

EmbSpatial-Bench: Benchmarking Spatial Understanding for Embodied Tasks with Large Vision-Language Models

Mengfei Du^{1*}, Binhao Wu^{1*}, Zejun Li¹, Xuanjing Huang², Zhongyu Wei^{1†}

¹School of Data Science, Fudan University, China

²School of Computer Science, Fudan University, China

{mfdu22, bhwu22}@m.fudan.edu.cn

{zejunli20, xjhuang, zywei}@fudan.edu.cn

Abstract

The recent rapid development of Large Vision-Language Models (LVLMs) has indicated their potential for embodied tasks. However, the critical skill of spatial understanding in embodied environments has not been thoroughly evaluated, leaving the gap between current LVLMs and qualified embodied intelligence unknown. Therefore, we construct EmbSpatial-Bench, a benchmark for evaluating embodied spatial understanding of LVLMs. The benchmark is automatically derived from embodied scenes and covers 6 spatial relationships from an egocentric perspective. Experiments expose the insufficient capacity of current LVLMs (even GPT-4V). We further present EmbSpatial-SFT, an instruction-tuning dataset designed to improve LVLMs' embodied spatial understanding.

1 Introduction

Embodied AI is the frontier direction of general-purpose AI systems, requiring intelligent agents to understand instructions, perceive physical environments, plan and execute actions to accomplish corresponding tasks (Anderson et al., 2018). Recently, LLM-based large vision-language models (LVLMs) have demonstrated powerful capabilities in following instructions and performing planning based on the visual contexts (Li et al., 2023b; Zhu et al., 2023; OpenAI, 2023), paving a promising path for the development of embodied AI systems.

However, recent studies have revealed significant deficiencies of LVLMs in understanding visual contents (Li et al., 2023c). In terms of embodied scenarios, the ability to understand spatial relationships between objects is particularly vital for agents to effectively interact with the environment (Anderson et al., 2018; Padmakumar et al., 2022). Evaluating and enhancing such capabilities of LVLMs is essential for constructing LVLM-driven embodied

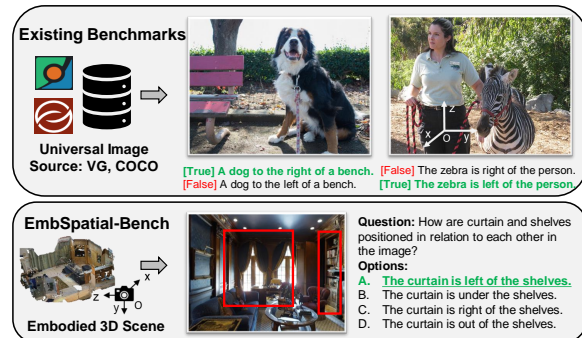


Figure 1: Comparison between EmbSpatial-Bench and existing benchmarks for spatial understanding. Existing benchmarks may determine spatial relationships based on a coordinate system centered on the subject in the image (upper right), whereas EmbSpatial-Bench consistently determines them from an egocentric perspective.

agents. Yet, existing benchmarks are not suitable for accurately assessing such capabilities.

In this paper, we argue that two important features should be considered for excellent evaluation of spatial understanding abilities in embodied tasks. First, the spatial relationships should be described from the egocentric perspective, for the reason that agents take themselves as the center of coordinates to follow instructions and infer decisions in embodied tasks. However, previous benchmarks for spatial understanding (Liu et al., 2023a) tend to depict spatial relationships from the perspective of subject within images, as illustrated in Figure 1. Second, the visual scenes for evaluation should be consistent with that in embodied tasks. Nevertheless, existing benchmarks (Liu et al., 2023a; Kamath et al., 2023) are mainly constructed from universal image-text datasets like MSCOCO (Lin et al., 2014) and VG (Krishna et al., 2017) which are weakly related to embodied scenarios.

