over prior state-of-the-art. Besides, we investigate another cause of poor performance of prior zero-shot approaches: the insufficient quality of candidate object proposals. Using state-of-the-art zero-shot 3D instance segmentation method, we generate high-quality proposals that serve as a strong basis for further VLM reasoning.

3DVG from images Most existing zero-shot methods assume access to explicit 3D representations, such as point clouds, depth maps, or pre-built scene reconstructions, which restricts their applicability in real-world settings. Fully supervised SPAR (Zhang et al., 2025) and VG LLM (Zheng et al., 2025a) claim to be image-based, but both require ground truth camera pose: SPAR uses them during the test phase, while VG LLM is exposed to camera poses during the training. Zero-shot 3DVG is addressed with modern VLMs, e.g., Qwen3-VL (Bai et al., 2025) and Seed1.5-VL (Guo et al., 2025), but only from single-view images, which is a critical limitation, since we aim for scene-level understanding. In this work, we study zero-shot 3D visual grounding from multi-view images alone and present Z3D, a universal grounding pipeline

[†]Corresponding author: kolodyazhnyi@my.msu.ru
Code available at <https://github.com/col14m/z3d>

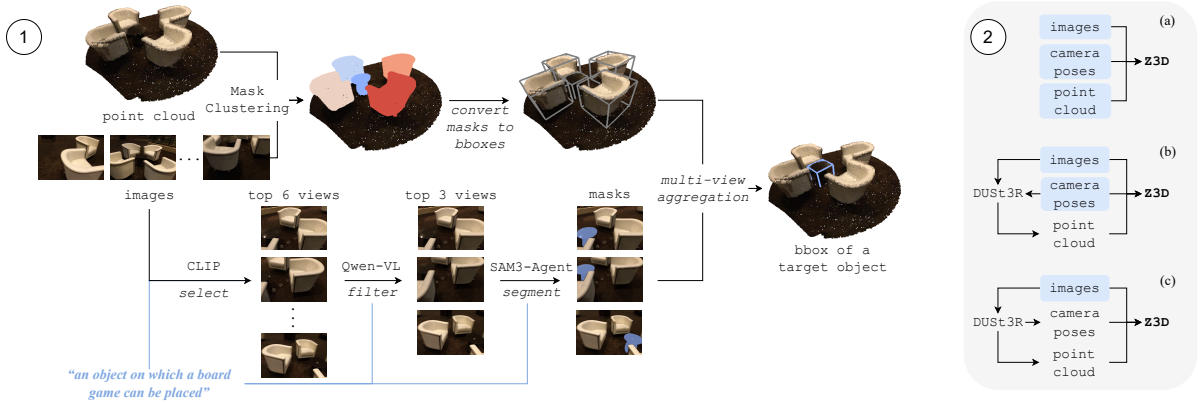


Figure 1: (1) Inference pipeline of Z3D given point cloud inputs. Z3D leverages MaskClustering (Yan et al., 2024) to predict class-agnostic 3D bounding boxes from a point cloud and images. The most relevant frames are chosen based on their CLIP embeddings and further filtered with VLM. Object masks are annotated using SAM3-Agent (Carion et al., 2025) and lifted to 3D. The final bounding box is selected from MaskClustering proposals through a mask voting scheme. (2) Our approach flexibly accommodates different input modalities, including images, camera poses, and depth maps. When camera poses and depth maps are unavailable, they are estimated with DUST3R (Wang et al., 2024).

that operates on images and can optionally incorporate camera poses and depth maps when available.

Our contribution is twofold:

- we improve components of the existing VLM-based 3DVG pipeline, achieving state-of-the-art results in 3DVG from point clouds;
- we extend our method to handle various inputs, thus being the first to address 3DVG in a zero-shot setting from images only.

2 Proposed Method

2.1 3DVG With Depth

3DVG implies estimating a 3D bounding box of a target object in a scene given a query in natural language. Existing zero-shot methods rely on VLMs to handle fuzzy indirect references and require images to proceed. The target object might be visible in a subset of frames, so the task of selecting the most relevant views arises naturally. Then, the target object must be located in those views and lifted to 3D space. Respectively, the VLM-based 3DVG workflow can be broadly decomposed into (i) view selection, (ii) 2D object segmentation, and (iii) 2D-to-3D lifting. The only existing zero-shot baseline showing reasonable performance, VLM-Grounder (Xu et al., 2024), follows this paradigm; in Z3D, we propose modifications of each step that push the quality to the state-of-the-art level.

