

Enhancing Spatial Reasoning Through Visual and Textual Thinking

Xun Liang^{1,2}, Xin Guo², Zhongming Jin², Weihang Pan³, Penghui Shang⁴, Deng Cai¹, Binbin Lin^{3,4*}, Jieping Ye²

¹State Key Lab of CAD&CG, Zhejiang University

²Alibaba Cloud Computing

³School of Software Technology, Zhejiang University

⁴Xihu, Hangzhou Zhiyuan Research Institute Co., Ltd

Abstract

The spatial reasoning task aims to reason about the spatial relationships in 2D and 3D space, which is a fundamental capability for Visual Question Answering (VQA) and robotics. Although vision language models (VLMs) have developed rapidly in recent years, they are still struggling with the spatial reasoning task. In this paper, we introduce a method that can enhance Spatial reasoning through Visual and Textual thinking Simultaneously (SpatialVTS). In the spatial visual thinking phase, our model is trained to generate location-related specific tokens of essential targets automatically. Not only are the objects mentioned in the problem addressed, but also the potential objects related to the reasoning are considered. During the spatial textual thinking phase, our model conducts long-term thinking based on visual cues and dialogues, gradually inferring the answers to spatial reasoning problems. To effectively support the model's training, we perform manual corrections to the existing spatial reasoning dataset, eliminating numerous incorrect labels resulting from automatic annotation, restructuring the data input format to enhance generalization ability, and developing thinking processes with logical reasoning details. Without introducing additional information (such as masks or depth), our model's overall average level in several spatial understanding tasks has significantly improved compared with other models.

Code — <https://github.com/whalelx/SpatialVTS.git>

1 Introduction

Spatial reasoning aims to understand spatial arrangements in 2D and 3D spaces, which is crucial for accurately interpreting complex visual environments (Cheng et al. 2024). The significance of enhancing the spatial reasoning of the model lies not only in general visual understanding but also has important practical significance for other downstream tasks such as embodied intelligent decision-making, human-computer interaction, and other active fields (Driess et al. 2023; Chen et al. 2024a). Visual language models (VLMs) have developed rapidly in recent years and made a significant progress in many basic visual tasks, such as image and video understanding (Liu et al. 2023b,a; Li et al. 2024b,a; Bai et al. 2023; Wang et al. 2024; Bai et al. 2025; Chen et al.

2024b,c,d; Achiam et al. 2023; Hurst et al. 2024; Grattafiori et al. 2024; Abdin et al. 2024). However, spatial reasoning tasks, such as determining the 3D positions of objects or their spatial relationships, are still complex for most modern VLMs (Liu, Emerson, and Collier 2023).

Previous research on spatial reasoning has typically concentrated on explicit spatial scene memories (Gervet et al. 2023; Gordon et al. 2018) or spatial scene graphs (Hemachandra et al. 2014; Hildebrandt et al. 2020; Wald et al. 2020; Walter et al. 2013). In order to address spatial problems presented in the VQA format, these approaches often need to treat the task explicitly as a path-finding problem within the context of such scene graphs. With the explosive growth of visual language models' capabilities, their spatial reasoning capabilities are expected to be substantially improved. Researchers have first focused on constructing benchmarks and enhancing the spatial understanding ability of large visual models in a data-driven manner (Liu, Emerson, and Collier 2023; Chen et al. 2024a). Meanwhile, some researchers (Cheng et al. 2024) focus on integrating more visual elements (such as depth, region, etc.) to enhance the model's understanding of spatial relationships and scales.

Our key insight is that spatial reasoning is not merely simple question-and-answer for the targets in the picture. Many spatial reasoning tasks require analysis combined with potential visual cues to give the correct answer. As shown in Fig.1, if we only focus on Region [0] and Region [1], it will be difficult for us to determine their actual vertical distance. In fact, the building where Region [0] is located (indicated by the red box) is an important visual reference in this question. We can first estimate that Region [0] is approximately at the height of a four-story building and then infer the apparent distance between Region [0] and Region [1]. That is to say, visual cues are not limited to the objects mentioned in the question. There are several potential related targets in the image that seem irrelevant to the question but can assist in answering spatial reasoning tasks. In the inference process of VLM, attention should be paid to these potential targets, thereby assisting in the answer to the final question. On the other hand, after obtaining more visual cues, the model should conduct further reasoning on the question and visual cues rather than simply providing a straightforward answer. A large number of visual reasoning works (Xu et al. 2024; Liu et al. 2025; Li et al. 2025; Thawakar et al. 2025) have

*Corresponding author



Could you measure the vertical distance between Region [0] and Region [1]?

Spatial Visual Thinking

Region [0]
[<x_1><y_0><x_2><y_1>]

Region [1]
[<x_3><y_6><x_4><y_6>]

Building
[<x_0><y_2><x_1><y_7>]

Spatial Text Thinking

The building is approximately four stories tall, with **each story being about 9.84 feet high**, which can be used to determine the vertical distance between Region [0] (spire) and Region [1] (car).

A vertical distance of **39.36 feet** exists between Region [0] and Region [1].