To meet aforementioned requirements, we establish EmbSpatial-Bench, a benchmark for evaluating spatial understanding abilities of LVLMs in embodied environments. As shown in Figure 1, we focus on six spatial relationships described from the ego-

*Equal contribution

†Corresponding author

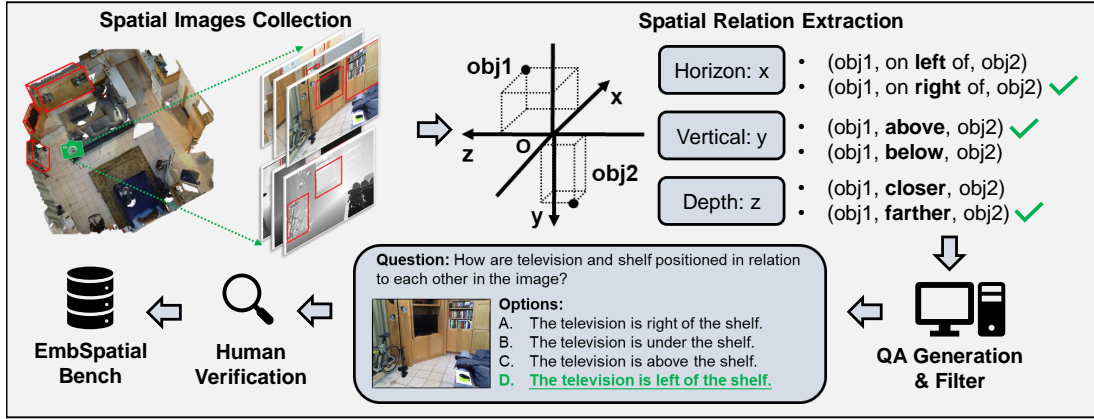


Figure 2: Overview of the construction pipeline for EmbSpatial-Bench based on existing annotated 3D environments.

centric perspective, including *above*, *below*, *left*, *right*, *close* and *far*, which completely covers three dimensions of the coordinates. The benchmark is organized into the format of multiple-choice questions. The images used for evaluation are directly collected from embodied 3D scenes, namely MP3D (Chang et al., 2017), AI2-THOR (Kolve et al., 2017) and ScanNet (Dai et al., 2017).

Based on EmbSpatial-Bench, various LVLMs have been assessed. Experimental results indicate the poor embodied spatial understanding of current LVLMs, including GPT-4V (OpenAI, 2023) and Qwen-VL-Max (Bai et al., 2023). To address the issue, we further construct an instruction-tuning dataset, EmbSpatial-SFT, to empower LVLMs with embodied spatial understanding ability. LVLMs fine-tuned on EmbSpatial-SFT consistently demonstrate improved spatial perception abilities across different scenarios.¹

2 EmbSpatial-Bench

Unlike existing benchmarks built on 2D images (Liu et al., 2023a), EmbSpatial-Bench is constructed from 3D scenes. Figure 2 illustrates the construction pipeline. We first generate target images from 3D scenes and extract spatial relations among objects. Then, we generate QA pairs and conduct filtering. Section 2.1 provides detailed explanations of each part, while Section 2.2 offers statistics of the benchmark.

2.1 Dataset Construction

Spatial Image Sources. Current embodied 3D simulators offer comprehensive annotations for tasks such as visual navigation (Chang et al., 2017) and room rearrangement (Weihs et al., 2021), mak-

ing them ideal for constructing a challenging benchmark to evaluate embodied spatial understanding. Therefore, we choose MP3D (Chang et al., 2017), ScanNet (Dai et al., 2017) and AI2-THOR (Kolve et al., 2017). Specifically, we utilize the test scenes from MP3D and validation scenes from ScanNet and A. Within each 3D scene, we randomly select viewpoints and capture the corresponding RGB-D images accordingly. In AI2-THOR, we select 7 types of household tasks from ALFRED (Shridhar et al., 2020), spanning 93 different scenes. During task execution, we identify key RGB-D images based on the dataset’s PDDL (Aeronautiques et al., 1998) annotations. (See Appendix A).

Spatial Relation Extraction. Instead of relying on object detectors (Tejas et al., 2023), we extract spatial relations directly from well-annotated 3D datasets. For each object in each image, we can utilize the camera parameters along with the corresponding 3D coordinates to obtain its 2D coordinates in the image (in the form of bounding boxes). With the 2D annotations, we extract the spatial relation triples with non-overlapping bounding boxes. We consider six spatial relationships from the viewer’s perspective: *above*, *below*, *left*, *right*, *close* and *far*. For the first four types, we determine the spatial relation based on position of the entire bounding boxes. For instance, if the entire bounding box of object A is located to the left of object B, we consider the relationship between A and B as *A is left of B*. For the other two types, we use the average depth within the bounding box to determine which object is farther or closer.