View selection. Processing all images of a scene with VLM is time-consuming, so optimizations

are inevitable. VLM-Grounder packs images in a grid to minimize the number of calls and iteratively narrows the search scope until finding the best views. In contrast, we employ a two-stage strategy to efficiently identify informative observations. First, views are preselected using CLIP so that the six most similar frames for a given query pass the filter. Then, VLM selects the best three views. Therefore, the search space is reduced using a lightweight model and simple selection strategy, while the accuracy benefits of a more powerful but computationally expensive approach are retained.

2D object segmentation. VLM-Grounder combines Grounding DINO (Ren et al., 2024) with SAM to localize and segment objects in a frame, respectively. This standard pipeline, powered by BERT (Devlin et al., 2019), is tailored for open-vocabulary segmentation with explicit object descriptions, while less specific 3DVG prompts might cause performance degradation. To overcome this limitation, Z3D employs a SAM3-Agent (Carion et al., 2025) for zero-shot high-quality object segmentation. Guided by VLM reasoning, the agent iteratively generates and refines segmentation prompts, enabling precise instance extraction without geometric supervision or object priors.

2D-to-3D lifting. Each object mask can be lifted into 3D space, giving a partial point cloud of an object. VLM-Grounder simply unites all partial point clouds and encloses them with a 3D bounding box to create a proposal; hence, outliers have a

large impact on the estimated size and shape of an object, making the whole procedure prone to noise. Differently, we first obtain class-agnostic 3D object proposals and use 2D object masks to select the best one. To this end, we leverage MaskClustering (Yan et al., 2024), a zero-shot 3D instance segmentation method that takes point clouds as inputs and produces object 3D masks, which we convert into 3D bounding boxes. Segmentation masks from the top-3 views are lifted to 3D and matched against the MaskClustering proposals. A proposal with the highest 3D IoU with a mask gets one vote; the final proposal is the most voted candidate, or, in ambiguous cases, the one voted by a mask in the frame with the highest CLIP and VLM relevance scores.

2.2 3DVG From Images

When depths or point clouds are unavailable, we bridge the gap between sole visual inputs and real geometry with a 3D reconstruction method. Specifically, we use DUS_t3R (Wang et al., 2024): with its ability to handle omnimodal inputs, it fits perfectly into both images-only and images + camera poses scenarios. With DUS_t3R, our processing pipeline remains purely zero-shot, since, contrary to some latest methods (Wang et al., 2025b) it was not trained on ScanNet (Dai et al., 2017). Given images, DUS_t3R returns dense depth maps and infers poses when they are not available. The depths are fused into a TSDF volume using ground truth or predicted camera poses. Finally, a point cloud is extracted using the marching cubes algorithm.

3 Experiments

We evaluate our approach on the ScanRefer (Chen et al., 2020) and Nr3D (Achlioptas et al., 2020) benchmarks. ScanRefer annotates ScanNet scenes with over 51K human-written query-object pairs, where the goal is to localize the target object by predicting its 3D bounding box from scene point clouds and language queries. Following standard practice, we report Acc@0.25 and Acc@0.5, defined as the percentage of predictions whose 3D IoU with ground truth exceeds 0.25 and 0.5, respectively. The Nr3D dataset contains 41K language queries over ScanNet scenes and provides ground-truth 3D bounding boxes without class labels. The task is to select the most relevant candidate object, which is evaluated using top-1 accuracy.

3DVG with depths 3DVG methods that use depth, point clouds, or other sources of spatial information represent the most well studied setting. These approaches can be categorized based on their exposure to 3D bounding boxes providing geometric supervision: (i) methods trained with ground-truth bounding boxes (e.g., approaches using Mask3D proposals), (ii) methods provided with bounding boxes at inference time (as in the Nr3D benchmark), and (iii) methods that are not exposed to bounding boxes at any stage. When using Mask3D as a proposal generator, Z3D demonstrates substantial improvements over prior methods that are exposed to bounding boxes in the training set (Tab. 1, row 9). In the inference-time 3D bounding box setting, Z3D achieves state-of-the-art top-1 accuracy on Nr3D (Tab. 2), indicating that the gain stems not only from proposal quality but also from the effectiveness of the other components of our pipeline. Finally, in the purely zero-shot setting without any bounding-box supervision, Z3D significantly outperforms all competitors, achieving +38.7 Acc@0.5 over OpenScene on ScanRefer (Tab. 1, row 12).

