

CoV: Chain-of-View Prompting for Spatial Reasoning

Haoyu Zhao^{1*} Akide Liu^{2*} Zeyu Zhang^{2*} Weijie Wang^{1*}
 Feng Chen³ Ruihan Zhu¹ Gholamreza Haffari² Bohan Zhuang^{1†}
¹ZIP Lab, Zhejiang University ²Monash University ³AIML, Adelaide University

*Equal contribution. †Corresponding author: bohan.zhuang@gmail.com.



Figure 1: The **Chain-of-View** prompting framework. Given a spatial query and its corresponding 3D scene, CoV facilitates a coarse-to-fine active reasoning process to derive the answer. **(Bottom)** v_1 to v_4 denote the task-relevant viewpoints strategically selected from the initial candidate view set by our View Selection Agent. **(Left and Right)** The interleaved action-reasoning chain demonstrates how the agent dynamically adjusts its perspective (e.g., rotation and movement) to gather discriminative visual evidence and resolve spatial ambiguities. **(Center)** The visualized camera frustums depict the autonomous exploration trajectory, where the agent bridges the gap between fragmented local views and global spatial context to reach a grounded conclusion.

Abstract

Embodied question answering (EQA) in 3D environments often requires collecting context that is distributed across multiple viewpoints and partially occluded. However, most recent vision-language models (VLMs) are constrained to a fixed and finite set of input views, which limits their ability to acquire question-relevant context at inference time and hinders complex spatial reasoning. We propose **Chain-of-View (CoV) prompting**, a *training-free, test-time* reasoning framework that transforms a VLM into an active viewpoint reasoner through a coarse-to-fine exploration process. CoV first employs a *View Selection* agent to filter redundant frames and identify question-aligned anchor views. It then performs *fine-grained view adjustment* by interleaving iterative reasoning with discrete camera actions, ob-

taining new observations from the underlying 3D scene representation until sufficient context is gathered or a step budget is reached. We evaluate CoV on OpenEQA across four mainstream VLMs and obtain an average +11.98% improvement in LLM-Match, with a maximum gain of +13.62% on Qwen3-VL-Flash. CoV further exhibits test-time scaling: increasing the minimum action budget yields an additional +2.54% average improvement, peaking at +3.73% on Gemini-2.5-Flash. On ScanQA and SQA3D, CoV delivers strong performance (e.g., 116 CIDEr / 31.9 EM@1 on ScanQA and 51.1 EM@1 on SQA3D). Overall, these results suggest that question-aligned view selection coupled with open-view search is an effective, model-agnostic strategy for improving spatial reasoning in 3D EQA without additional training. Code: <https://github.com/>

1 Introduction

As artificial intelligence transitions from digital domains to physical reality, Embodied Question Answering (EQA) has emerged as a critical capability for enabling intuitive, human-centric interaction with the environment. It holds significant potential across domains like robotics, autonomous navigation, and human-computer interaction. In EQA settings, the agent processes a textual question based on a sequence of egocentric images (optionally with a 3D scene representation such as point clouds or 3D meshes). The agent must perceive and reason within the real environment to derive the correct answer.

However, existing methods (Mo and Liu, 2024; Fu et al., 2024; Zhu et al., 2024; Li et al., 2024c) encounter a substantial limitation: conventional methods use a limited and fixed set of viewpoints as input (see Figure 2), making it difficult for VLMs to acquire sufficient question-relevant views. In complex embodied QA tasks, answers are not immediately apparent, and a question often requires multi-step reasoning to solve. For example, for the question “Where can I get some pop drinks?”, the scene does not directly show soda. The model must invoke world knowledge and navigate autonomously to locate objects like a refrigerator. Answering such complex real-world questions requires sufficient question-relevant context and cannot be accomplished through one-step answer generation.

In this paper, we propose the **chain-of-view (CoV)** prompting framework (see Figure 1), a two-stage agent system designed to shift from passive observation to active exploration and iterative reasoning. Specifically, the framework operates in a coarse-to-fine manner: in the coarse-grained view selection stage, the **View Selection Agent** selects a small set of the most question-relevant views from the available viewpoints as anchor views for subsequent exploration. Meanwhile, we provide the agent with a bird’s-eye view of the entire scene to facilitate a global understanding of the environment. In the fine-grained view adjustment stage, the **CoV Agent** executes an action-reasoning loop. At each step, the VLM generates an action instruction based on the current observation and the question (e.g., forward-movement or right-rotation). The action is mapped to a rigid-body camera transformation, producing the next observation, which is fed back

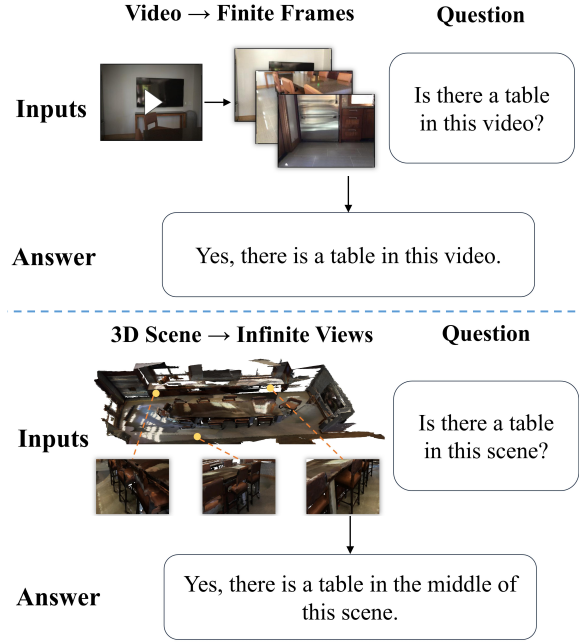


Figure 2: **Video VLM vs. CoV.** Unlike prior approaches (**top**) that rely on fixed-frame video inputs and answer from a limited temporal window, our chain-of-view framework (**bottom**) explores an open-ended view space constructed from a 3D scene. CoV dynamically selects informative viewpoints and performs step-by-step reasoning during inference, enabling more complete and grounded answers without additional training.

to the VLM for the next reasoning step. The process terminates when the CoV agent determines that sufficient information has been acquired or when a predefined limit on action steps is reached, at which point the final answer is produced.