QA Generation. The format of our benchmark is multiple-choice questions, a widely adopted approach in various LVLm benchmarks (Liu et al., 2023c; Li et al., 2023d). For the relations *above*,

¹<https://github.com/mengfeidu/EmbSpatial-Bench>

Model	Generation	Likelihood
BLIP2 (2023b)	37.99	35.71
InstructBLIP (2023)	38.85	33.41
Cheetor (2023a)	24.56	32.80
Lynx (2023)	29.09	41.62
mPlugOwl (2023)	24.12	27.42
ImagebindLLM (2023)	26.46	33.46
Shikra (2023b)	28.38	34.75
MiniGPT4 (2023)	23.54	31.70
MiniGPT-v2 (2023a)	23.93	43.85
LLaVA-1.6 (2023b)	35.19	38.84
GPT-4V (2023)	36.07	-
Qwen-VL-Max (Bai et al., 2023)	49.11	-
Human	90.33	-

Table 2: Zero-shot performance (Acc%) of LVLMs in EmbSpatial-Bench. **Bold** indicates the best results.

4 Experiments

4.1 Experimental Setup

Based on EmbSpatial-Bench, we conduct zero-shot evaluation of current LVLMs, using accuracy as the metric. Two evaluation strategies are employed. The first one is the generation-based strategy, which directly uses predicted options from the textual outputs of models. Considering the insufficient instruction-following ability of some LVLMs, we also employed a likelihood strategy, using the option with the highest probability generated by the model (Li et al., 2023d). Please refer to Appendix B for more evaluation details.

4.2 Zero-shot Performance

Table 2 presents the zero-shot performance of 10 open-source LVLMs and 2 closed-source models. The results indicate that current LVLMs, including powerful closed-source models like GPT-4V and Qwen-VL-Max, have not demonstrated satisfactory spatial understanding abilities in embodied scenes. The best performance among all LVLMs merely reaches an accuracy of 49.11% (Generation) or 43.85% (Likelihood) which is significantly lower than human performance (90.33%). We present failure cases of GPT-4V in Appendix C, revealing its poor abilities of both object localization and spatial relation identification. The versions of these models can be found in Appendix B.3.

4.3 Instruction Tuning on EmbSpatial-SFT

Furthermore, we fine-tune MiniGPT-v2 on EmbSpatial-SFT, to explore whether the data could further enhance the model’s spatial understanding capabilities. The trainable parameters include the visual connection module and LoRA (Hu et al., 2021) modules in the LLM backbone.

Model	In-Domain		Out-Domain		All
	MP3D	AI2-THOR	ScanNet		
Generation					
MiniGPT-v2 (2023a)	23.31	20.58	28.00	23.93	
Finetuned MiniGPT-v2	31.64	34.06	33.17	32.97	
w/o LoRA	26.81	25.26	23.25	25.11	
w/o OL	34.22	31.40	31.92	32.50	
Likelihood					
MiniGPT-v2 (2023a)	46.71	41.97	42.92	43.85	
Finetuned MiniGPT-v2	80.52	73.69	80.25	78.10	
w/o LoRA	48.38	38.90	44.17	43.76	
w/o OL	80.35	72.15	79.67	77.34	

Table 3: Performance (Acc%) of MiniGPT-v2 tuned on EmbSpatial-SFT. OL stands for object localization while w/o LoRA indicates that only the connection module is fine-tuned. **Bold** indicates the best results.

Main Results. According to Table 3, under the likelihood evaluation strategy, learning from EmbSpatial-SFT consistently improves the performance across both in-domain and out-domain environments, with an increase of 34.25% in the overall accuracy. Though not as significant as that under likelihood strategy, the evaluated results under generation strategy still demonstrate an adequate performance improvement (+9.04% overall) after instruction-tuning. The improvement in AI2-THOR is less than in ScanNet, which we attribute to AI2-THOR primarily consisting of simulated scenes, unlike the real-world scenarios in MP3D and ScanNet.

Ablations. We further validate the effectiveness of finetuning LLM backbone with LoRA and the auxiliary object localization data. As shown in Table 3, tuning the LLM backbone with LoRA significantly contributes to the performance across all scenarios compared to the variant with a frozen LLM backbone. This phenomenon implies the necessity for the LLM backbone to learn corresponding reasoning abilities for spatial understanding, rather than solely adjusting the input visual representations. The auxiliary data also contribute to the performance across different embodied environments, leading to an overall improvement of 0.47% and 0.76% under generation strategy and likelihood strategy, respectively.