3DVG from images For image-based 3DVG, both with and without camera poses, existing approaches are fully supervised; therefore, we report their results for reference only. To establish a meaningful baseline, we combine DUS_t3R with a state-of-the-art point cloud-based 3DVG method. Specifically, we adopt SeeGround (Li et al., 2025), which is highly competitive in depth-aware settings. While the original SeeGround uses Mask3D to generate proposals, in this series of experiments, we replace it with MaskClustering to keep the whole pipeline zero-shot. As shown in Tab. 1 (rows 15, 18), Z3D consistently outperforms SeeGround on DUS_t3R reconstructions, demonstrating that its advantages are preserved regardless of the reconstruction approach. Notably, Z3D establishes a new state of the art in the posed-images setting, surpassing the previous state-of-the-art fully supervised method, SPAR (Zhang et al., 2025).

Ablation study To quantify the contribution of each component, we conduct an ablation study by progressively building our pipeline from a simple baseline. The original MaskClustering method, designed for open-vocabulary 3D instance segmentation, already achieves 27.6 Acc@0.5. Incorporating CLIP-based view selection (top-1 view) and SAM3-Agent for object segmentation increases

Method	Venue	Supervision		Unique		Multiple		Overall	
		bboxes	texts	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5
<i>Images + camera poses + depths</i>									
LLaVA-3D	ICCV'25	✓	✓	-	-	-	-	50.1	42.7
Video-3D LLM	CVPR'25	✓	✓	86.6	77.0	50.9	45.0	57.9	51.2
ROSS3D	ICCV'25	✓	✓	87.2	77.4	54.8	48.9	61.1	54.4
LIBA	AAAI'25	✓	✓	88.8	74.3	54.4	44.4	59.6	49.0
GPT4Scene	-	✓	✓	90.3	83.7	56.4	50.9	62.6	57.0
ZSVG3D	CVPR'24	Mask3D	✗	63.8	58.4	27.7	24.6	36.4	32.7
CSVG	BMVC'25	Mask3D	✗	68.8	61.2	38.4	27.3	49.6	39.8
SeeGround	CVPR'25	Mask3D	✗	75.7	68.9	34.0	30.0	44.1	39.4
Z3D	-	Mask3D	✗	82.3	74.8	51.5	45.7	58.9	52.7
OpenScene	CVPR'23	✗	✗	20.1	13.1	11.1	4.4	13.2	6.5
LLM-Grounder	ICRA'24	✗	✗	-	-	-	-	17.1	5.3
Z3D	-	✗	✗	73.9	64.0	47.8	40.3	54.2	46.0
<i>Images + camera poses</i>									
SPAR	NIPS'25	✓	✓	-	-	-	-	31.9	12.4
DUST3R → SeeGround	-	✗	✗	44.5	27.4	21.1	12.6	26.8	16.2
DUST3R → Z3D	-	✗	✗	56.7	32.0	38.4	22.6	42.8	24.8
<i>Images</i>									
VG LLM	NIPS'25	✓	✓	-	-	-	-	41.6	14.9
DUST3R → SeeGround	-	✗	✗	35.2	17.5	13.9	5.2	19.0	8.2
DUST3R → Z3D	-	✗	✗	42.7	21.9	27.5	10.1	31.2	12.9

Table 1: Results on ScanRefer. Queries are categorized as “*Unique*”, with only one object of the target class in the scene, or “*Multiple*”, with other objects of the same class. Results of our method are blue, our baselines are marked with a lighter hue.