Figure 1: Our model analyzes the original problem and the image to find both the evident and potential visual cues related to the problem in the image and provides corresponding rationales. Based on the visual cues and rationales, our model can effectively enhance the ability of spatial reasoning.

proved that allowing the model to think further can enhance the model’s understanding of the task. In the spatial reasoning task, we believe that stimulating the model’s long-term thinking ability can help further explore and utilize visual cues. For the example in Fig.1, textual reasoning can further establish the connections between various visual targets. After further reflection on visual cues, the model finally determines the vertical height between Region [0] and Region [1].

Inspired by the recent CoT-based(Shao et al. 2024; Xu et al. 2025) and reasoning-based VLMs(Thawakar et al. 2025), we propose a method that can enhance Spatial reasoning through Visual and Textual thinking Simultaneously (SpatialVTS). Our model can be summarized as two phases: Spatial Visual Thinking and Spatial Textual Thinking. In the first phase, the VLM is trained to generate the regional location tokens related to important targets. In the second phase, the VLM takes the question and visual cues into consideration. By stimulating the model’s reasoning ability, SpatialVTS can further establish the connections between various visual targets and ultimately determine the final answer. To achieve this goal, we construct a dataset consisting of questions, answers, relevant target boxes, and rationales. The relevant target boxes are represented using special tokens of discrete positions. The rationales tell us what is in the target box and why it relates to spatial reasoning. Experiments demonstrate that our model can improve spatial reasoning capabilities in average performance across various datasets, even without any assistance of additional visual elements, such as depth estimation maps. Our contributions can be summarized as follows.

- We propose a method that simultaneously enhances spatial reasoning through visual and textual thinking. Spatial visual thinking effectively explores spatial visual cues and outputs the candidate reference targets. Spatial textual thinking conducts further thinking based on visual cues, rationalizes, and answers the question step by step.

- To effectively support model training, we analyze the deficiencies of existing spatial reasoning datasets. We clean and correct the data with relevant target special tokens and rationales through manual annotation.
- The results of the experiment demonstrate the effectiveness of our approach. Without incorporating additional information, such as masks or depth data, our model demonstrates a significant improvement in the average performance across various spatial understanding tasks.

2 Related Work

Spatial Reasoning Benchmarks Early works on spatial reasoning of large models focus on building the spatial reasoning benchmarks and exploring the performance of the VLMs in spatial reasoning tasks. Visual Spatial Reasoning (VSR)(Liu, Emerson, and Collier 2023) presents a dataset containing more than 10k natural text-image pairs with 66 types of spatial relations (such as under, in front of, facing). It performs a by-relation analysis and finds that the models’ performances on certain relations have little correlation with the number of training examples, and certain relations are inherently more challenging. SpatialVLM(Chen et al. 2024a) introduces an automated 3D spatial VQA data generation framework capable of scaling up to 2 billion VQA examples across 10 million real-world images. The authors explore various factors in the training process, including data quality, training pipeline design, and VLM architecture. Training on these datasets can significantly improve models’ performance in both qualitative and quantitative aspects of spatial VQA. SpatialRGPT(Cheng et al. 2024) introduces a scalable data pipeline that constructs region-aware spatial reasoning QAs from existing datasets. It presents SpatialRGPT-Bench, a comprehensive benchmark based on 3D annotations that span indoor, outdoor, and simulated environments.

Visual Elements Enhancement Many researchers consider incorporating more elements into models to enhance

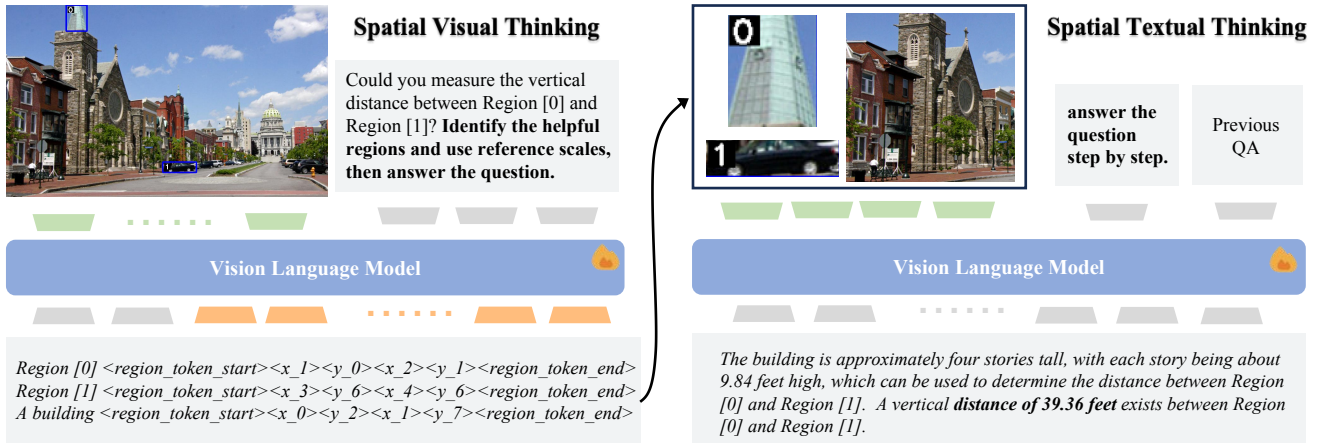


Figure 2: The framework of SpatialVTS. Our model contains two phases: the Spatial Visual Thinking phase generates the region’s special token, and the Spatial Textual Thinking phase generates the rationales and the final answer.