Unlike prior 2D VLMs that rely on predetermined views (as illustrated in Figure 2), our chain-of-view framework addresses complex embodied QA problems through multi-step reasoning. Moreover, we improve the alignment between visual content and the question via explicit view selection and fine-grained adjustment.

To evaluate the efficacy of our methodology, we conduct comprehensive experiments on the latest EQA benchmark OpenEQA (Majumdar et al., 2024), which serve as standard metrics for assessing 3D scene understanding and question answering capabilities. We evaluate both mainstream open-source and proprietary VLMs, and applying the CoV framework yields an average improvement of 11.98%, with a maximum gain of 13.62% on Qwen3-VL-Flash. We further empirically verify the test-time scaling capability of CoV: as the number of action steps increases, the agent’s score

improves by an average of 2.54%, with a maximum gain of 3.73% on Gemini-2.5-Flash.

Qualitative analyses further validate that our CoV framework produces more coherent and interpretable reasoning chains, particularly in complex or cluttered environments. These results demonstrate the potential of test-time scaling strategies to enhance scene understanding without requiring additional model training or dataset-specific tuning, making our framework robust and adaptable across diverse 3D tasks and domains.

The main contributions of our work can be summarized as follows:

- We propose Chain-of-View Prompting, a test-time reasoning framework that enhances VLMs’ ability to handle complex spatial reasoning in embodied question answering. By leveraging coarse-to-fine view selection and camera adjustment, the agent acquires sufficient question-relevant views to answer spatially complex questions, thereby improving performance on EQA tasks.
- CoV prompting enables test-time scaling. As the number of exploration steps increases, the agent’s performance improves gradually.
- Experimental results on the latest embodied QA benchmarks demonstrate significant improvements through our systematic view exploration approach. Our method achieves up to a 13.62% improvement on the OpenEQA benchmark.

2 Related Work

Recent advances in 3D scene understanding have unified perception and language. Methods like Vote2Cap-DETR (Chen et al., 2023a), D3Net (Chen et al., 2022), and SpaCap3D (Wang et al., 2022) integrate object localization and description generation, enabling more grounded scene understanding for robotics, AR/VR, and embodied AI. 3D vision-language models such as LLaVA-3D (Zhu et al., 2024) and LL3DA (Chen et al., 2024) further advance scene understanding by synthesizing 2D multimodal perception with 3D spatial context. These architectures leverage multi-view images augmented with 3D positional embeddings, facilitating more context-aware reasoning without dependency on external object proposals or segmentation mechanisms.

Test-time reasoning. Large models such as Qwen, ChatGPT and Gemini (Bai et al., 2025; OpenAI, 2023; Georgiev et al., 2024) show strong performance in multi-modal reasoning tasks. Due to high fine-tuning costs, recent work explores efficient adaptation methods that keep pretrained weights. In-context learning (Brown et al., 2020; Sahoo et al., 2024), prompt engineering, and chain-of-thought prompting (Wei et al., 2022) guide model behavior at inference time. Recent works like Simple Scaling (Muennighoff et al., 2025a), adaptive compute (Snell et al., 2024), and calibration (McKenna and Carse, 2024) offer practical, training-free improvements.

Scene understanding. 3D scene understanding primarily encompasses tasks such as 3D question answering (Ma et al., 2023; Azuma et al., 2022) and 3D dense captioning (Chen et al., 2020; Achlioptas et al., 2020). Early 3D dense captioning used a detect-then-describe pipeline (Chen et al., 2021; Wang et al., 2022). Newer methods adopt end-to-end transformers (Chen et al., 2023a, 2022; Huang et al., 2025a,b) to predict object-caption pairs directly. Methods for embodied QA often incorporate multi-modal fusion (Mo and Liu, 2024), navigation-conditioned reasoning (Zheng et al., 2024), and LLM grounding in 3D scenes (Hong et al., 2023) to improve spatial and semantic understanding. Agent3D-Zero (Zhang et al., 2024a) introduces a zero-shot framework in which a VLM actively selects viewpoints to understand 3D scenes, conceptually similar to our approach; however, it relies on VLM-predicted grid-based locations on the bird’s-eye view, which can be inaccurate and differs from our action-instruction-based method.

3D VLMs. Recent 3D VLM works integrate point clouds (Huang et al., 2023; Chen et al., 2024; Yang et al., 2023; Zhang et al., 2024b) and multi-view images (Fu et al., 2024; Hong et al., 2023; Qi et al., 2025) into large language models for scene reasoning. LL3DA (Chen et al., 2024) encodes global features from scene-level point clouds. LEO (Yang et al., 2023) and Chat-Scene (Zhang et al., 2024b) segment and encode object-level features. 3D-LLM (Hong et al., 2023) and Scene-LLM (Fu et al., 2024) use object-centric patches from multi-view images. LLaVA-3D (Zhu et al., 2024) builds on 2D LLMs with 3D positional embeddings to structure image patches spatially.

3 Method

3.1 Problem Setting

The goal of the embodied question answering task is to answer questions related to a 3D scene. The questions can span multiple categories, such as object recognition, attribute recognition, object localization, and spatial reasoning etc. (Majumdar et al., 2024). Unlike conventional 2D VQA with fixed frame sequences, embodied question answering utilizes either multi-view RGB-D images with camera poses or reconstructed 3D meshes, enabling reasoning from arbitrary viewpoints aligned with question semantics.