Method	Easy	Hard	Dep.	Indep.	Overall
<i>Fully supervised</i>					
MiKASA	69.7	59.4	65.4	64.0	64.4
ViL3DRel	70.2	57.4	62.0	64.5	64.4
SceneVerse	72.5	57.8	56.9	67.9	64.9
MPEC	-	-	-	-	66.7
<i>Zero-shot (use gt object class)</i>					
CSVG	67.1	51.3	53.0	62.5	59.2
EaSe	-	-	-	-	67.8
Transcrib3D	79.7	60.3	60.1	75.4	70.2
<i>Zero-shot</i>					
ZSVG3D	46.5	31.7	36.8	40.0	39.0
SeeGround	54.5	38.3	42.3	48.2	46.1
LaSP	60.7	45.3	49.2	54.7	52.9
EaSe	-	-	-	-	52.9
SPAZER	62.4	46.9	49.9	56.8	54.3
Z3D	62.6	47.5	50.7	57.1	54.8

Table 2: Results of methods using depths on Nr3D, specified by the query type: “*Easy*” (one distractor) “*Hard*” (multiple distractors), “*View-Dependent*” or “*View-Independent*” based on viewpoint requirements for grounding.

Module	Acc@0.25	Acc@0.5
MaskClustering (Yan et al., 2024)	32.0	27.6
+ CLIP + SAM3-Agent	51.0	42.8
+ VLM view selection	53.0	44.8
+ multi-view aggregation	54.2	46.0

Table 3: Ablation study of model components on ScanRefer, with images, camera poses, and depths as inputs.

performance to 42.8. When the number of views selected by CLIP is increased to 6, and followed by top-1 view selection with VLM, accuracy further improves to 44.8. Finally, aggregating predictions across the top-3 views selected by the VLM results in the best performance, reaching 46.0.

4 Conclusion

We presented Z3D, a universal pipeline for zero-shot 3D visual grounding. We identified proposal quality and underutilization of VLMs as key bottlenecks in prior methods and addressed them with zero-shot 3D instance segmentation and VLM reasoning. Our approach flexibly accommodates different input modalities, including multi-view images, camera poses, and depth maps. Evaluations on ScanRefer and Nr3D demonstrate that Z3D achieves state-of-the-art performance among zero-shot approaches across multiple settings. We hope this work encourages further research on image-based and supervision-free 3D visual grounding, paving the way toward more practical and scalable 3D scene understanding systems.

This work was supported by the The Ministry of Economic Development of the Russian Federation in accordance with the subsidy agreement (agreement identifier 000000C313925P4H0002; grant No 139-15-2025-012).

Limitations

While introducing advanced VLM reasoning about the selected frames, our method still uses CLIP to pre-select frame candidates and therefore can be limited by CLIP’s ability to analyze complex concepts from subtle cues rather than direct descriptions. Moreover, in image-only scenarios, the performance of our method heavily depends on the quality of underlying 3D reconstruction. While DUS3R is known to perform robustly on ScanNet captures, the similar quality is not guaranteed for other scenes. In terms of performance, one of the processing bottlenecks of Z3D is MaskClustering, which adds a significant computation overhead; the component-wise time analysis can be found in the supplementary materials.

