Toward Ambulatory Vision: Learning Visually-Grounded Active View Selection

Juil Koo* Daehyeon Choi* Sangwoo Youn* Phillip Y. Lee Minhyuk Sung
KAIST
{63days,daehyeonchoi,andy2884,phillip0701,mhsung}@kaist.ac.kr

Project Page: https://active-view-selection.github.io

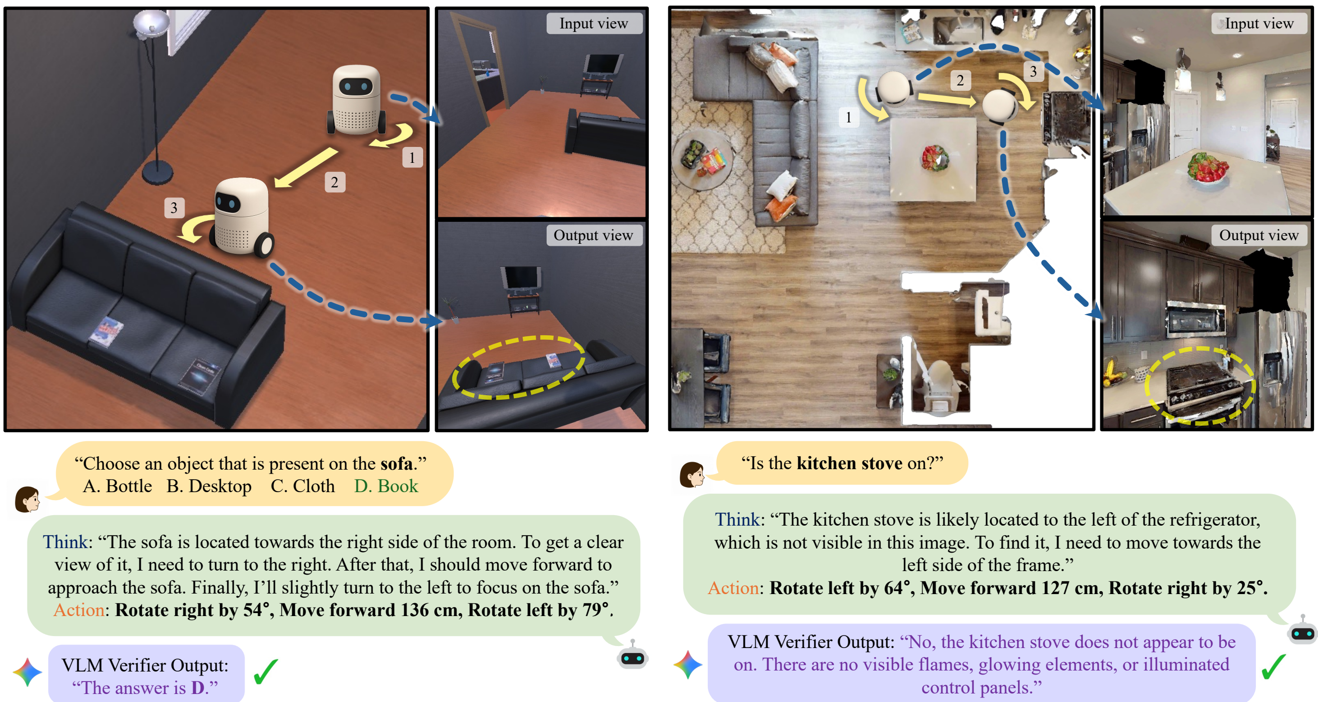

Figure 1. **An overview of our proposed Visually-Grounded Active View Selection (VG-AVS).** Given a 3D environment from synthetic (left) to real scenes (right) and a question, our learning-based active view selection (AVS) framework predicts continuous actions to refine the agent's viewpoint. The refined view is then fed into a VLM verifier, which answers the question based on the improved observation.

Abstract

Vision Language Models (VLMs) excel at visual question answering (VQA) but remain limited to snapshot vision, reasoning from static images. In contrast, embodied agents require ambulatory vision, actively moving to obtain more informative views. We introduce Visually Grounded Active View Selection (VG-AVS), a task that selects the most informative next viewpoint using only the visual information in the current image, without relying on scene memory or external knowledge. To support this task, we construct a synthetic dataset with automatically generated paired query–target views and question–answer prompts. We also propose a framework that fine-tunes pretrained VLMs through supervised fine-tuning (SFT) followed by RL-based policy optimization. Our approach achieves strong question answering performance based on viewpoint selection and generalizes robustly to unseen synthetic and real scenes. Furthermore, incorporating our learned VG-AVS framework into existing scene-exploration-based EQA systems improves downstream question-answering accuracy.

1. Introduction

> *"We look around, walk up to something interesting and move around it so as to see it from all sides, and go from one vista to another. That is natural vision."*
>
> JAMES J. GIBSON [10]

*Equal contribution.

We are now witnessing the rise of vision language models (VLMs) that have effectively mastered visual question answering (VQA), seemingly closing the chapter on this line of research. What, then, should come next?