the comprehension ability of the vision language model. Visual CoT(Shao et al. 2024) proposes a novel multi-turn processing pipeline for MLLMs that can dynamically focus on visual inputs and provide interpretable thoughts. Their work is oriented towards interpretability analysis and does not explicitly enhance spatial vision information. The recent work(Yu, Ma, and Wang 2025) proposes to use region selection tokens to select the region related to the query and Vision Re-Encoding Tokens to re-encode the original image using re-encode the original image. However, it does not focus on potential reasoning clues and spatial tasks. RegionGPT(Guo et al. 2024) proposes a general framework that harnesses the capabilities of LLMs to tackle complex region-level captioning and understanding tasks. Osprey(Yuan et al. 2024) proposes a mask-text instruction tuning approach to extend MLLMs by incorporating fine-grained mask regions into language instruction, aiming at achieving pixel-wise visual understanding. For spatial reasoning, SpatialRGPT(Cheng et al. 2024) presents a framework that enhances region-level spatial reasoning by enabling the representation of regional information and the acquisition of spatial knowledge. They also integrate depth information, significantly improving 3D perception ability.

VLM Reasoning The trend of large language models reasoning (Jaech et al. 2024; Guo et al. 2025) has quickly spread to the field of visual language models(Team 2024). LLaVA-CoT(Xu et al. 2024) independently engages in sequential stages of summarization, visual interpretation, logical reasoning, and conclusion generation. This structured approach enables LLaVA-CoT to achieve marked improvements in precision on reasoning-intensive tasks. Visual-RFT(Liu et al. 2025) introduces a Visual Reinforcement Fine-tuning framework, which adapts the GRPO-based reinforcement learning strategy to improve the visual perception and grounding ability of LVLMs. MVoT(Li et al. 2025) leverages multimodal-native architectures to transcend text-form thinking into multimodal-native reasoning by generating image visualizations of their reasoning traces. This

reasoning paradigm enables the model to create reasoning traces and ‘think’ in words and images seamlessly while avoiding the potential errors introduced in captioning the images. LlamaV-o1(Thawakar et al. 2025) proposes a comprehensive approach for advancing multimodal reasoning by introducing a new benchmark, a novel metric, and an innovative model trained using curriculum learning.

Our model can be regarded as a region-based and reasoning-based method, which uses Spatial Visual Thinking and Spatial Textual Thinking to explicitly increase the model’s understanding of spatial reasoning.

3 Methodology

Our method can be divided into two stages. The first stage is Spatial Visual Thinking. It analyzes the original problem and the image to find both the evident and potential target positions related to the problem in the image. The second stage is Spatial Textual Thinking. In this stage, the images of the relevant regions are re-input into the VLM as additional visual cues. The model is encouraged to conduct long reasoning based on visual cues and provide the rationales and the final answer.

3.1 Spatial Visual Thinking

Spatial visual thinking aims to identify and extract potential targets related to a given spatial reasoning problem. Unlike most visual VQA questions, where answers can be easily given based on the images, the answers to spatial reasoning tasks (such as the scale of the target, the relative position of the target, etc.) cannot be obtained directly.

To prevent VLM from struggling to interpret and generate precise bounding boxes, following the work of VPT (Yu, Ma, and Wang 2025), we use the region’s approximate location to represent the target’s position in the image. As in Fig.3, we divide the $h \times w$ image evenly into a grid of $k \times k$ rectangular cells, with each cell sized $h/k \times w/k$. Each cell can be indexed by its row and column, with the top-left cell indexed as $(0, 0)$ and the top-right cell as $(k - 1, 0)$. We use

the indices of the cells containing the top-left and bottom-right pixels of the region R to describe its location. In our implementation, we set $k = 8$.

The most significant difference between our method and VPT is that we not only focus on the goals directly mentioned in the problem but also on the potential goals related to spatial reasoning tasks. Secondly, the candidate regions of our method are encoded as visual cues by the same visual encoder and serve as one of the reasoning bases for the next stage. In contrast, the candidate regions in VPT will undergo feature enhancement through other visual feature extractors (such as DINOv2) to improve the fine-grained understanding ability of the model. The construction details of the training data will be introduced in detail in Sec.3.3.




Figure 3: The region’s bounding box is represented with the indices of the cells containing the top-left and bottom-right pixels. The car in the blue bounding box can be represented as $\langle x_0 \rangle \langle y_3 \rangle \langle x_3 \rangle \langle y_6 \rangle$.

3.2 Spatial Textual Thinking

For spatial reasoning tasks, logical reasoning is critical. Firstly, the questions of most spatial reasoning tasks are difficult to answer directly. They require analysis in combination with multiple targets in the picture and the overall environment. Secondly, the answers to many spatial reasoning tasks are simple responses such as “up, down, left, right, front, or back” and “yes or no”. Using overly simplistic answers without the reasoning process in training data will easily cause pattern collapse during the model training process and lead to hallucinations in the final model.