Formally, let S denote the 3D scene representation (e.g., a point cloud or mesh) and Q be the natural language query. A 3D scene implies a continuous space of potential viewpoints Ω . However, directly reasoning over the raw global geometry S is often inefficient and lacks fine-grained visual details. In practice, the input for embodied question answering is combined with a sequence of frames $\mathcal{V} = \{v_1, \dots, v_T\}$ sampled from a video episode of the scene. The objective is to generate a textual answer A that accurately addresses the query Q based on these observations. This process is modeled as maximizing the conditional probability:

$$A^* = \underset{A}{\operatorname{argmax}} P(A | S, \mathcal{V}, Q). \quad (1)$$

3.2 Method Overview

We propose chain-of-view prompting, a test-time search framework for embodied visual question answering. The goal of CoV prompting is to infer the correct answer A from a given question text Q and the corresponding 3D scene. The overall pipeline consists of two primary stages: (i) Coarse-grained View Selection, and (ii) Fine-grained View Adjustment. Initially, we filter the full set of input view frames to a subset of candidate views based on their relevance to the question. We dynamically adjust the perspectives within this subset to produce a refined sequence of viewpoints. Finally, an answer is synthesized based on these optimized viewpoints.

3.3 Coarse-Grained View Selection

\mathcal{V} is sampled from the video episode of the corresponding scene, containing many redundant frames with low information density. Only a small number of key frames are relevant to the question Q , while a large amount of irrelevant information can

interfere with the agent’s correct judgment. Therefore, we employ a View Selection Agent to filter the initially available views and extract the frames most relevant to the question. The prompt template of the view selection agent is provided in Section A. Given the input (Q, \mathcal{V}) , the view selection agent selects a reduced subset of frames $\mathcal{V}' = \{v_{i_1}, \dots, v_{i_K}\}$, where $K \ll T$.

This initial filtering step significantly reduces redundancy in the input data and focuses on views that are most likely to contain information relevant to answering the question. By narrowing down the search space, we enable more efficient processing in the subsequent fine-grained phase, effectively addressing the challenge of searching through the vast view space of 3D scenes.

3.4 Fine-Grained View Adjustment

In prior EQA approaches, the agent passively infers an answer over a limited and fixed set of input frames and cannot actively acquire informative observations from the environment. This one-step generation paradigm overlooks environmental details that may be relevant to the question and constrains the agent’s performance. Inspired by chain-of-thoughts (Wei et al., 2022), we aim to provide more detailed environmental information through fine-grained viewpoint adjustment while eliciting deeper thinking. Therefore, we employ the Chain-of-View Agent for fine-grained view adjustment. The prompt template is provided in Section A.

We leverage the 3D scene representation to dynamically generate new perspectives that reveal information not visible from the initially selected views. For each initial view $v_i \in \mathcal{V}'$, the CoV agent generates a sequence of actions $\{a_1, \dots, a_L\}$, where $a_t \in \mathcal{A}$ for $t = 1, \dots, L$, and \mathcal{A} denotes the agent’s action space. It consists of discrete translational and rotational actions:

- **Translational actions:** move forward, move backward, move left, move right, move up, and move down. Each action corresponds to a fixed displacement along the respective axis of the agent’s local coordinate frame.
- **Rotational actions:** yaw (rotate left/right), pitch (look up/down), and roll (tilt clockwise/counterclockwise). Each rotation adjusts the agent’s orientation around its local coordinate axes.
- **View switch actions:** switch to $v_i, v_i \in \mathcal{V}'$.



Figure 3: **Action-reasoning chain of the CoV agent.** The CoV agent executes an iterative action–reasoning chain. For the question “What should I do to cool down?”, the agent first selects view 6 from the input images as an anchor. It then adjusts the viewpoint at each reasoning step to acquire new observations. Once the agent determines that sufficient information has been obtained, it outputs the answer “turn on the air conditioner.”

Agent can switch to any view anchor obtained through the coarse view selection stage.

Specifically, the agent’s context is defined as $C_0 = \{Q, \mathcal{V}'\}$ at start. At each step, the CoV agent thinks over the current observation and the question to generate an action instruction a_t . We convert a_t into an SE(3) transformation matrix, which updates the camera pose and yields a new viewpoint v_i^{t+1} :

$$v_i^{t+1} = \text{Transform}(v_i; a_t, S). \quad (2)$$

v_i^{t+1} is appended into the agent’s context, which is then fed into the model for the next step:

$$C_{t+1} = \{Q, \mathcal{V}', v_i^1, \dots, v_i^{t+1}\}. \quad (3)$$

The reasoning process terminates when the CoV agent determines that sufficient information has been collected to answer the question or when a predefined action step limit is reached, at which point the final answer is produced.

Figure 3 presents a multi-step reasoning example of the CoV agent. Fine-grained view adjustment enables the agent to observe regions that were previously occluded or blurred, allowing it to acquire richer question-relevant environmental details and thereby answer the question more accurately.

4 Experiments

4.1 Benchmarks and Metrics

We evaluate our method on the OpenEQA (Majumdar et al., 2024), ScanQA (Azuma et al., 2022), and SQA3D (Ma et al., 2023) benchmarks, covering both mainstream open-source and proprietary models.