5 Related Works

Large Vision-Language Models The prevalent LVLMs (Dai et al., 2023; Zeng et al., 2023) learn visual representations from abundant image-text interleaved datasets with a lightweight connection module. Further works (Tsai et al., 2023; Zheng et al., 2023) fine-tunes LVLMs-based architecture

and obtain acceptable performance on embodied tasks, which preliminarily reveal the potential of LVLMs as embodied intelligence. However, these works neither evaluate nor empower LVLMs with spatial understanding ability, which is essential for various embodied tasks.

Benchmarks for Spatial Understanding. While there are numerous universal benchmarks available for LVLMs (Xu et al., 2023; Fu et al., 2023; Li et al., 2023d), dedicated benchmarks for evaluating spatial understanding remain scarce. VSR (Liu et al., 2023a) typically examines spatial relationships from the perspective of the subject within the image. What’sUp (Kamath et al., 2023) addresses data bias and generates uncluttered images to eliminate interference from unrelated objects. SR_{2D} (Tejas et al., 2023) focuses on evaluating text-to-image generative model. However, all of them are built on COCO (Veit et al., 2016) or VG (Krishna et al., 2017) which are not consistent with the embodied scenarios. This lack of specialized benchmarks leaves the spatial understanding capabilities of LVLMs in embodied tasks unexplored.

6 Conclusion

In this work, we propose EmbSpatial-Bench, a benchmark to evaluate embodied spatial understanding of LVLMs. The evaluation results reveal the weak spatial understanding ability of current popular LVLMs. We further propose EmbSpatial-SFT, an instruction tuning dataset to enhance the capacity of LVLMs. Extensive experiments valid the effectiveness of each data component in our EmbSpatial-SFT, with the goal of empowering the spatial understanding ability of LVLMs.

Limitations

Spatial understanding in embodied environments is a crucial aspect of LVLMs’ capabilities for embodied tasks. In this study, we advance towards this goal by constructing benchmark and instruction-tuning datasets from well-annotated 3D embodied datasets. These datasets are derived from three widely used indoor embodied datasets, which may restrict their suitability for outdoor environments. Additionally, our study only investigates the English language, thus limiting the generalizability of the benchmark and findings to other languages.

Ethical Considerations

The benchmark and instruction-tuning data are built from publicly available embodied datasets, which include either photorealistic scenes or generated rendered scenes without any copyright issues. Besides, our data source does not contain any personal data, uniquely identifiable individuals, or offensive content.

Acknowledgements

This work is supported by National Natural Science Foundation of China (No. 62176058) and National Key R&D Program of China (2023YFF1204800). The project’s computational resources are supported by CFFF platform of Fudan University.

References

- Constructions Aeronautiques, Adele Howe, Craig Knoblock, ISI Drew McDermott, Ashwin Ram, Manuela Veloso, Daniel Weld, David Wilkins Sri, Anthony Barrett, Dave Christianson, et al. 1998. Pddl the planning domain definition language. *Technical Report, Tech. Rep.*
- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3674–3683.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv:2308.12966*.
- Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. 2017. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*.
- Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. 2023a. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. 2023b. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan

- Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality.](#)
- Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. [Instructblip: Towards general-purpose vision-language models with instruction tuning.](#)
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, et al. 2023. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.
- Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. 2023. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*.
- Jiaming Han, Renrui Zhang, Wenqi Shao, Peng Gao, Peng Xu, Han Xiao, Kaipeng Zhang, Chris Liu, Song Wen, Ziyu Guo, et al. 2023. Imagebind-llm: Multi-modality instruction tuning. *arXiv preprint arXiv:2309.03905*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Amita Kamath, Jack Hessel, and Kai-Wei Chang. 2023. What's "up" with vision-language models? investigating their struggle with spatial reasoning. *arXiv preprint arXiv:2310.19785*.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798.
- Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. 2017. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73.
- Jun Cheng Li, Kaihang Pan, Zhiqi Ge, Minghe Gao, Hanwang Zhang, Wei Ji, Wenqiao Zhang, Tat-Seng Chua, Siliang Tang, and Yueting Zhuang. 2023a. Fine-tuning multimodal llms to follow zero-shot demonstrative instructions. *arXiv preprint arXiv:2308.04152*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023b. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023c. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Zejun Li, Ye Wang, Mengfei Du, Qingwen Liu, Binhao Wu, Jiwen Zhang, Chengxing Zhou, Zhihao Fan, Jie Fu, Jingjing Chen, et al. 2023d. Reform-eval: Evaluating large vision language models via unified re-formulation of task-oriented benchmarks. *arXiv preprint arXiv:2310.02569*.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Fangyu Liu, Guy Emerson, and Nigel Collier. 2023a. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 11:635–651.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023b. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023c. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Aishwarya Padmakumar, Jesse Thomason, Ayush Srivastava, Patrick Lange, Anjali Narayan-Chen, Spanana Gella, Robinson Piramuthu, Gokhan Tur, and Dilek Hakkani-Tur. 2022. Teach: Task-driven embodied agents that chat. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2017–2025.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10740–10749.