In order to stimulate the thinking ability of the model, we further take the visual cues obtained from the visual encoding of the candidate targets in the previous stage as input. The candidate objects are cropped as new images and inputted into the same vision encoder. Then, we elaborately design a prompt to stimulate the reasoning ability of the model. We use the public visual language model to automatically generate the reasoning process as a part of the training data. In particular, we employ the “seeking the cause by grasping the result” strategy to ensure the rationality and correctness of the generating reasoning process. The question, visual cues, and answer are inputted into the large model simultaneously to find the correct reasoning path. During the training and inference phase, the model is trained to rationalize only based on the problem and visual cues.

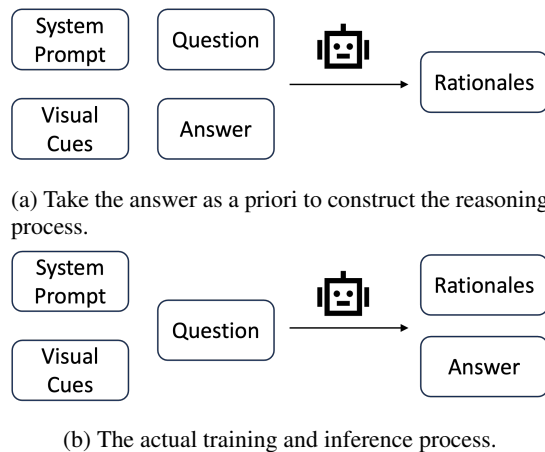


Figure 4: Construct the reasoning process via “seeking the cause by grasping the result”.

3.3 Dataset Generation

Our dataset is built based on SpatialRGPT and VPT. We find that the data needs to be reconstructed to assist the model in deeply exploring the correlations among visual targets.

1) *The input form is limited.* The SpatialRGPT model is set to require users to explicitly specify masks or bounding boxes as the target’s input to eliminate the ambiguity of target reference. However, this also limits the model’s ability for generalized spatial reasoning tasks, since it cannot work without bounding boxes. The SpatialRGPT dataset is tailor-made for the SpatialRGPT model, which means its data needs to be modified to fit the general setting.

2) *The data quality is poor.* The scale of the SpatialRGPT training set is quite large owing to the automatic data pipeline. However, the prompt template mechanism of this pipeline have brought about some data quality issues. As in Fig.5, A large proportion of the answers to the questions are incorrect. Meanwhile, since the objects and questions are randomly select from the template, a considerable number of questions are uncorrelated with the images. These problems seriously undermine the quality of model training and thus need to be manually checked and corrected.

3) *Data lacks process-oriented thinking.* The raw SpatialRGPT and VPT dataset provides answers directly based on the questions without intermediate visual or textual thinking. Therefore, based on the original dataset, we cannot obtain the potentially visual targets related to the spatial reasoning task, nor can we know the specific reasoning process.

Based on the above considerations, for the SpatialRGPT dataset, we replace the $\langle \text{mask} \rangle$ and $\langle \text{depth} \rangle$ placeholders in the user prompt using the corresponding region labels, such as Region [0] and Region [1]. The region bounding boxes are drawn on the original images and are marked with the regional tags in the corresponding questions and answers. The above corrections transformed the dataset into the most generalized input format of images and plain text for VLMs. Then, we use manual annotation for data cleaning. First, we filter out the data with unreasonable question settings and



Figure 5: The original Q&A from SpatialRGPT and our corrected Q&A.

correct the incorrect answers. Secondly, we add the bounding box annotations for the potential targets of the problem and provide relevant evidence. Finally, for both the SpatialRGPT and VPT datasets, we use the public vision language model to conduct long-term objective reasoning based on the above-mentioned goals and questions and automatically generate a long-term reasoning process. In particular, we adopt a strategy of “seeking the cause by grasping the result” to generate the thinking process. We simultaneously input the problems, visual cues, and results into the large model for the reasoning process. This can help the model output the correct reasoning process with high probability.

4 Experiments

4.1 Experiment Setup

Training Details We construct the training dataset to endow VLMs with the ability to answer the spatial reasoning questions quantitatively and qualitatively. We sample 100k samples of data from the original SpatialRGPT dataset and 340k samples of data from the VPT dataset (Yu, Ma, and Wang 2025), including 140k spatial relation questions and 200k general questions. The merged datasets are then reconstructed following Sec.3.3. We choose Qwen2-VL as our base model. We promise to publish our model and datasets after the review is completed.

Benchmarks To evaluate qualitative performance, we select the most widely adopted benchmarks, including What-

sUP (Kamath, Hessel, and Chang 2023), VSR (Liu, Emerson, and Collier 2023), BLINK-Spatial (Fu et al. 2024), and the qualitative split of the SRGPT-Bench (Battaglia et al. 2018), which is denoted as *SRGPT-QUAL*. These benchmarks contain hundreds of spatial relation concepts, such as “up/down” and “front/behind”. Meanwhile, we adopt Q-Spatial++ (Liao et al. 2024) and the qualitative split of the SRGPT-Bench, which is denoted as *SRGPT-QUAN*, to examine the perception of the model of absolute scales.