To address this question, we revisit the fundamental purpose of the visual perception system. James J. Gibson, one of the most influential perceptual psychologists of the twentieth century, offered foundational insights through his *ecological* approach to perception [10]. According to Gibson, current VLMs operate within what he described as *snapshot vision*, the ability to interpret a single static retinal image. Much of early computer vision research similarly focused on this static image interpretation. Yet Gibson emphasized that the real problem of vision in animals extends far beyond a fixed view. As he wrote, "natural vision depends on the eyes in the head on a body supported by the ground". Vision in the real world requires looking around and moving toward objects of interest, an ability he referred to as *ambulatory vision*.

In this work, we study active perception for VQA by assuming an embodied agent situated in a scene and training it to select the most informative viewpoint for answering a question. Unlike prior approaches, which rely on memorized scene representations [32, 42] or external commonsense knowledge [21, 31], we develop a system that operates solely on the visual information available in the current observation. We refer to this task as Visually-Grounded Active View Selection (VG-AVS).

Active perception has previously been studied under the Embodied Question Answering (EQA) framework [6]. EQA typically involves four components: (1) scene exploration, (2) scene memorization, (3) commonsense reasoning, and (4) localization and perception. Most recent work [21, 31, 32, 42] focuses on exploration, memorization, and reasoning, corresponding respectively to planning, representation learning, and language understanding. The fourth component—localization and perception—remains less explored, despite being the core computer-vision challenge. Our study focuses precisely on this aspect and provides a concrete setup for learning active visual perception.

Compared with previous efforts, our formulation advances active perception in three key aspects. (1) We aim to learn the optimal viewpoint rather than only the agent's position in the scene [21, 31]. To achieve this, we consider the full set of control parameters for a mobile agent: heading rotation, forward translation, and view rotation. (2) We pursue fine-grained, continuous control of these parameters instead of coarsely discretizing them into actions such as turn-left, turn-right or move-forward. [41]. Our model predicts continuous rotation angles and translation distances, enabling precise viewpoint adjustment. This formulation avoids the complex multi-turn navigation settings used in prior EQA methods [21, 31, 32, 42], which makes learning continuous, high-resolution control policies feasible. (3) Most importantly, unlike previous zero-shot EQA methods [21, 31, 32, 42], we propose a fine-tuning-based active perception framework and demonstrate its strong generalization to unseen environments and diverse question types.

For training, as one of our key contributions, we introduce a curated synthetic dataset, called AVS dataset. Built on ProcTHOR [7], each sample consists of a rendered query view and target view. The query view partially includes objects visible in the target view while omitting others, simulating incomplete visual observations. A corresponding question prompt asks about the missing information, requiring the agent to infer the appropriate mobility parameters needed to reach the target viewpoint.

We fine-tune pretrained VLMs using two complementary strategies: supervised fine-tuning (SFT) to learn ground-truth transformations, and reinforcement learning (RL)-based unsupervised fine-tuning, where the model predicts mobility parameters at the end of its reasoning process without explicit supervision. Combining the two in a sequential SFT-then-RL scheme first grounds the model through supervision and then refines it via unsupervised policy optimization, improving both performance and generalization.

Our experiments show that our training-based methods significantly outperform zero-shot approaches in VG-AVS, confirming the benefit of learning active visual perception beyond pretrained VLM capabilities. Despite being trained only on a small-scale synthetic dataset, our method generalizes to real scenes with diverse question types and can serve as a plug-and-play component that improves existing EQA frameworks.

2. Related Work

Active Visual Question Answering. Beyond solving VQA from static images, several works explore active visual question answering (Active VQA), where the model interacts with visual inputs before answering. PixelReasoner [35], ToA [18], and Directional Guidance [20] all operate in the 2D image space, performing actions such as cropping, selecting key frames, or predicting coarse directions toward regions of interest to obtain more informative views. However, these methods operate strictly in the 2D image space with limited action types, remaining confined to the given frame rather than exploring new viewpoints in the underlying 3D scene.

MindJourney [41] aims to enhance the spatial reasoning of VLMs by generating new observations with a generative model, but its action space is restricted to discrete primitives and actions are selected via beam search rather than by learning a policy. In contrast, we directly learn a policy for an active perception system in physically grounded 3D scenes, operating in a fine-grained, continuous action space.

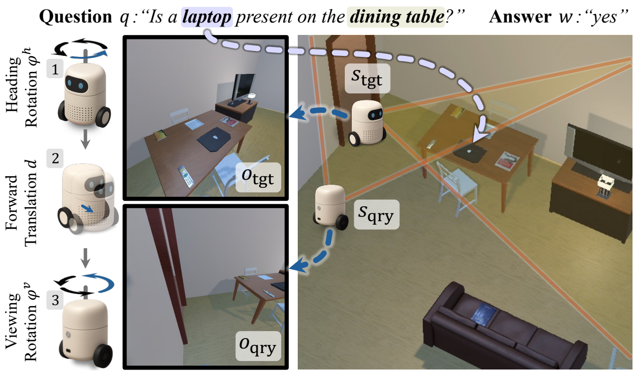

Figure 2. **Our action sequence (left) and AVS dataset sample (right).** The action controls how the agent adjusts its viewpoint. On the right, the target object (*laptop*) is visible only in the target view, while the query view shows only the supporting object (*dining table*), which may serve as a visual clue that motivates active view selection to gather sufficient visual evidence for answering.