- Tejas, Hamid Gokhale, Besmira Palangi, Vibhav Nushi, Eric Vineet, Ece Horvitz, Chitta Kamar, Yezhou Baral, and Yang. 2023. Benchmarking spatial relationships in text-to-image generation. *arXiv preprint arXiv:2212.10015*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Yao-Hung Hubert Tsai, Vansh Dhar, Jialu Li, Bowen Zhang, and Jian Zhang. 2023. Multimodal large language model for visual navigation. *arXiv preprint arXiv:2310.08669*.
- Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. 2016. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*.
- Luca Weihs, Matt Deitke, Aniruddha Kembhavi, and Roozbeh Mottaghi. 2021. Visual room rearrangement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao, Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang, Yu Qiao, and Ping Luo. 2023. Lvlm-ehub: A comprehensive evaluation benchmark for large vision-language models. *arXiv preprint arXiv:2306.09265*.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. 2023. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*.
- Yan Zeng, Hanbo Zhang, Jiani Zheng, Jiangnan Xia, Guoqiang Wei, Yang Wei, Yuchen Zhang, and Tao Kong. 2023. What matters in training a gpt4-style language model with multimodal inputs? *arXiv preprint arXiv:2307.02469*.
- Duo Zheng, Shijia Huang, Lin Zhao, Yiwu Zhong, and Liwei Wang. 2023. Towards learning a generalist model for embodied navigation. *arXiv preprint arXiv:2312.02010*.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.

Appendix A Dataset Details

A.1 AI2-THOR Image Selection

Due to the significant similarity between many images in the observation sequences for each task in AI2-THOR, filtering is necessary. Based on the detailed PDDL annotations from ALFRED (Shridhar et al., 2020), we select key images that show significant content changes after each sub-goal is reached as our benchmark image resources.

A.2 Dataset Statistics

The wordcloud of object categories can be observed in Figure 5. The distribution of questions for each spatial relation is illustrated in Figure 6. The diversity and balance of the data enhance to the reliability of our benchmark.

A.3 Data Cases

Three samples of EmbSpatial-Bench constructed from MP3D (Chang et al., 2017), AI2-THOR (Kolve et al., 2017) and ScanNet (Dai et al., 2017) are shown in Fig. 7, Fig. 8 and Fig. 9.

A.4 Filtering and Verification

Initially, we will implement two primary filtering processes to enhance the robustness and quality of our benchmark. First, we filter out objects with excessively large or small bounding boxes. To exclude improperly displayed objects, we filter out spatial relationship triplets where the length or width of the bounding box is less than 50 or greater than half the length of the corresponding dimension of the image.

After automated construction and filtering processes, the human verification is implemented to further ensure the correctness of our benchmark. Specifically, the correctness of each sample is examined by human from several aspects: 1) the objects involved in the question can be identified in the image uniquely and clearly; 2) the target object conforms to the described spatial relationship; 3) the negative options are indeed incorrect objects or relationships. Any sample that does not meet either of these conditions is discarded.

Appendix B Experiments


B.1 Experimental Details

Implementation details. We use MiniGPT-v2 (Chen et al., 2023a) as a baseline LLM for investigation. The architecture of MiniGPT-v2 comprises three components, including a vision encoder

Model	In-Domain						Out-Domain											
	above	below	MP3D		close	far	above	below	AI2THOR		close	far	above	below	ScanNet		close	far
			left	right					left	right					left	right		
Generation																		
MiniGPT-v2	31.22	26.90	24.76	20.48	21.29	15.54	25.60	23.41	18.93	16.19	25.24	13.93	31.00	33.00	30.00	22.50	29.50	22.00
Finetuned MiniGPT-v2	38.62	46.19	22.86	23.33	34.65	25.39	36.23	49.76	28.16	26.67	34.76	28.86	39.00	48.00	27.50	23.00	28.50	33.00
Likelihood																		
MiniGPT-v2	91.01	76.65	30.95	30.48	25.74	29.53	79.71	62.93	30.58	25.24	32.38	20.9	78.50	73.50	28.00	32.50	27.00	18.00
Finetuned MiniGPT-v2	92.59	91.88	84.29	82.38	71.78	60.10	93.72	88.78	83.50	80.95	50.00	44.77	90.50	89.00	89.50	90.50	56.50	65.50

Table 4: Performance (Acc%) of MiniGPT-v2 and fine-tuned MiniGPT-v2 across different spatial relations.