Metrics To evaluate the model accuracy on qualitative QA benchmarks, we assign a score in $[0, 1]$ to answers using DeepSeek-V3-0324, and accept answers with scores higher than 0.5. For quantitative QA benchmarks, the maximum ratio between the ground truth and the answer can be denoted as: $\delta = \max(\frac{\hat{d}}{d^*}, \frac{d^*}{\hat{d}})$, where d^* is the ground truth and \hat{d} is the estimated distance. Following Q-Spatial++, we accept the answers satisfying $\delta \leq 2$, denoted as $\delta_{\leq 2}$.

Baselines We compare SpatialVTS with generalized VLMs and VLMs specifically designed for spatial reasoning. *VPT* (Yu, Ma, and Wang 2025) is a 7B VLM with an additional CLIP encoder. *SpaceThinker* (AI 2023) is the newest open-source reproduction of SpatialVLM (Chen et al. 2024a), and achieves better performance than SpaceLLaVA (Li et al. 2024b). *SpatialRGPT* (Battaglia et al. 2018) uses refined depth and mask information to assist in training, while our model is only trained using RGB images and text. Thus, SpatialRGPT is a challenging baseline.

4.2 Main Results

Table.1 shows the performance comparison on the most commonly used *qualitative* spatial reasoning benchmarks. We can find that our SpatialVTS achieves the best or sub-optimal results in most indicators. In terms of the average effect, our model significantly outperforms other baseline models, proving our method’s superiority in spatial reasoning. Though some models can reach a high score on a specific benchmark, they can not maintain performance on all tasks. This is because spatial reasoning benchmarks are always vague and biased. A question can have multiple correct answers, even from a single perspective, and the observation frame changes a lot. Meanwhile, the answers are always simple, making it easy for VLMs to guess. Therefore, achieving high accuracy on a single benchmark does not mean it substantially understands the spatial relationship. When evaluated on another data distribution, it may make mistakes. Only if a model fully understands the spatial concept can it perform well on multiple benchmarks.

Fig.6 shows the comparison results on the *quantitative* dataset. On both the Q-Spatial++ and SRPGTBench testing datasets, we find that SpatialVTS outperforms all VLMs trained with text-image pairs, and its performance is nearly on par with SpatialRGPT, which is trained on a large amount of depth and segmentation mask pairs. Since SpatialRGPT utilizes depth information during training, it can naturally make more accurate judgments regarding precise distances. However, as the depth information is incorporated into the token list, this comes at the cost of sacrificing part of the

Models	CoT	Avg.	VSR	SpatialRGPT-QUAL	WhatsUp	BLINK-spatial
GPT-4o(Hurst et al. 2024)	✗	73.49	60.02	58.14	99.58	76.20
Qwen2-VL-7B(Wang et al. 2024)	✗	63.38	63.11	51.06	82.03	57.34
Qwen2.5-VL-7B(Bai et al. 2025)	✗	69.26	53.96	63.38	94.66	65.03
VPT(Yu, Ma, and Wang 2025)	✗	73.77	79.00	58.45	76.45	<u>81.18</u>
SpatialRGPT(Cheng et al. 2024)	✗	59.00	65.09	69.71	39.32	61.88
SpaceLLaVA(AI 2023)	✓	57.55	55.94	55.31	55.33	63.63
SpatialVTS(vision)	✗	<u>75.96</u>	74.50	<u>74.87</u>	69.17	85.31
SpatialVTS(vision/text)	✓	82.03	<u>78.21</u>	78.32	<u>95.38</u>	76.22

Table 1: Performance comparison on the qualitative spatial reasoning benchmarks. The numbers in the table represent the accuracy. The best performance is highlighted in bold, and the second performance is underlined. Due to the inherent complexity and ambiguity of the spatial reasoning task, all the models exhibited performance fluctuations across benchmarks. SpatialVTS reaches the best average result.

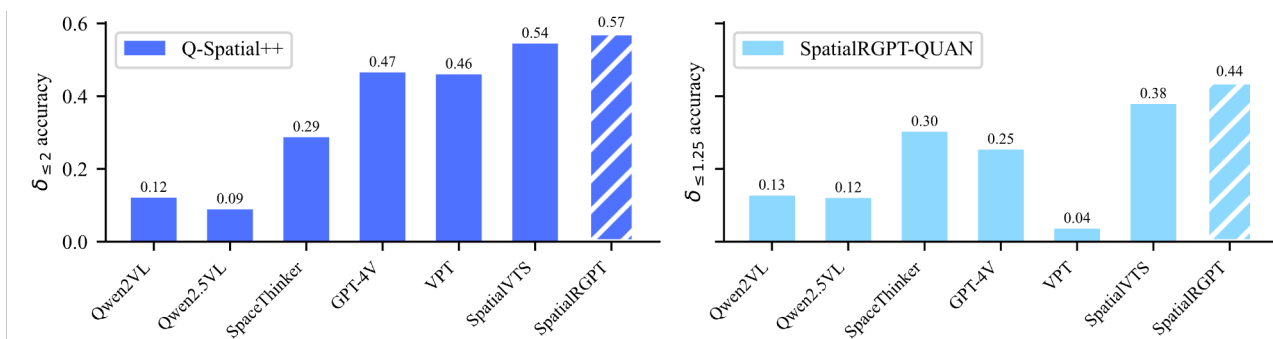


Figure 6: Performance comparison on quantitative spatial QA benchmarks. The white hatch means SpatialRGPT is trained with additional data information (depth and segmentation masks). SpatialVTS outperforms all VLMs trained with text and images, and its performance is almost on par with the SpatialRGPT model that uses depth and mask maps.

model’s QA capabilities, leading to a decline in qualitative metrics. In addition, we notice that all models except VPT demonstrate similar capabilities across the two benchmarks. This is because VPT lacks the concept of units, only providing meaningless numbers on SpatialRGPT-QUAN, and is thus rejected by the evaluator model.