Sourced from over 180 real-world environments (ScanNet and HM3D (Dai et al., 2017; Ramakrishnan et al., 2021)), OpenEQA is a challenging open-vocabulary benchmark, designed to evaluate embodied question answering capabilities in the era of foundation models. Besides, we include ScanQA and SQA3D, two representative datasets for spatial reasoning challenges. ScanQA is a large-scale 3D question answering dataset comprising over 41,000 question-answer pairs. It focuses on object-grounded QA by linking natural language queries to specific 3D objects within richly annotated RGB-D scans, facilitating spatial reasoning and object localization. In contrast, SQA3D emphasizes situated reasoning, requiring agents to understand their position and orientation within a 3D scene. It includes 33,400 reasoning questions spanning 6,800 unique situations from 650 ScanNet scenes, presenting complex challenges such as

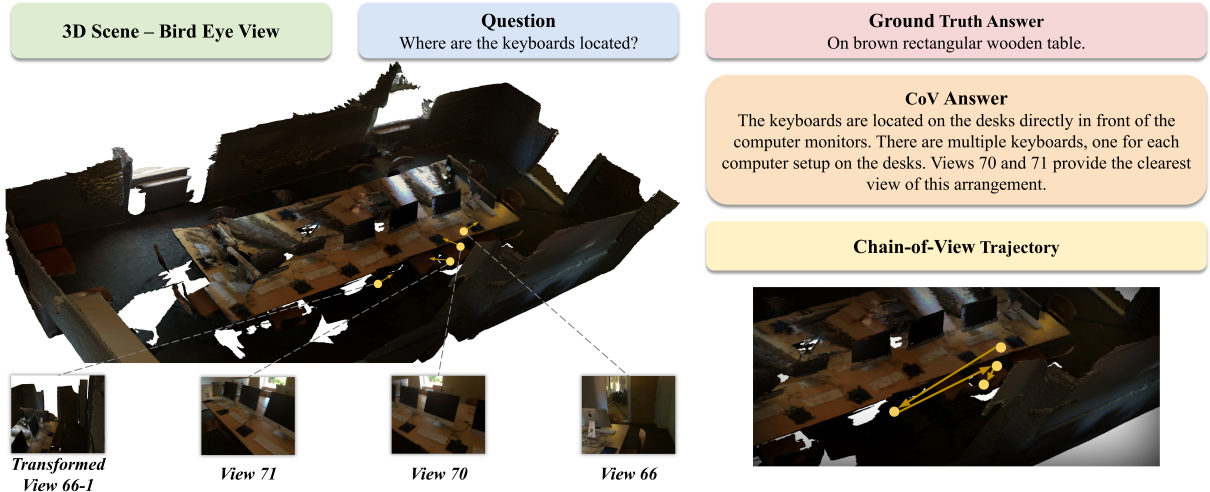


Figure 4: Visualization of CoV reasoning results. Our method selects informative views and produces coherent multi-step answers grounded in the spatial context.

spatial relationships, commonsense understanding, and multi-hop inference.

We use the **LLM-Match** metric proposed in OpenEQA. This metric utilizes an LLM judge to compare a predicted answer against the ground-truth answer, assigning a score $\gamma_i \in \{1, \dots, 5\}$. The final metric is computed by normalizing and averaging them to a percentage scale over question number N :

$$\text{LLM-Match} = \frac{1}{N} \sum_{i=1}^N \left(\frac{\gamma_i - 1}{4} \right) \times 100\%. \quad (4)$$

For ScanQA and SQA3D, we adopt a comprehensive set to assess answer accuracy. Specifically, CIDEr (C) measures consensus with human annotations; BLEU-4 (B-4) captures n-gram overlap; METEOR (M) considers both precision and recall with synonym matching; ROUGE-L (R) evaluates the longest common subsequence; and Exact Match at top-1 (EM1) reflects the strict correctness of generated answers.

4.2 Implementation Details

The input for each question–answer pair consists of video frames uniformly sampled from a scene video episode at a ratio of 10:1, together with the textual question Q . For the baseline, we provide all images in a single pass and let the model directly generate an answer. For CoV, we first feed all images to the View Selection agent for coarse filtering, and the selected views are then passed to the CoV agent to answer the question; both agents use the same underlying VLM. The prompt templates used are provided in Section A.

4.3 Main Results

For ScanQA and SQA3D, we instantiate CoV with three representative VLMs: GPT-4.1, Gemini Pro Flash, and InternVL (OpenAI, 2023; Georgiev et al., 2024; Chen et al., 2023b). We run all evaluations in the val-unseen set for fair, zero-shot comparison with 3D VLM baselines. Table 1 shows that our method achieves state-of-the-art performance on ScanQA, with notable improvements in CIDEr and EM@1. On SQA3D, our method remains competitive with 51.1 EM@1. These results demonstrate our method’s effectiveness in both generating human-like responses and providing accurate answers across 3D QA benchmarks.

For OpenEQA (Majumdar et al., 2024), we evaluate Qwen3-VL-Flash (Bai et al., 2025), GLM-4.6V (Team et al., 2025), Gemini-2.5-Flash (Cohan et al., 2025), and GPT-4o-mini (OpenAI et al., 2024). The evaluation results are shown in Table 2. Compared to the setting without CoV, our method achieves an average improvement of 11.98% and a maximum gain of 13.62% (on Qwen3-VL-Flash) in a training-free manner.

Evaluation results confirm that our chain-of-view framework produces higher quality answers by iteratively refining visual understanding through strategically selected views, enabling deep spatial reasoning in complex 3D environments.

4.4 Test-Time Scaling

We investigate the test-time scaling behavior of chain-of-view prompting by quantitatively analyzing the relationship between the number of action steps and performance. When the number of action

Table 1: **Quantitative comparison with SOTA models on ScanQA (val) and SQA3D (test).** ‘‘C’’ stands for ‘‘CIDEr’’, ‘‘B-4’’ for ‘‘BLEU-4’’, ‘‘M’’ for ‘‘METEOR’’, ‘‘R’’ for ‘‘ROUGE’’, and ‘‘EM@1’’ for top-1 exact match. BEV denotes whether a model takes bird’s-eye view as an input.