References

- Panos Achlioptas, Ahmed Abdelreheem, Fei Xia, Mohamed Elhoseiny, and Leonidas Guibas. 2020. Referit3d: Neural listeners for fine-grained 3d object identification in real-world scenes. In *European conference on computer vision*, pages 422–440. Springer.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, and 1 others. 2025. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.
- Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, and 1 others. 2025. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*.
- Chun-Peng Chang, Shaoxiang Wang, Alain Pagani, and Didier Stricker. 2024. Mikasa: Multi-key-anchor & scene-aware transformer for 3d visual grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14131–14140.
- Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. 2020. Scanrefer: 3d object localization in rgb-d scans using natural language. In *European conference on computer vision*, pages 202–221. Springer.
- Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. 2022. Language conditioned spatial relation reasoning for 3d object grounding. *Advances in neural information processing systems*, 35:20522–20535.
- Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Dong Guo, Faming Wu, Feida Zhu, Fuxing Leng, Guang Shi, Haobin Chen, Haoqi Fan, Jian Wang, Jianyu Jiang, Jiawei Wang, and 1 others. 2025. Seed1. 5-vl technical report. *arXiv preprint arXiv:2505.07062*.
- Baoxiong Jia, Yixin Chen, Huangyue Yu, Yan Wang, Xuesong Niu, Tengyu Liu, Qing Li, and Siyuan Huang. 2024. Sceneverse: Scaling 3d vision-language learning for grounded scene understanding. In *European Conference on Computer Vision*, pages 289–310. Springer.
- Zhao Jin, Rong-Cheng Tu, Jingyi Liao, Wenhao Sun, Xiao Luo, Shunyu Liu, and Dacheng Tao. 2025. Spacer: Spatial-semantic progressive reasoning agent for zero-shot 3d visual grounding. *arXiv preprint arXiv:2506.21924*.
- Rong Li, Shijie Li, Lingdong Kong, Xulei Yang, and Junwei Liang. 2025. Seeground: See and ground for zero-shot open-vocabulary 3d visual grounding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3707–3717.
- Boyu Mi, Hanqing Wang, Tai Wang, Yilun Chen, and Jiangmiao Pang. 2025a. Evolving symbolic 3d visual grounder with weakly supervised reflection. *arXiv preprint arXiv:2502.01401*.
- Boyu Mi, Hanqing Wang, Tai Wang, Yilun Chen, and Jiangmiao Pang. 2025b. Language-to-space programming for training-free 3d visual grounding. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 3844–3864.
- Zhangyang Qi, Zhixiong Zhang, Ye Fang, Jiaqi Wang, and Hengshuang Zhao. 2025. Gpt4scene: Understand 3d scenes from videos with vision-language models. *arXiv preprint arXiv:2501.01428*.
- Tianhe Ren, Qing Jiang, Shilong Liu, Zhaoyang Zeng, Wenlong Liu, Han Gao, Hongjie Huang, Zhengyu Ma, Xiaoke Jiang, Yihao Chen, and 1 others. 2024. Grounding dino 1.5: Advance the "edge" of open-set object detection. *arXiv preprint arXiv:2405.10300*.
- Zachary Teed and Jia Deng. 2021. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569.
- Haochen Wang, Yucheng Zhao, Tiancai Wang, Haoqiang Fan, Xiangyu Zhang, and Zhaoxiang Zhang. 2025a. Ross3d: Reconstructive visual instruction tuning with 3d-awareness. *arXiv preprint arXiv:2504.01901*.

- Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. 2025b. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306.
- Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. 2024. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709.
- Yan Wang, Baoxiong Jia, Ziyu Zhu, and Siyuan Huang. 2025c. Masked point-entity contrast for open-vocabulary 3d scene understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14125–14136.
- Yuan Wang, Ya-Li Li, WU Eastman ZY, and Shengjin Wang. 2025d. Liba: Language instructed multi-granularity bridge assistant for 3d visual grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 8114–8122.
- Runsen Xu, Zhiwei Huang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. 2024. Vlm-grounder: A vlm agent for zero-shot 3d visual grounding. *arXiv preprint arXiv:2410.13860*.
- Mi Yan, Jiazhao Zhang, Yan Zhu, and He Wang. 2024. Maskclustering: View consensus based mask graph clustering for open-vocabulary 3d instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28274–28284.
- Jianing Yang, Xuweiyi Chen, Shengyi Qian, Nikhil Madaan, Madhavan Iyengar, David F Fouhey, and Joyce Chai. 2024. Llm-grounder: Open-vocabulary 3d visual grounding with large language model as an agent. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7694–7701. IEEE.
- Qihao Yuan, Jiaming Zhang, Kailai Li, and Rainer Stiefelwagen. 2024a. Solving zero-shot 3d visual grounding as constraint satisfaction problems. *arXiv preprint arXiv:2411.14594*.
- Zhihao Yuan, Jinke Ren, Chun-Mei Feng, Hengshuang Zhao, Shuguang Cui, and Zhen Li. 2024b. Visual programming for zero-shot open-vocabulary 3d visual grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20623–20633.
- Jiahui Zhang, Yurui Chen, Yanpeng Zhou, Yueming Xu, Ze Huang, Jilin Mei, Junhui Chen, Yu-Jie Yuan, Xinyue Cai, Guowei Huang, and 1 others. 2025. From flatland to space: Teaching vision-language models to perceive and reason in 3d. *arXiv preprint arXiv:2503.22976*.
- Duo Zheng, Shijia Huang, Yanyang Li, and Liwei Wang. 2025a. Learning from videos for 3d world: Enhancing mllms with 3d vision geometry priors. *arXiv preprint arXiv:2505.24625*.
- Duo Zheng, Shijia Huang, and Liwei Wang. 2025b. Video-3d llm: Learning position-aware video representation for 3d scene understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8995–9006.
- Chenming Zhu, Tai Wang, Wenwei Zhang, Jiangmiao Pang, and Xihui Liu. 2025. Llava-3d: A simple yet effective pathway to empowering llms with 3d capabilities. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4295–4305.