Embodied Question Answering. There has been substantial effort to tackle embodied question answering (EQA), both in constructing datasets [6, 21, 24] and in designing sophisticated methods [21, 31, 32, 42]. Since EQA requires multiple abilities, such as large-scale scene exploration, scene memorization, common sense reasoning, and active perception, prior work has mainly focused either on efficient scene exploration [21, 31] or on better 3D scene representations for memory construction [32, 42]. In contrast, the crucial link between partial visual observations and fine-grained view selection, that is, how an agent should refine its viewpoint based on existing visual clues, remains underexplored. Our work explicitly targets this computer vision challenge and can be adopted as a plug-in module that fills this missing component in EQA, leading to improved overall performance.

Post-Training in VLMs. Motivated by the strong success of post-training in the LLM literature [5, 27, 28, 33, 38, 39], many works have focused on post-training vision-language models (VLMs) [1, 12, 17, 34, 36, 37] to enhance their capabilities, including spatial understanding [2–4, 22, 23, 40] and multi-view or 3D understanding [8, 9, 47]. In particular, GRPO-based reinforcement learning approaches [33, 43, 46] have demonstrated that post-training can improve model performance by optimizing over the model's own reasoning process without requiring additional human annotations or explicit supervision. Another line of work fine-tunes VLMs to endow them with action-oriented decision-making abilities [13, 15, 16, 19, 45]. However, none of these works seeks to improve active perception, the ability to compensate for missing visual evidence in the current view.

3. Visually-Grounded Active View Selection

We study the problem of learning an *active perception system* that adjusts its egocentric viewpoint to acquire more informative observations for answering a language query.

Beyond passive answering from static images, our objective is to cultivate an active perception ability, namely the capacity to recognize when the current view lacks sufficient visual evidence, identify informative cues from the current visual-linguistic context, and actively adjust the viewpoint to acquire the missing information. Through this process, the model learns to perform question-driven view selection in a learning-based manner, ultimately enhancing its visual grounding and improving VQA performance.

Unlike prior works on Active VQA [18, 20, 22], which typically define the action space in the 2D image domain and apply simple visual operators such as cropping or zooming to focus on salient regions, we consider more spatially grounded scenarios where the agent is situated in a 3D environment. Our formulation enables the agent to refine its viewpoint through explicit locomotion and viewpoint rotation, demanding fine-grained control within a substantially larger action space.

To realize this active perception capability, we introduce three key components: a dedicated training dataset in Section 3.1, a systematically designed framework for continuous active view selection in Section 3.2, and an effective two-stage training strategy in Section 3.2.3.

3.1. AVS Dataset

To enable such visually-grounded active view selection, it is essential to construct training data that explicitly supports this capability. To this end, we curate a dataset, called AVS, designed to train models that actively adjust their viewpoints to gather sufficient visual evidence for answering a given question.

Each sample in the dataset is represented as a tuple $(q, w, s_{\text{tgt}}, s_{\text{qry}}, o_{\text{tgt}}, o_{\text{qry}})$, where q denotes the language query, w its corresponding answer, s_{tgt} and s_{qry} indicate the target and query camera poses, and o_{tgt} and o_{qry} are their corresponding rendered observations.

The target view o_{tgt} contains sufficient visual information to answer the query, while the query view o_{qry} provides only partial visual evidence, encouraging the model to actively adjust its viewpoint to reveal the missing information. To capture this contextual difference between the two views within each training sample, we build the AVS dataset using the ProcTHOR [7] environment, which provides richly annotated indoor 3D scenes with instance-level object labels and explicit surface–object relations, where certain objects are designated as *supporting* objects that serve as surfaces for other assets, such as countertops holding cups or beds supporting pillows, and these supporting objects could offer visual cues about the unseen *target* object referred to by the question.

Below is an overview of our automatic data curation pipeline. We first select an object of interest and generate a corresponding question–answer pair (q, w) referring to that object. We then sample an answerable target view o_{tgt} and a contextual (query) view o_{qry}. To determine these views, we leverage the instance labels to render each viewpoint together with its instance segmentation mask, which provides pixel-level visibility information for every object in the scene.