Model	Venue	BEV	ScanQA (val)					SQA3D (test)
			C	B-4	M	R	EM@1	EM@1
Task-specific models								
Scan2Cap (Chen et al., 2021)	-	✗	-	-	-	-	-	41.0
ScanRefer+MCAN (Yu et al., 2019)	ECCV2020	✗	55.4	7.9	11.5	30.0	18.6	-
ClipBERT (Lei et al., 2021)	CVPR2021	✗	-	-	-	-	-	43.3
ScanQA (Azuma et al., 2022)	CVPR2022	✗	64.9	10.1	13.1	33.3	21.1	47.2
3D-VisTA (Zhu et al., 2023)	ICCV2023	✗	69.6	10.4	13.9	35.7	22.4	48.5
3D LMMs								
3D-LLM (FlanT5) (Hong et al., 2023)	NeurIPS2023	✗	69.4	12.0	14.5	35.7	20.5	-
LL3DA (Chen et al., 2024)	CVPR2024	✗	76.8	13.5	15.9	37.3	-	-
Chat-3D v2 (Huang et al., 2024a)	-	✗	87.6	14.0	-	-	-	54.7
LEO (Huang et al., 2024c)	ICML2024	✗	101.4	13.2	20.0	49.2	24.5	50.0
Scene-LLM (Fu et al., 2024)	WACV2025	✗	80	12.0	16.6	40.0	27.2	54.2
ChatScene (Huang et al., 2024b)	CVPR2024	✗	87.7	14.3	18.0	41.6	21.6	54.6
Zero-shot 2D LMMs								
VideoChat2 (Li et al., 2024b)	CVPR2024	✓	49.2	9.6	9.5	28.2	19.2	37.3
LLaVA-NeXT-Video (Li et al., 2024a)	-	✓	46.2	9.8	9.1	27.8	18.7	34.2
LLaVA-Video (Zhang et al., 2024c)	-	✓	88.7	-	-	-	-	48.5
Agent3D-Zero (Zhang et al., 2024a)	ECCV2024	✓	71.8	4.4	16.0	37.0	17.5	-
CoV (Ours)	-	✓	116	16.9	24.5	50.4	31.9	51.1

steps is not constrained, the statistics of question counts and scores at different action steps for the CoV agent on the OpenEQA dataset are shown in Figure 5. Questions that require more action steps exhibit a clear upward trend in scores. However, most questions involve only a small number of action steps, typically between one and three. If the agent were able to execute more action steps, its performance on embodied QA tasks would be expected to further improve.

Inspired by the budget forcing strategy used in S1 (Muennighoff et al., 2025b), we set a lower bound on the number of action steps in the CoV agent’s prompt template (see Section A). As shown in Figure 6, increasing the number of action steps gradually improves. Compared to setting the minimum number of action steps to 1, increasing the action-step limit yields an average improvement of 2.54%.

These results demonstrate that multi-step reasoning enhances agent performance on EQA tasks and highlight the potential of our method for training-free test-time scaling.

4.5 Ablation Studies

We conduct an ablation study to examine the role of coarse-grained view selection. We evaluate the agent’s performance on OpenEQA (Majumdar

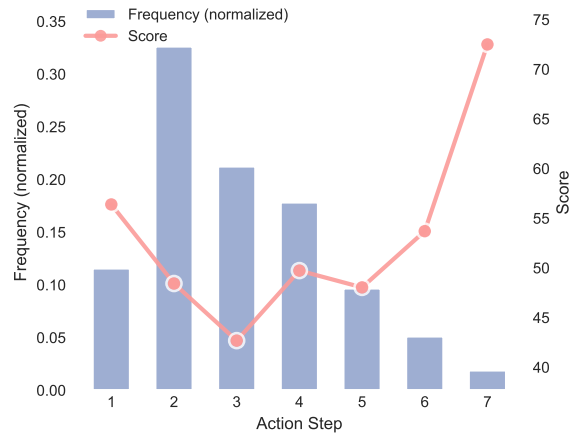


Figure 5: **Action Step Analysis.** Distribution of action steps for Qwen3-VL-Flash (Bai et al., 2025) on the OpenEQA (Majumdar et al., 2024) dataset.

et al., 2024) with and without the coarse-grained view selection agent. Without the coarse-grained stage to filter views, the CoV agent is exposed to a large number of redundant and low-information-density frames, making it harder to identify a question-relevant anchor to initiate actions.

As shown in Table 3, removing the view selection agent from our method degrades VLM performance, with an average drop of 2.71%. This result demonstrates that coarse-grained view selection is

Table 2: **Performance comparison on the OpenEQA benchmark.** During evaluation, we fix the temperature of each model to 0. LLM-Match is measured using the evaluation protocol provided by OpenEQA. n step denotes that the minimum number of action steps for the model is set to n . **Scaling Improvement** measures the gain obtained by increasing the number of action steps, defined as the difference between the best CoV result and CoV-1. **Boldface** indicates the **best** results.

Method	Qwen3-VL-Flash	GLM-4.6V	GPT-4o-mini	Gemini-2.5-Flash
Baseline	52.65	62.40	45.87	52.30
CoV(1 Step)	58.75	67.00	49.85	57.10
2 Step	59.07	67.33	50.15	57.15
3 Step	59.03	67.20	49.60	58.80
4 Step	59.82	66.30	50.80	58.15
5 Step	59.03	67.23	51.56	58.82
6 Step	59.39	67.70	51.41	59.23
7 Step	59.77	67.64	51.62	59.16
Improvement	↑ 13.62%	↑ 8.50%	↑ 12.54%	↑ 13.25%
Scaling Improvement	↑ 1.82%	↑ 1.04%	↑ 3.55%	↑ 3.73%




Figure 6: **Test-time scaling ability of CoV.** On OpenEQA, we evaluate different VLMs and observe that the performance of all models gradually improves as the number of action steps increases.

an essential component of our approach.