A Quantitative Results

Some recent methods follow the alternative evaluation protocol, proposed in VLM-Grounder (Xu et al., 2024). This protocol implies testing on 250-scene subsets of ScanRefer and Nr3D rather than their full validation splits. The results on ScanRefer and Nr3D are reported in Tab. 4 and 5, respectively; clearly, Z3D scores the best in both benchmarks. Z3D shines in the pure zero-shot scenario (w/o access to ground truth 3D bounding boxes), achieving +18.8 Acc@0.5 w.r.t. VLM-Grounder on ScanRefer. On Nr3D, our method outperforms previous state-of-the-art SPAZER (Jin et al., 2025) using the same Qwen2.5-VL-72B, and even beats VLM-Grounder based on much more powerful proprietary GPT-4o.

B Ablation Studies

In this Section, all results are reported on 250-scenes subsets.

VLM size In Tab. 5, we vary the size of Qwen3-VL-Thinking serving as our VLM reasoner, and report the quality achieved with each model size. Even with a 30B model, Z3D outperforms prior methods, and using a larger 235B model pushes the quality even further.

Mask3D vs. MaskClustering The key difference between MaskClustering and Mask3D is that the first is a pure training-free approach, while the second is trained with ground truth 3D bounding box annotations. In Tab. 4, we demonstrate that even with less exposure to the training data, Z3D outperforms methods that source object proposals with Mask3D. When using Mask3D, Z3D shows +2.8 Acc@0.5 on ScanRefer w.r.t. the best competing approach in the respective category.

Number of images In image-base scenarios, the number of input images is a crucial aspect of the

Method	Venue	Supervision		Unique		Multiple		Overall	
		bboxes	texts	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5
ZSVG3D	CVPR'24	Mask3D	✗	55.3	55.3	25.6	25.6	31.2	31.2
SeqVLM	ACMMM'25	Mask3D	✗	77.3	72.7	47.8	41.3	55.6	49.6
SPAZER	NIPS'25	Mask3D	✗	80.9	72.3	51.7	43.4	57.2	48.8
Z3D	-	Mask3D	✗	87.9	81.8	51.7	44.6	61.2	54.4
LLM-Grounder	ICRA'24	✗	✗	12.1	4.0	11.7	5.2	12.0	4.4
VLM-Grounder	CoRL'24	✗	✗	66.0	29.8	48.3	33.5	51.6	32.8
Z3D	-	✗	✗	78.8	71.2	50.5	44.6	58.0	51.6

Table 4: Evaluations of 3DVG from point clouds on 250 scenes subset from ScanRefer.

Method	VLM	Easy	Hard	Dep.	Indep.	Overall
SeeGround	Qwen2-VL-72B	51.5	37.7	44.8	45.5	45.2
VLM-Grounder	GPT-4o	55.2	39.5	45.8	49.4	48.0
SeqVLM	Doubao-1.5-vision-pro	58.1	47.4	51.0	54.5	53.2
SPAZER	Qwen2.5-VL-72B	60.3	50.9	54.2	57.1	56.0
Z3D	Qwen2.5-VL-72B	66.9	44.7	58.3	55.8	56.8
Z3D	Qwen3-VL-8B-Thinking	59.6	47.4	54.2	53.9	54.0
Z3D	Qwen3-VL-30B-Thinking	66.2	47.4	57.3	57.8	57.6
Z3D	Qwen3-VL-235B-Thinking	68.4	47.4	58.3	59.1	58.8

Table 5: Evaluations of 3DVG from point clouds on 250 scenes subset from Nr3D.

model’s performance. According to the experiments on ScanRefer, the more images, the better (Tab. 7). Since the reconstruction quality is highly correlated with the coverage of a scene, this conclusion can be expected. Existing approaches use a comparable number of images, e.g., VLM-Grounder takes up to 60 images and VG LLM uses 24 images. Still, after the view selection procedure, all methods reason based on fewer images: 3 in Z3D, 7 in VLM-Grounder, or 6 in VG LLM.