Let x_{tgt} and x_{sup} denote the target object of interest and its supporting object, respectively. We use $N_p(o, x)$ to represent the number of pixels in view o that belong to object x, and $c(o, x)$ to denote the normalized distance between the projected centroid of x and the image center. We define three thresholds $\epsilon_{\text{vis}}^{\text{sup}} > \epsilon_{\text{vis}}^{\text{obj}} > \epsilon_{\text{inv}}^{\text{obj}}$, corresponding to the visibility thresholds for the supporting and target objects, and the invisibility threshold for the target object, respectively. Then, the target and query views are defined as:

$$o_{\text{tgt}} \sim \{o \mid N_p(o, x_{\text{tgt}}) > \epsilon_{\text{vis}}^{\text{obj}},\ c(o, x_{\text{tgt}}) < \delta_{\text{center}}\} \quad (1)$$

$$o_{\text{qry}} \sim \{o \mid N_p(o, x_{\text{tgt}}) < \epsilon_{\text{inv}}^{\text{obj}},\ N_p(o, x_{\text{sup}}) > \epsilon_{\text{vis}}^{\text{sup}}\} \quad (2)$$

where δ_{center} specifies the maximum allowed distance for the target to be regarded as centered. Given the selected target and supporting objects, we construct the corresponding question–answer pair by instantiating a predefined question template. The illustration of the dataset example is in the right of Figure 2, and further implementation details are provided in the **supplementary material**.

The resulting training set consists of 1,320 scenes with 1,867 tuple samples focused solely on binary *existence* questions.

For evaluation, we construct two additional benchmark datasets that cover more diverse question types and real indoor scenes. AVS-ProcTHOR extends beyond the binary object existence questions used for training to include three question types: *existence*, *counting*, and *state*. Furthermore, AVS-HM3D is built from real-world indoor scenes in the Habitat–Matterport 3D dataset [29]. Details of these benchmarks are presented in Section 4.

3.2. AVS Framework

As discussed in Section 3, we cast active view selection as learning a policy that moves the agent from a contextual viewpoint toward the *answerable observation space*.

Formally, we model this as a continuous decision-making problem, where the agent predicts a real-valued action that adjusts its egocentric viewpoint to acquire sufficient visual evidence for answering a given question. Let the state $s \in \mathcal{S}$ represent the agent's 3D position and orientation in the 3D environment, and the corresponding observation $o = \Omega(s)$ be the egocentric RGB view obtained from s by the observation function $\Omega(\cdot)$. The policy $\pi_\theta(\cdot|o, q)$ outputs an action $a \in \mathcal{A}$, such as rotation angles and translation distances, that determines how the agent should move to acquire sufficient visual evidence for answering the language query q. Executing a in the environment updates the state via $s' = \mathcal{T}(s, a)$ and produces a new observation $o' = \Omega(s')$. This continuous formulation enables the agent to directly reach an answerable viewpoint with a single fine-grained action, rather than relying on multiple discrete steps, thereby simplifying training from a multi-step process to a single-step policy optimization. We next detail the design of our action space below.

Action Space Design. We represent the state s as a triplet $s = (x, y, \varphi)$, where (x, y) denotes the agent's 2D position and φ is the azimuth orientation. We fix the height and elevation angle for simplicity. The action space is parameterized to compactly span all physically executable movements of the embodied agent using a minimal set of continuous components, resulting in a triplet $a = (\varphi^h, d, \varphi^v)$:

- **Heading rotation** $\varphi^h \in (-180°, 180°]$: azimuthal turning angle that determines the moving direction relative to the agent's current orientation φ.
- **Forward translation** $d \geq 0$: distance to move forward along the rotated heading.
- **View rotation** $\varphi^v \in (-180°, 180°]$: final azimuthal offset of the agent relative to the moving direction, modeling a fine-grained head turning.

All azimuthal angles follow the same convention, where positive angles indicate right turns.

The action is performed sequentially: the agent first determines its heading direction by φ^h, then moves forward by d along that heading, and finally adjusts its viewing direction by applying the view rotation φ^v to the current azimuth orientation. Accordingly, the state transition is given by:

$$s' = \begin{pmatrix} x' \\ y' \\ \varphi' \end{pmatrix} = \mathcal{T}(s, a) = \begin{pmatrix} x + d\sin(\varphi + \varphi^h) \\ y + d\cos(\varphi + \varphi^h) \\ \varphi + \varphi^h + \varphi^v \end{pmatrix}. \quad (3)$$

Given two states, specifically $s' = s_{\text{tgt}}$ and $s = s_{\text{qry}}$, the ground-truth action a_{tgt} that moves the agent from the query state s_{qry} to the target state s_{tgt} can be analytically computed as follows:

$$a_{\text{tgt}} = \begin{pmatrix} \varphi_{\text{tgt}}^h \\ d_{\text{tgt}} \\ \varphi_{\text{tgt}}^v \end{pmatrix} = \begin{pmatrix} \operatorname{atan2}(\Delta x, \Delta y) - \varphi \\ \sqrt{(\Delta x)^2 + (\Delta y)^2} \\ \varphi' - (\varphi + \varphi_{\text{tgt}}^h) \end{pmatrix}, \quad (4)$$

where $\Delta x = x' - x$ and $\Delta y = y' - y$ denote the positional differences. See the left of Figure 2 for an illustration of action space design.

3.2.1. Supervised Fine-Tuning

Under the problem formulation discussed in Section 3.2, the most straightforward approach is to adopt a teacher-forcing objective for learning the policy, where the model is supervised to predict the ground-truth action given the input observation and language query.