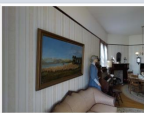
Question: How are curtain and shelves positioned in relation to each other in the image?



Options:

- A. **The curtain is left of the shelves.**
- B. The curtain is under the shelves.
- C. The curtain is right of the shelves.
- D. The curtain is out of the shelves.

Question: From your perspective, which object in the image is at the shortest distance?




Options:

- A. table.
- B. chair.
- C. **sculpture.**
- D. fireplace.

Figure 7: Data samples from Matterport3D.

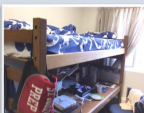
Question: How are television and shelf positioned in relation to each other in the image?



Options:

- A. The television is right of the shelf.
- B. The television is under the shelf.
- C. The television is above the shelf.
- D. **The television is left of the shelf.**

Question: Which object from the list is situated at the shortest distance from your point of view within the image?




Options:

- A. window.
- B. **lamp.**
- C. curtain.
- D. clothes.

Figure 9: Data samples from ScanNet.


Question: What is the spatial arrangement of pan and pepper shaker in the image concerning each other?



Options:

- A. **The pan is left of the pepper shaker.**
- B. The pan is blocking the pepper shaker.
- C. The pan is inside the pepper shaker.
- D. The pan is right of the pepper shaker.

Question: Which object from the list is situated at the largest distance from your point of view within the image?




Options:

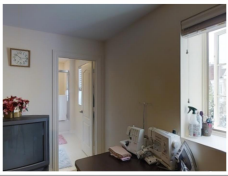
- A. potato.
- B. bowl.
- C. pot.
- D. **bin.**

Figure 8: Data samples from AI2-THOR.

both two processes.

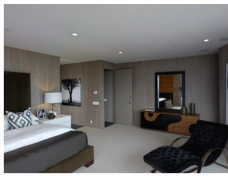
whether object localization or spatial relationships determination becomes a bottleneck. As shown in Figure 10, the GPT-4V not only makes mistakes in object positioning, but also misjudge their spatial relationship when successfully localizing the objects involved. In the first case (left part), GPT-4V mistakenly positions the clock in top left corner to the top right corner, further leading to the incorrect selection of option with the word "right". In the second case (right part), GPT-4V successfully locates the positions of all object referred in the question, but incorrectly choose the pillow as the nearest object rather than the bed. The case study demonstrate the potential room for improvement in

 **System:** Assume you are a viewer seeing current observation. You are supposed to understand the spatial relationships among several objects. The spatial relationships should be described in the viewer's perspective. You need to select the option to answer the question below:




Question: In the image, how do the positions of television and clock interact with each other?
Options:
 (A) The television is under the clock.
 (B) The television is outside the clock.
 (C) The television is blocking the clock.
 (D) The television is at the right side of the clock.

1. Please first describe the position of **television and clock** respectively in the image.
 2. Please choose the option to answer the question above with the objects' positions determined in the previous step.



Question: Which object, in relation to your current position, holds the nearest placement in the image?
Options:
 (A) Pillow. (B) Headboard.
 (C) Door. (D) Bed.

1. Please first describe the position of **pillow, headboard, door and bed** respectively in the image.
 2. Please choose the option to answer the question above with the objects' positions determined in the previous step.



1. In the image, the television is situated near the bottom left corner. **The clock is visible on the wall closer to the top right corner, above and to the right of the television.**

2. Based on the positions determined in the previous step, the correct option is: **(D) The television is at the right side of the clock.**

1. In the image, **the pillow is positioned on the bed, closest to the bottom left of the frame.** The headboard is directly behind the pillow, attached to the bed. The door is located across the room, near the center of the image in the background. **The bed is situated in the foreground of the image, occupying most of the bottom half of the visual field.**

2. To answer the question, **the pillow (A) is the nearest object** in relation to the current position, as it is the closest to the camera's perspective at the bottom left of the image.

Figure 10: Case study of GPT-4V on our benchmark. It not only makes mistakes in object positioning, but also misjudge their spatial relationship when successfully localizing the objects involved. **The text in red** means the wrong answers generated by GPT-4V.