DUST3R vs. DROID-SLAM To investigate how dependent is our pipeline on the reconstruction quality, we replace DUST3R with DROID-SLAM (Teed and Deng, 2021). According to Tab. 8, this leads to a massive drop of scores: apparently, DROID-SLAM cannot deliver the sufficient quality produce to localize and recognize 3D objects reliably.

Inference time We measure inference time component-wise and report the performance in Tab. 6. The most time-consuming part of our pipeline is MaskClustering, while other components are executed relatively fast. Overall, Z3D is on par with the zero-shot baseline VLM-Grounder.

C Qualitative Results

ScanRefer Fig. 2 depicts Z3D predictions on ScanRefer from all types of inputs: images solely, images with poses, and images with poses and

Method	Step	Time (s)	Total (s)
Z3D	MaskClustering	56.3	61.0
	CLIP view selection	0.001	
	VLM view selection	1.5	
	SAM3-Agent	2.8	
	multi-view aggregation	0.4	
SPAZER	view selection	5.2	23.5
	candidate object screening	8.5	
	3D-2D decision-making	9.8	
VLM-Grounder	-	-	50.3

Table 6: Inference time of each step in Z3D.

# Images	Images		Images + camera poses	
	Acc@0.25	Acc@0.5	Acc@0.25	Acc@0.5
15	20.3	8.5	32.6	17.1
45	30.0	12.7	41.1	24.0

Table 7: Results of Z3D from images with and without poses on ScanRefer 250-scenes subset with varying number of images.

Method	Acc@0.25	Acc@0.5
DROID-SLAM → Z3D	14.3	4.8
DUST3R → Z3D	30.0	12.7

Table 8: Results of Z3D from images on ScanRefer 250-scenes subset with different pose estimation methods.




Figure 2: Qualitative results of Z3D on ScanRefer.

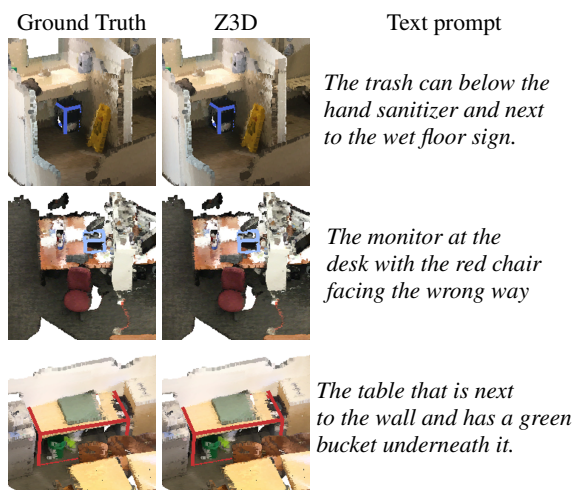


Figure 3: Qualitative results of Z3D on Nr3D.

depths. Comparison on the same scene shows how additional inputs contribute to the quality.

Nr3D Predictions on Nr3D given images with poses and depths are shown in Fig. 3. Nr3D benchmark provides ground truth 3D bounding boxes, from which the only one should be selected as an answer; accordingly, predicted boxes is strictly equal to ground truth ones if the guess is correct.

Failure cases We analyzed failure cases and identified a typical pattern. As can be observed in Fig. 4, our model sometimes fails to select the object in the presence of multiple similar objects in a scene: monitors (top row), pillows (bottom row).

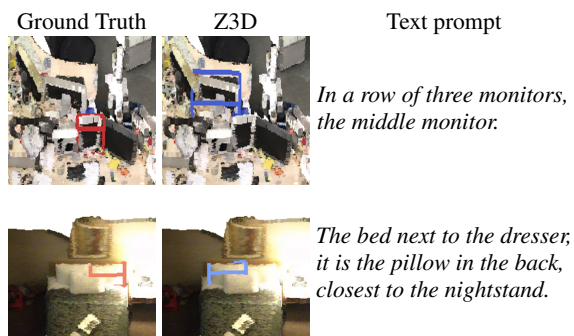


Figure 4: Failure cases of Z3D on Nr3D.