4.3 Results in Complex Real-world Scenarios

To further examine our model’s ability in complex real-world scenarios, we evaluated our model on RoboSpatial-Home(Song et al. 2025) and RealWorldQA(xAI 2024). Both of these leaderboards feature densely arranged objects and complex background environments, making them suitable for evaluating robotic tasks. The results can be found in Table.4 and Table.5. For the RoboSpatial-Home benchmark in Table.4, SpatialVTS obtains the highest accuracy on the compatibility split, which requires fine distance perception; however, its performance drops considerably on the configuration split, likely due to ambiguity in the reference frame. For the RealWorldQA benchmark, SpatialVTS surpasses leading commercial models and achieves performance on par with Claude 3.5 Sonnet.

4.4 Ablation Study

To verify the effectiveness of the text thinking, we removed the CoT of the training dataset and trained it without thinking in text. The results are shown in Table.2. With text-based reasoning, the performance increases by 3% and 7%, respectively. It shows that the references in the CoT can greatly improve the quantitative reasoning ability. In order to assess the effectiveness of our training dataset, we construct a tiny dataset with 250k samples, including 120k CoT samples. Then we compare it to the model trained using the full-size dataset. From Table.3, even with just the tiny dataset, our model’s performance is already comparable to VPT, despite the latter being trained on 800k data points and utilizing an additional visual encoder. Performance can be further improved when more CoT data is used. This validates the effectiveness of our model and dataset.

4.5 Case Study

We provide a detailed case analysis to give a more concrete understanding of the model’s ability in Fig.7. In the upper-left case, SpatialRGPT directly provides answers with significant deviations without analysis. At the same time, our model is capable of thoroughly analyzing visual cues and

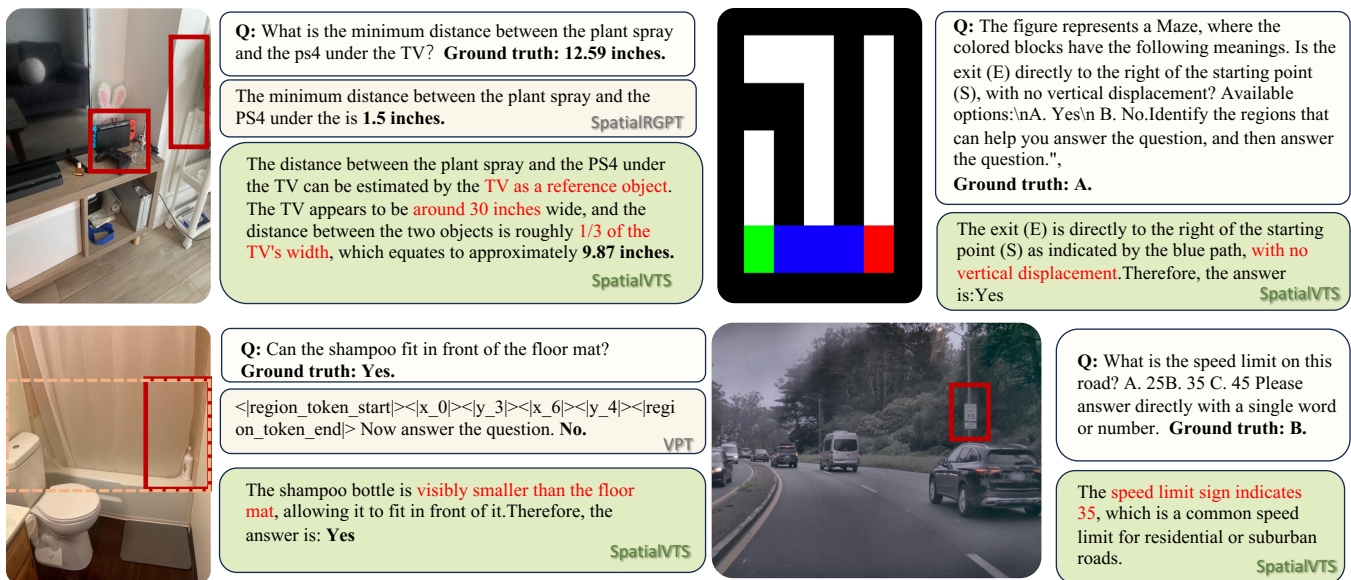


Figure 7: Some reasoning examples. The special regions are bounded by red boxes. SpatialVTS is able to use fine-grained visual grounding and reference scales. The region generated by VPT is bounded by the orange box in the third example.

providing logical reasoning and reasonable answers. It uses the common knowledge of the television’s height as the reference scale to calculate the distance. In the bottom-left case, both VPT and our model can provide key target boxes. However, the box given by VPT pays more attention to the ‘floor mat’ concept, while SpatialVTS can focus on the shampoo and the gap between the mat and the wall.

	without text	with text
SpatialRGPT-QUAN ($\delta_{\leq 1.25}$)	34.78	37.85
Q-Spatial++ ($\delta_{\leq 2}$)	48.10	54.45

Table 2: Ablation study of the textual thinking.