We further analyze the number of coarse-selected views K on Qwen3-VL-Flash. As shown in Table 4, selecting $K = 5$ views achieves the best performance, while using fewer views misses relevant context and using more views introduces redundant observations.

We also examine whether the gains of CoV can be explained by view selection alone. In Table 5, a static VLM equipped with coarse view selection still performs much worse than CoV, indicating that informative anchor views are helpful but insufficient without fine-grained multi-step view search.

Table 3: **Ablation study on view selection component.**

This table shows the performance comparison between CoV and CoV without the coarse view selection mechanism across different models. CVS denotes coarse view selection.

Model	CoV(3 Step)	CoV w/o CVS
Qwen3-VL-Flash	59.03	57.50
GLM-4.6V	67.20	62.43
GPT-4o-mini	49.60	46.74
Gemini-2.5-Flash	58.80	57.11

Table 4: Ablation study on the number of coarse-selected views K .

Number of K	Scores on OpenEQA
$K = 3$	48.16
$K = 5$	59.03
$K = 7$	52.55

Finally, we test whether simply exposing a static VLM to the new views generated during CoV exploration is enough. Results in Table 6 show that this setting remains clearly below CoV, suggesting that the improvement comes not only from accessing additional views, but from the sequential reasoning process that actively gathers and uses them.

4.6 Qualitative Result

Figure 4 presents a qualitative example illustrating how CoV selects and reasons over a sequence of informative views. The agent progressively explores

Table 5: Ablation study on fine-grained view search.

Method	Scores on OpenEQA
CoV	59.03
CVS + static VLM	35.77

Table 6: Ablation study on intermediate results obtained during CoV search.

Method	Scores on OpenEQA
CoV	59.03
Static VLM + new views	47.69

the scene, identifying relevant object locations with high spatial precision and integrating these observations into its reasoning process. Through multi-step reasoning, CoV generates answers that are not only semantically aligned with the question but also consistent with the spatial layout of the scene. The predicted answer closely matches the ground truth, highlighting the agent’s capability to perform detailed spatial reasoning.

Both qualitative and quantitative analyses demonstrate the effectiveness of the CoV method. Through coarse-grained view selection and multi-step actions, the agent autonomously acquires more question-relevant information and is able to perform complex embodied QA tasks. Additional qualitative visualization examples are provided in Section B.

5 Conclusion

In conclusion, our work rethinks embodied question answering through the lens of viewpoint-aware reasoning. By adopting a coarse-to-fine viewpoint adjustment strategy, the view selection agent and the CoV agent can acquire question-relevant observations and address complex spatial reasoning problems through multi-step reasoning. We further examine the test-time scaling capability of the proposed method and observe that the agent’s performance improves as the number of action steps increases. Beyond performance gains, our framework represents a conceptual shift—emphasizing not just what the model sees, but how it sees, remembers, and reasons. We believe this chain-of-view prompting framework will open new possibilities for embodied AI systems that must act, adapt, and explore in complex real-world spaces.

Limitations

While the CoV Prompting approach offers significant advancements, it is not without limitations. The coarse-to-fine paradigm may struggle in highly dynamic or cluttered environments where rapid context shifts occur, potentially leading to misinterpretation of scene elements. Additionally, the effectiveness of CoV relies on the quality and relevance of the selected views; suboptimal view selection can impair reasoning accuracy. Another limitation is that CoV does not currently employ an explicit stopping condition: the decision to terminate exploration mainly relies on the VLM’s inner reflection about whether sufficient information has been gathered. While this design preserves flexibility, it may lead to unstable stopping behaviors or unnecessary exploration. Developing advanced view selection algorithms, adaptive reasoning mechanisms, and more explicit termination criteria represents a promising direction for future research.

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. *16th European Conference on Computer Vision (ECCV)*.
- Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. 2022. *Scanqa: 3d question answering for spatial scene understanding*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19107–19117.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025. *Qwen3-vl technical report*. *Preprint*, arXiv:2511.21631.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, and 1 others. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. 2020. Scanrefer: 3d object localization in rgb-d scans using natural language. *16th European Conference on Computer Vision (ECCV)*.
- Dave Zhenyu Chen, Ali Gholami, Matthias Nießner, and Angel X. Chang. 2022. D3net: A unified speaker-listener architecture for 3d dense captioning and visual grounding. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