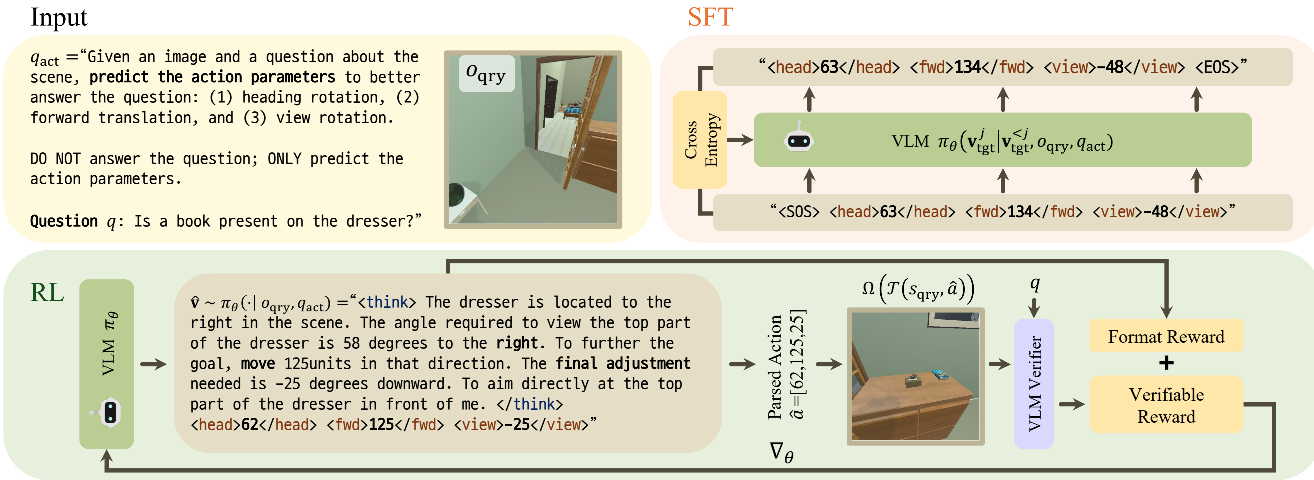

Figure 3. **Overview of training strategies.** For a given query image o_{qry} and question q_{act}, the action model is trained to predict action to obtain next view. In SFT, the model is trained to directly predict ground-truth actions from the input. In contrast, under RL, the model first generates a token sequence that includes its own reasoning process, followed by the final action prediction, and the model is optimized so that outputs leading to higher rewards become more likely.

Note that such supervised fine-tuning is feasible in our setup since our synthetic dataset provides access to ground-truth target views, enabling an analytic computation of the ground-truth action between paired query and target viewpoints. In contrast, existing EQA setups lack such view-level supervision, as ground-truth actions are difficult to define or annotate due to long-horizon navigation trajectories.

Specifically, we parameterize the policy with a VLM. To enable the VLM to predict real-valued actions, we introduce a function $\texttt{str}(\cdot)$ that maps real-valued numbers to their string representations. Given an action $a = (\varphi^h, d, \varphi^v)$, we convert it into the following formatted string representation: $\mathbf{v}$ =<H>str(φ^h)</H> <D>str(d)</D> <V>str(φ^v)</V>, where <H>, <D>, and <V> are special tags to explicitly delimit each scalar component.

Given a tuple $(q, w, s_{\text{tgt}}, s_{\text{qry}}, o_{\text{tgt}}, o_{\text{qry}})$ sampled from our AVS dataset, we first compute the ground-truth target action a_{tgt} following Equation 4, which specifies how the agent should move from the query viewpoint to the target one. The teacher-forcing objective for SFT is expressed as follows:

$$\mathcal{L}_{\text{SFT}} = \sum_{j=1}^{L} \log \pi_\theta(\mathbf{v}^j_{\text{tgt}} \mid \mathbf{v}^{<j}_{\text{tgt}}, o_{\text{qry}}, q_{\text{act}}), \tag{5}$$

where $\mathbf{v}_{\text{tgt}}$ denotes the string representation of a_{tgt}, and q_{act} is an action-specific instruction expanded from q, which explicitly prompts the model to predict the action rather than directly answering the question.

While SFT enables the model to learn which action to predict from a multimodal context, it confines the VLM to human-annotated actions.

3.2.2. Reinforcement Learning

Another direction for learning the policy is to employ reinforcement learning (RL), which leverages the model's internal chain-of-thought reasoning process to arrive at the final action prediction without explicit supervision. Once the model generates a token sequence through its reasoning process, denoted as $\hat{\mathbf{v}}$ =<think>...</think> <H>h</H> <D>d</D> <V>v</V>, we denote the executable real-valued action parsed from this output as $\hat{a}$. The policy is optimized to maximize rewards computed from the generated sequence. We employ a verifiable reward r^{ver} that converts each predicted view into a binary feedback signal, together with a format reward r^{fmt} that encourages the model to produce a correctly formatted action string.