	250k	440k
VSR (acc)	72.89	78.21
Q-Spatial++ ($\delta_{\leq 2}$)	51.02	54.45

Table 3: Ablation study of the data amount.

In the upper right case, SpatialVTS demonstrates the ability to read maps and navigate in mazes. It can understand the semantic meaning of color blocks in the maze and reason about their spatial arrangement. In the bottom right case, which is a complex and noisy road driving environment, SpatialVTS shows visual grounding and OCR capabilities, enabling detection and recognition of speed limit signs.

5 Conclusion

We propose SpatialVTS, a framework that enhances spatial reasoning through visual and textual thinking simultaneously. During the visual thinking phase, our model can

	RoboSpatial-Home		Avg
	Configuration	Compatibility	
GPT-4o	77.20	58.10	67.65
Qwen2-VL-7B	76.42	42.85	59.64
SpaceLLaVA	61.00	61.00	61.00
VPT	77.23	43.80	60.52
SpatialVTS	70.73	67.60	69.17

Table 4: Accuracies on the configuration and compatibility splits of the RoboSpatial-Home benchmark.

Model	RealWorldQA
Gemini 1.5 Flash	64.8
Claude 3.5 Sonnet	51.9
GPT-4o	47.6
SpatialVTS	54.3

Table 5: Accuracies on the RealWorldQA leaderboard. Due to space limitations, only a subset of the leaderboard(Stanford Center for Research on Foundation Models 2024) is listed here.

identify implicit reference objects as visual cues. Then, in the textual thinking phase, it utilizes the text rationale and visual cues together to reason about the spatial relationship. We further construct a dataset to stimulate the model’s reasoning ability when there is visual information in the chain-of-thought. Experiments demonstrate that our model can substantially improve the spatial reasoning capabilities even without extra depth information.

Acknowledgements

This work was supported in part by the Key R&D Program of Zhejiang Province (2025C01212), in part by the National Nature Science Foundation of China (Grant No: 62273303), in part by Yongjiang Talent Introduction Programme (2022A-240-G), in part by Ningbo Key R&D Program (2023Z229).

References

- Abdin, M.; Aneja, J.; Awadalla, H.; Awadallah, A.; Awan, A. A.; Bach, N.; Bahree, A.; Bakhtiari, A.; Bao, J.; Behl, H.; et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- AI, R. 2023. VQASynth.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. *arXiv preprint arXiv:2308.12966*.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*.
- Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Chen, B.; Xu, Z.; Kirmani, S.; Ichter, B.; Sadigh, D.; Guibas, L.; and Xia, F. 2024a. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14455–14465.
- Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; et al. 2024b. Expanding Performance Boundaries of Open-Source Multimodal Models with Model, Data, and Test-Time Scaling. *arXiv preprint arXiv:2412.05271*.
- Chen, Z.; Wang, W.; Tian, H.; Ye, S.; Gao, Z.; Cui, E.; Tong, W.; Hu, K.; Luo, J.; Ma, Z.; et al. 2024c. How Far Are We to GPT-4V? Closing the Gap to Commercial Multimodal Models with Open-Source Suites. *arXiv preprint arXiv:2404.16821*.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024d. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.
- Cheng, A.-C.; Yin, H.; Fu, Y.; Guo, Q.; Yang, R.; Kautz, J.; Wang, X.; and Liu, S. 2024. SpatialRGPT: Grounded Spatial Reasoning in Vision Language Models. *arXiv preprint arXiv:2406.01584*.
- Driess, D.; Xia, F.; Sajjadi, M. S. M.; Lynch, C.; Chowdhery, A.; Ichter, B.; Wahid, A.; Tompson, J.; Vuong, Q.; Yu, T.; Huang, W.; Chebotar, Y.; Sermanet, P.; Duckworth, D.; Levine, S.; Vanhoucke, V.; Hausman, K.; Tous-saint, M.; Greff, K.; Zeng, A.; Mordatch, I.; and Florence, P. 2023. PALM-E: An Embodied Multimodal Language Model. *arXiv preprint arXiv:2303.03378*.
- Fu, X.; Hu, Y.; Li, B.; Feng, Y.; Wang, H.; Lin, X.; Roth, D.; Smith, N. A.; Ma, W.-C.; and Krishna, R. 2024. BLINK: Multimodal Large Language Models Can See but Not Perceive. *arXiv preprint arXiv:2404.12390*.
- Gervet, T.; Chintala, S.; Batra, D.; Malik, J.; and Chaplot, D. S. 2023. Navigating to objects in the real world. *Science Robotics*, 8(79): eadf6991.
- Gordon, D.; Kembhavi, A.; Rastegari, M.; Redmon, J.; Fox, D.; and Farhadi, A. 2018. Iqa: Visual question answering in interactive environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4089–4098.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Guo, Q.; De Mello, S.; Yin, H.; Byeon, W.; Cheung, K. C.; Yu, Y.; Luo, P.; and Liu, S. 2024. Regiongpt: Towards region understanding vision language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13796–13806.
- Hemachandra, S.; Walter, M. R.; Tellex, S.; and Teller, S. 2014. Learning spatial-semantic representations from natural language descriptions and scene classifications. In *2014 IEEE international conference on robotics and automation (ICRA)*, 2623–2630. IEEE.
- Hildebrandt, M.; Li, H.; Koner, R.; Tresp, V.; and Günnemann, S. 2020. Scene graph reasoning for visual question answering. *arXiv preprint arXiv:2007.01072*.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Jaech, A.; Kalai, A.; Lerer, A.; Richardson, A.; El-Kishky, A.; Low, A.; Helyar, A.; Madry, A.; Beutel, A.; Carney, A.; et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Kamath, A.; Hessel, J.; and Chang, K.-W. 2023. What’s “up” with vision-language models? Investigating their struggle with spatial reasoning. In *EMNLP*.

- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; et al. 2024a. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Li, C.; Wu, W.; Zhang, H.; Xia, Y.; Mao, S.; Dong, L.; Vulić, I.; and Wei, F. 2025. Imagine while Reasoning in Space: Multimodal Visualization-of-Thought. *arXiv preprint arXiv:2501.07542*.
- Li, F.; Zhang, R.; Zhang, H.; Zhang, Y.; Li, B.; Li, W.; Ma, Z.; and Li, C. 2024b. LLaVA-NeXT-Interleave: Tackling Multi-image, Video, and 3D in Large Multimodal Models. *arXiv preprint arXiv:2407.07895*.
- Liao, Y.-H.; Mahmood, R.; Fidler, S.; and Acuna, D. 2024. Reasoning Paths with Reference Objects Elicit Quantitative Spatial Reasoning in Large Vision-Language Models. *arXiv:2409.09788*.
- Liu, F.; Emerson, G.; and Collier, N. 2023. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 11: 635–651.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2023a. Improved Baselines with Visual Instruction Tuning.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023b. Visual Instruction Tuning.
- Liu, Z.; Sun, Z.; Zang, Y.; Dong, X.; Cao, Y.; Duan, H.; Lin, D.; and Wang, J. 2025. Visual-rft: Visual reinforcement fine-tuning. *arXiv preprint arXiv:2503.01785*.
- Shao, H.; Qian, S.; Xiao, H.; Song, G.; Zong, Z.; Wang, L.; Liu, Y.; and Li, H. 2024. Visual cot: Unleashing chain-of-thought reasoning in multi-modal language models. *arXiv e-prints*, arXiv–2403.
- Song, C. H.; Blukis, V.; Tremblay, J.; Tyree, S.; Su, Y.; and Birchfield, S. 2025. RoboSpatial: Teaching Spatial Understanding to 2D and 3D Vision-Language Models for Robotics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Oral Presentation.
- Stanford Center for Research on Foundation Models. 2024. HELM VHELM RealWorldQA Leaderboard. https://nlp.stanford.edu/helm/vhelm/?group=real_world_qa. Last updated: 2024-11-08; Accessed: 2025-08-02.
- Team, Q. 2024. QVQ: To See the World with Wisdom.
- Thawakar, O.; Dissanayake, D.; More, K.; Thawkar, R.; Heakl, A.; Ahsan, N.; Li, Y.; Zumri, M.; Lahoud, J.; Anwer, R. M.; et al. 2025. Llamav-o1: Rethinking step-by-step visual reasoning in llms. *arXiv preprint arXiv:2501.06186*.
- Wald, J.; Dhama, H.; Navab, N.; and Tombari, F. 2020. Learning 3d semantic scene graphs from 3d indoor reconstructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3961–3970.
- Walter, M. R.; Hemachandra, S.; Homberg, B.; Tellex, S.; and Teller, S. J. 2013. Learning Semantic Maps from Natural Language Descriptions. In *Robotics: science and systems*, volume 2.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Fan, Y.; Dang, K.; Du, M.; Ren, X.; Men, R.; Liu, D.; Zhou, C.; Zhou, J.; and Lin, J. 2024. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution. *arXiv preprint arXiv:2409.12191*.
- xAI. 2024. Grok-1.5 Vision Preview: RealWorldQA Dataset. <https://x.ai/news/grok-1.5v>. Dataset released under CC BY-ND 4.0.
- Xu, G.; Jin, P.; Hao, L.; Song, Y.; Sun, L.; and Yuan, L. 2024. Llava-o1: Let vision language models reason step-by-step. *arXiv preprint arXiv:2411.10440*.
- Xu, G.; Jin, P.; Li, H.; Song, Y.; Sun, L.; and Yuan, L. 2025. LLaVA-CoT: Let Vision Language Models Reason Step-by-Step. *arXiv:2411.10440*.
- Yu, R.; Ma, X.; and Wang, X. 2025. Introducing Visual Perception Token into Multimodal Large Language Model. *arXiv preprint arXiv:2502.17425*.
- Yuan, Y.; Li, W.; Liu, J.; Tang, D.; Luo, X.; Qin, C.; Zhang, L.; and Zhu, J. 2024. Osprey: Pixel understanding with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28202–28211.