- Sijin Chen, Xin Chen, Chi Zhang, Mingsheng Li, Gang Yu, Hao Fei, Hongyuan Zhu, Jiayuan Fan, and Tao Chen. 2024. L13da: Visual interactive instruction tuning for omni-3d understanding reasoning and planning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26428–26438.
- Sijin Chen, Hongyuan Zhu, Xin Chen, Yinjie Lei, Gang Yu, and Tao Chen. 2023a. End-to-end 3d dense captioning with vote2cap-detr. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11124–11133.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, and 1 others. 2023b. Internvl: Scaling up vision foundation models and aligning for generic vision-language tasks. *arXiv preprint arXiv:2312.14238*.
- Zhenyu Chen, Ali Gholami, Matthias Nießner, and Angel X. Chang. 2021. Scan2cap: Context-aware dense captioning in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3195–3204.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, and 3416 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *Preprint*, arXiv:2507.06261.
- 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.
- Rao Fu, Jingyu Liu, Xilun Chen, Yixin Nie, and Wenhan Xiong. 2024. Scene-llm: Extending language model for 3d visual understanding and reasoning. *arXiv preprint arXiv:2403.11401*.
- Petko Georgiev and 1 others. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 2023. 3d-llm: Injecting the 3d world into large language models. In *Advances in Neural Information Processing Systems*, volume 36, pages 20482–20494. Curran Associates, Inc.
- Haifeng Huang, Yilun Chen, Zehan Wang, Rongjie Huang, Runsen Xu, Tai Wang, Luping Liu, Xize Cheng, Yang Zhao, Jiangmiao Pang, and Zhou Zhao. 2024a. Chat-scene: Bridging 3d scene and large language models with object identifiers. *Preprint*, arXiv:2312.08168.
- Haifeng Huang, Yilun Chen, Zehan Wang, Rongjie Huang, Runsen Xu, Tai Wang, Luping Liu, Xize Cheng, Yang Zhao, Jiangmiao Pang, and Zhou Zhao. 2024b. Chat-scene: Bridging 3d scene and large language models with object identifiers. *Preprint*, arXiv:2312.08168.
- Haifeng Huang, Zehan Wang, Rongjie Huang, Luping Liu, Xize Cheng, Yang Zhao, Tao Jin, and Zhou Zhao. 2023. Chat-3d v2: Bridging 3d scene and large language models with object identifiers. *arXiv preprint arXiv:2312.08168*.
- Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, and Siyuan Huang. 2024c. An embodied generalist agent in 3d world. *Preprint*, arXiv:2311.12871.
- Ting Huang, Zeyu Zhang, Yemin Wang, and Hao Tang. 2025a. 3d coca: Contrastive learners are 3d captioners. *arXiv preprint arXiv:2504.09518*.
- Ting Huang, Zeyu Zhang, Ruicheng Zhang, and Yang Zhao. 2025b. Dc-scene: Data-centric learning for 3d scene understanding. *arXiv preprint arXiv:2505.15232*.
- Jie Lei, Linjie Li, Luwei Zhou, Zhe Gan, Tamara L. Berg, Mohit Bansal, and Jingjing Liu. 2021. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7331–7341.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024a. Llava-onevision: Easy visual task transfer. *Preprint*, arXiv:2408.03326.
- Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. 2024b. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22195–22206.
- Yifan Li, Yikang Wang, Yang Zhao, Ziang Zhang, and Zhou Zhao. 2024c. Video-3d llm: Learning position-aware video representation for 3d scene understanding. *arXiv preprint arXiv:2412.00493*.
- Xiaojian Ma, Silong Yong, Zilong Zheng, Qing Li, Yitao Liang, Song-Chun Zhu, and Siyuan Huang. 2023. Sqa3d: Situated question answering in 3d scenes. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mccvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, and 5 others.

2024. Openeqa: Embodied question answering in the era of foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16488–16498.
- Stephen McKenna and Jacob Carse. 2024. Calibrating where it matters: Constrained temperature scaling. *arXiv preprint arXiv:2406.11456*.
- Wentao Mo and Yang Liu. 2024. Bridging the gap between 2d and 3d visual question answering: A fusion approach for 3d vqa. *arXiv preprint arXiv:2402.15933*.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. 2025a. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori B Hashimoto. 2025b. s1: Simple test-time scaling. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20286–20332.
- OpenAI. 2023. Gpt-4 technical report. <https://openai.com/research/gpt-4>.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. Gpt-4 technical report. *Preprint*, arXiv:2303.08774.
- 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*.
- Santhosh K Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alex Clegg, John Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, and 1 others. 2021. Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai. *arXiv preprint arXiv:2109.08238*.
- Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv preprint arXiv:2402.07927*.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*.
- V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihang Wang, Yan Wang, Yean Cheng, Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, and 69 others. 2025. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *Preprint*, arXiv:2507.01006.
- Heng Wang, Jianhui Yu, and Weidong Cai. 2022. Spacap3d: Spatiality-guided transformer for 3d dense captioning on point clouds. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Jianren Yang, Yikang Wang, Yang Zhao, Ziang Zhang, and Zhou Zhao. 2023. Leo: An embodied generalist agent in 3d world. *arXiv preprint arXiv:2311.12871*.
- Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. 2019. Deep modular co-attention networks for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sha Zhang, Di Huang, Jiajun Deng, Shixiang Tang, Wanli Ouyang, Tong He, and Yanyong Zhang. 2024a. Agent3d-zero: An agent for zero-shot 3d understanding. In *European Conference on Computer Vision*, pages 186–202. Springer.
- Yifan Zhang, Yikang Wang, Yang Zhao, Ziang Zhang, and Zhou Zhao. 2024b. Chatscene: Knowledge-enabled safety-critical scenario generation for autonomous vehicles. *arXiv preprint arXiv:2405.14062*.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024c. Video instruction tuning with synthetic data. *Preprint*, arXiv:2410.02713.
- Duo Zheng, Shijia Huang, Lin Zhao, Yiwu Zhong, and Liwei Wang. 2024. Towards learning a generalist model for embodied navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chenming Zhu, Tai Wang, Wenwei Zhang, Jiangmiao Pang, and Xihui Liu. 2024. Llava-3d: A simple yet effective pathway to empowering llms with 3d-awareness. *arXiv preprint arXiv:2409.18125*.
- Ziyu Zhu, Xiaojian Ma, Yixin Chen, Zhidong Deng, Siyuan Huang, and Qing Li. 2023. 3d-vista: Pre-trained transformer for 3d vision and text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2911–2921.

You are a helpful assistant that can answer questions about 3D scenes.

You will be given images from a 3D scene and a question about that scene.

Your goal is to find the answer to the question based on the images of the scene.

The question is: {{ question }}

The images will be provided by the user below.

Figure 7: Prompt template for baseline.