Concretely, the verifiable reward r^{ver} is measured using a frozen pre-trained VLM v_ϕ that acts as an external verifier and checks whether the question can be correctly answered from the predicted view, defined as:

$$r^{\text{ver}}(o, q, w) = \begin{cases} 1, & \text{if } v_\phi(o, q) = w, \\ 0, & \text{otherwise}, \end{cases} \tag{6}$$

where (q, w) is a question-answer pair and o denotes the input observation.

Formally, the RL objective maximizes the expected reward over sampled tokens:

$$\max_\theta \mathbb{E}_{\hat{\mathbf{v}} \sim \pi_\theta(\cdot | o_{\text{qry}}, q_{\text{act}})} \left[r^{\text{fmt}}(\hat{\mathbf{v}}) + r^{\text{ver}}(\hat{o}, q, w) \right], \tag{7}$$

where $\hat{o} = \Omega(\mathcal{T}(s_{\text{qry}}, \hat{a}))$ with $\hat{a}$ denoting the real-valued action parsed from the sampled token sequence $\hat{\mathbf{v}}$. We use Group Relative Policy Optimization (GRPO) [33] for policy optimization.

Table 1. **Quantitative comparison on our AVS benchmark.** We report VQA accuracy on AVS-ProcTHOR and LLM-Match scores on AVS-HM3D, normalized to a percentage scale. The best except for *'No Action'* in each column is in **bold** and second best is underlined.

Action Model		AVS-ProcTHOR [7]				AVS-HM3D [29]					
		Existence	Counting	State	Average	Existence	Counting	State	Attribute	Object	Average
No Action	Query view	49.22	16.36	61.57	42.38	67.50	56.15	54.59	66.67	48.33	58.65
	Target view	93.02	69.14	92.58	84.91	86.67	77.50	80.00	74.48	71.79	78.09
Backbone Model	Qwen2.5-VL-7B [1]	64.34	29.74	56.55	50.21	66.25	46.15	66.49	49.17	46.86	54.98
Spatial VLMs	ViLaSR [40]	57.95	25.46	52.84	45.42	68.33	48.46	52.70	50.00	50.00	53.90
	SpatialReasoner [22]	54.65	22.68	52.62	43.32	70.42	50.00	47.30	37.50	40.00	49.04
EQA Framework	Fine-EQA [21]	63.57	31.97	64.41	53.32	70.00	52.31	52.70	54.17	44.44	54.72
Proprietary Models	GPT-5 [26]	81.01	55.58	79.69	72.09	76.67	54.62	65.95	65.83	60.00	64.91
	Gemini-2.5-Pro [11]	82.95	52.79	81.00	72.25	74.17	59.23	60.81	64.17	60.00	63.67
AVS Framework	SFT	91.28	57.06	83.84	77.39	67.50	70.77	62.16	66.67	55.56	64.53
(Backbone:	RL	86.82	65.24	83.41	78.49	**81.25**	70.00	72.97	69.17	60.00	70.68
Qwen2.5-VL-7B [1])	**SFT+RL (Ours)**	**91.47**	**69.52**	**90.17**	**83.72**	74.58	**71.54**	**73.78**	**70.83**	**62.78**	**70.70**

3.2.3. Bridging SFT and RL

While SFT enables the model to efficiently learn plausible actions from visual–linguistic inputs through explicit supervision from paired query–target viewpoints in our dataset, we empirically observe that SFT alone quickly saturates and provides limited further gains once the basic mapping from observations to actions is learned.

On the other hand, when the model is trained with RL from scratch, the resulting policy typically underperforms the SFT-trained model and is less stable in the continuous action space, even with carefully designed rewards.

Empirically, we find that combining the two objectives in a staged manner is crucial: warming up the policy with SFT provides a good initialization that captures plausible action magnitudes and directions, and subsequent fine-tuning with RL brings additional improvements by refining these actions under task-specific rewards through the model's own reasoning process. This two-stage training strategy yields better performance than using either SFT or RL alone.

4. Experiments

See the **supplementary material** for training details, extended comparisons, and cross-dataset generalization.

4.1. Experiments on VG-AVS

Experiment Setups. For evaluation on the Visually-Grounded Active View Selection (VG-AVS) task, we use two benchmark datasets: AVS-ProcTHOR and AVS-HM3D. AVS-ProcTHOR is automatically generated following the data curation pipeline introduced in Section 3.1, but includes more diverse question types than the training set, covering 516, 538, and 458 samples for *existence*, *counting*, and *state*, respectively.

AVS-HM3D consists of real indoor scenes constructed from triplets of (question, answer, ground-truth view) in the validation split of Fine-EQA [21]. Due to the absence of object visibility checks in real scenes, we first generate query views by randomly perturbing the ground-truth viewpoints so that the target object becomes invisible, and then manually collect 208 samples where the query view still provides sufficient contextual visual clues. AVS-HM3D spans five question types: *existence*, *counting*, *state*, *attribute*, and *object*. See the **supplementary material** for more details on the data construction.