Given a question about a 3D scene and a set of available images from different camera angles, select viewpoint IDs that provide the best visual evidence to answer the question.

- Question: {{ question }}

- Available view IDs: {{ view_ids }}

1. Target Localization: Identify the key objects or regions mentioned in the question.

2. View Selection Logic:

- For Object Properties: Choose close-up views or views with the clearest line-of-sight (minimal occlusion).

- For Spatial Relationships: Choose views or multiple views that capture all relevant objects.

- For Occluded Objects: Choose views from opposing angles to "see behind" obstacles.

3. Redundancy Reduction: Avoid selecting multiple views that offer nearly identical perspectives; prioritize diversity in angles.

Output Format

selected views: A, B, C, ...;

(Note: Replace A, B, C with the actual numeric or string IDs from the list)

Figure 8: Prompt template for coarse-grained view selection agent.

A Prompt Templates

The prompt templates used in our experiments are shown in Figures 7 to 9. For the baseline, we directly provide all view frames along with the question and let the model generate an answer. For the view selection agent, we input all view frames, the question, and an index for each frame, and the model is required to output the IDs of the views relevant to the question. For the CoV agent, we input the coarsely selected views and the question, and the model produces the answer after multi-step action execution.

Given a question and multiple viewpoints (including a bird-eye view), you must actively explore the scene by switching views and adjusting the camera to find the correct answer.

You are REQUIRED to execute at least {{ min_action_steps }} action steps before answering.

You may perform ONLY ONE action per step.

Available actions:

View switch: switch to view N, switch to bird-eye-view

Movement: forward-movement+N, backward, left, right, upward and downward.

Rotation: left-rotation+N, right-rotation+N

Before answering, you MUST enter a verification phase and state:

"I'm now verifying my answer by..."

and use camera actions to satisfy any remaining steps.

Output format

done+[your answer]

Question: {{ question }}

Available view IDs: {{ view_ids }}

Figure 9: Prompt template for fine-grained CoV agent.

B Result Visualization

The qualitative results illustrate CoV's ability to accurately localize objects, reason about spatial relationships, and align its predictions with the semantics of the question. Across diverse indoor scenes—bathroom 10, classroom 11, office 12, and kitchen 13, CoV consistently identifies relevant details, even when targets are partially occluded or distributed across views. These examples demonstrate CoV's strength in multi-step, viewpoint-aware reasoning.




Figure 10: CoV identifies the soap and paper towel dispensers between two sinks and reasons about material and spatial attributes such as the metal trash can and wall-mounted sink.




Figure 11: CoV accurately localizes furniture such as brown cushion sofas, chairs beside a table, and a tall bookshelf positioned near the chalkboard, demonstrating spatial alignment and object identification.




Figure 12: Given a desk scene, CoV correctly identifies multiple items on the desk and reasons about their positions, such as the shelf to the right of the monitor and office chairs nearby.

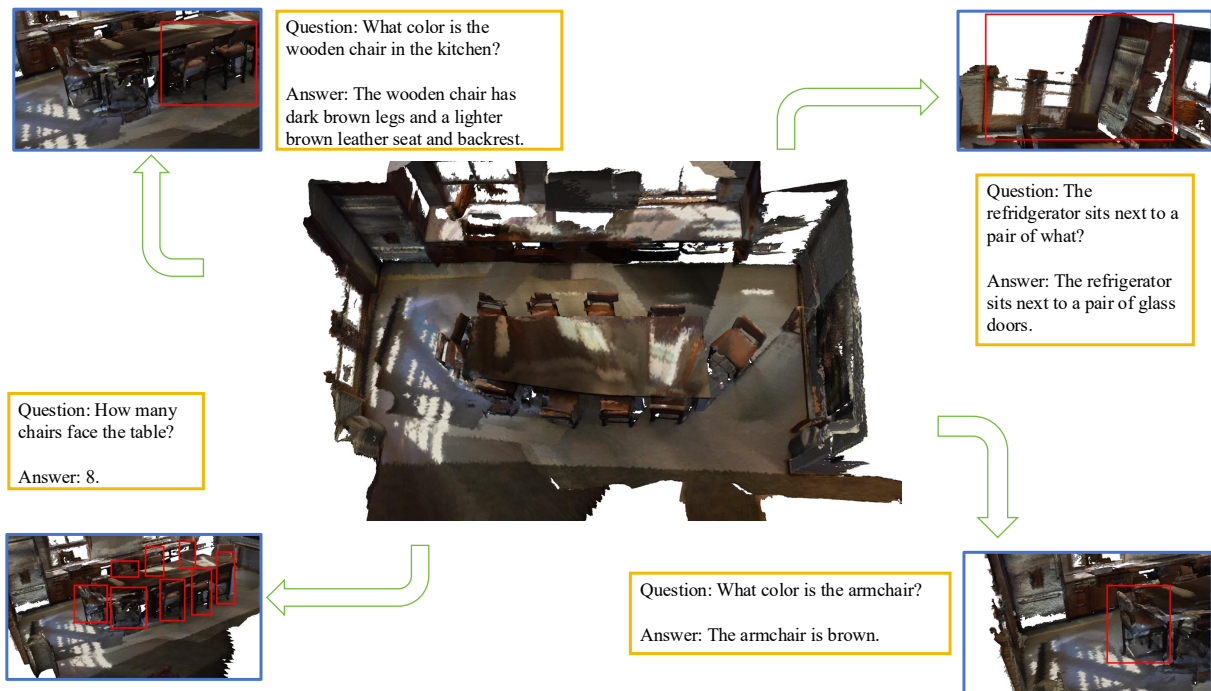


Figure 13: CoV answers layout and appearance questions, including identifying brown-toned wooden chairs, the refrigerator's proximity to glass doors, and counting the number of chairs facing the table.