As evaluation metrics, we report VQA accuracy on AVS-ProcTHOR, where all questions are in multiple-choice format so accuracy can be computed directly. For AVS-HM3D, which consists of open-ended questions, we use LLM-Match [24], an LLM-based correctness metric where an LLM assigns a discrete score from 1 to 5 by comparing the predicted answer with the annotated ground-truth answer. We use Gemini-2.5-Flash [11] for VQA accuracy and GPT-4 [25] for LLM-Match following prior work [24].

Quantitative Results. We present a quantitative comparison in Table 1. To provide loose reference upper and lower bounds that reflect the difficulty of the task, we report the VQA accuracy and LLM-Match scores when the VLM verifier is directly given the query view and the target view, respectively. For all action models, we instead first predict actions and feed the resulting predicted views to the VLM verifier to compute the metrics.

As shown, open-source models, including Qwen2.5-VL-7B [1] and recent spatial VLMs such as ViLaSR [40] and SpatialReasoner [22], all fail to achieve meaningful active view selection. In contrast, our AVS framework, which also uses Qwen2.5-VL-7B [1] as its backbone, consistently outperforms these models by large margins across different training strategies. This highlights that off-the-shelf open VLMs are insufficient to perform informative viewpoint changes from partial observations, and that training on our curated dataset is crucial for acquiring such active perception ability.

Figure 4. **Qualitative results in AVS-ProcTHOR (top two rows) and AVS-HM3D (bottom three rows).** Blue denotes the object of interest, and gray denotes surrounding objects that may serve as visual clues for reasoning. ✓ corresponds to correct answers (or LLM-Match scores $= 5$), ✗ corresponds to wrong answers (or LLM-Match scores ≤ 2). The red arrow in the final column marks the region of interest.

In our AVS framework, across different training strategies, all variants outperform other approaches, while the two-stage training scheme (SFT+RL) yields a further improvement of more than 5% in VQA accuracy compared to either method alone. Moreover, they show strong generalizability to diverse question types that are not seen during training, including *counting* and *state* questions. Additional training strategy variants and ablations are provided in the **supplementary material**.

Fine-EQA [21], an EQA framework proposed for a task closely related to ours, performs significantly worse on our benchmark. This indicates that EQA frameworks aiming for long-horizon navigation lack fine-grained viewpoint control, as their policies rely on coarse, discrete actions instead of continuous, precise adjustments.

When comparing against proprietary models, the advantage of our learned active perception system becomes even more pronounced. Despite using only a 7B-parameter backbone [1], our model surpasses substantially larger proprietary models, including GPT-5 [26] and Gemini-2.5-Pro [11]. This clearly highlights that active perception for ambulatory vision requires an explicit training procedure for precise view refinement, rather than relying solely on large-scale pretraining.

Beyond the synthetic scenes from ProcTHOR [7], the results on AVS-HM3D reported in Table 1 clearly demonstrate that our AVS framework generalizes well to real indoor environments with diverse question types. Despite be-

Table 2. **Comparison on Fine-EQA benchmark.** Each column reports normalized LLM-Match score for different question types. Plugging our AVS framework into Fine-EQA improves performance. In each coolumn, best in **bold**, second best <u>underlined</u>.

Method	Attr.	Count.	Exist.	Obj.	State	Loc.	Avg.
Fine-EQA [21]	49.23	38.75	67.92	**57.41**	<u>63.20</u>	41.14	52.94
w/ SFT	**59.08**	<u>45.00</u>	67.29	<u>56.30</u>	58.40	<u>45.71</u>	<u>55.30</u>
w/ RL	52.00	44.38	**69.17**	47.41	62.00	42.29	52.87
w/ Ours	<u>53.23</u>	**52.50**	<u>68.13</u>	55.93	**65.40**	**50.86**	**57.67**

ing trained only on a relatively small-scale synthetic 3D dataset with binary object existence questions, our approach outperforms all other baselines by clear margins, achieving 70.70 with SFT+RL compared to 64.91 with GPT-5 [26] in average score. This shows that a well-curated dataset that compactly supervises active view selection is sufficient to enable strong transfer to real-world scenes, even under shifts in both scene distribution and question format. Moreover, since our AVS module is orthogonal to the choice of backbone, these gains suggest that even stronger VLMs could further benefit from being equipped with our learned active perception system.

Qualitative Results. As shown in Figure 4, when the target object in the query view is only partially visible or appears too small, our AVS framework successfully refines the viewpoint to make the object fully observable and properly scaled. In the first row, for example, when the dining table is only partially visible in the query view, our framework moves the agent closer to the table so that all mug cups become fully observable. In the fourth row, we show a real indoor scene where the closet is only partially visible in the query view. Our method adjusts the viewpoint to fully reveal the closet, enabling the VLM verifier to answer correctly, demonstrating the generalizability of ours to real-world environments.

4.2. Experiments on EQA

Our AVS framework can naturally serve as a plug-and-play component to enhance existing EQA pipelines. Since EQA inherently requires multi-task abilities, prior work [6, 21, 24, 31] has largely underexplored the final step of fine-grained view refinement once the agent stops near the target, which is crucial for gathering the visual evidence needed to answer the question. Our AVS framework directly complements this stage: it takes the terminal view produced by the EQA policy, refines the viewpoint, and feeds the refined observation into the VLM verifier for answering.

As shown in Table 2, this simple plug-in yields clear LLM-Match improvements over the base Fine-EQA [21] framework, increasing the average score from 52.94 to 57.67 when combined with our AVS framework trained using SFT+RL, denoted as (w/ Ours) in the table.

Beyond demonstrating generalization to real scenes in VG-AVS task, this result also shows that our AVS module transfers well to EQA setups with real-world environments. This suggests that our method can serve as a complementary component that strengthens existing EQA frameworks.

Figure 5. **Qualitative results in the Fine-EQA benchmark [21].** Blue highlights denote object-of-interest. ✓ corresponds to high scores (= 5) and ✗ corresponds to low LLM-Match scores (≤ 2). The red arrow in the final column marks the region of interest.

As shown in Figure 5, when the final view produced by the base EQA framework is insufficient for answering a question, our module successfully identifies the region of interest and adjusts the viewpoint to obtain a more informative observation. For example, in the first row, our model moves the agent inside the room to better capture the laundry room mentioned in the question. Comparing different training variants, in the second row, SFT and RL produce views where the left side of the bar counter is only partially visible, making the chair count less clear, whereas our two-stage training method moves to a viewpoint where the entire counter is clearly captured.

5. Conclusion

We advance toward ambulatory vision by reframing VQA as an active perception problem and introducing Visually Grounded Active View Selection (VG-AVS). By focusing on viewpoint selection from a single image and enabling VLM fine-tuning through a novel synthetic dataset and an SFT+RL fine-tuning framework, our approach achieves significant gains in view-selection-based question answering. It also generalizes well to unseen scenes and provides meaningful improvements when incorporated into existing scene-exploration-based EQA pipelines.

References

[1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 3, 6, 7, 11, 13, 16

[2] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brian Ichter, Danny Driess, Pete Florence, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. SpatialVLM: Endowing vision-language models with spatial reasoning capabilities. In *CVPR*, 2024. 3

[3] Jiawei Chen, Yiran Wang, Yucheng Xu, Chang Liu, Xizhou Li, Jifeng Wang, and Jiwen Lu. Spatial-MLLM: Enhancing spatial reasoning in multi-modal large language models. In *NeurIPS*, 2025.

[4] An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. SpatialRGPT: Grounded spatial reasoning in vision-language models. In *NeurIPS*, 2024. 3

[5] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. In *JMLR*, 2024. 3

[6] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *CVPR*, 2018. 2, 3, 8

[7] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Jordi Salvador, Kiana Ehsani, Winson Han, Eric Kolve, Ali Farhadi, Aniruddha Kembhavi, and Roozbeh Mottaghi. ProcTHOR: Large-Scale Embodied AI Using Procedural Generation. In *NeurIPS*, 2022. 2, 3, 6, 7, 12, 13

[8] Jiajun Deng, Tianyu He, Li Jiang, Tianyu Wang, Feras Dayoub, and Ian Reid. 3d-llava: Towards generalist 3d lmms with omni superpoint transformer. In *CVPR*, pages 3772–3782, 2025. 3

[9] Mohsen Gholami, Ahmad Rezaei, Zhou Weimin, Yong Zhang, and Mohammad Akbari. Spatial reasoning with vision-language models in ego-centric multi-view scenes. *arXiv preprint arXiv:2509.06266*, 2025. 3

[10] James J. Gibson. *The Ecological Approach to Visual Perception*. Houghton Mifflin, 1979. 1, 2

[11] Google. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities, 2025. 6, 7

[12] Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, et al. Glm-4.1 v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv e-prints*, 2025. 3

[13] Chi-Pin Huang, Yueh-Hua Wu, Min-Hung Chen, Yu-Chiang Frank Wang, and Fu-En Yang. Thinkact: Vision-language-action reasoning via reinforced visual latent planning. *arXiv preprint arXiv:2507.16815*, 2025. 3

[14] Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, and Dorsa Sadigh. Prismatic vlms: Investigating the design space of visually-conditioned language models. In *ICML*, 2024. 11

[15] Dongyoung Kim, Sumin Park, Huiwon Jang, Jinwoo Shin, Jaehyung Kim, and Younggyo Seo. Robot-r1: Reinforcement learning for enhanced embodied reasoning in robotics. In *NeurIPS*, 2025. 3

[16] Jason Lee, Jiafei Duan, Haoquan Fang, Yuquan Deng, Shuo Liu, Boyang Li, Bohan Fang, Jieyu Zhang, Yi Ru Wang, Sangho Lee, Winson Han, Wilbert Pumacay, Angelica Wu, Rose Hendrix, Karen Farley, Eli VanderBilt, Ali Farhadi, Dieter Fox, and Ranjay Krishna. Molmoact: Action reasoning models that can reason in space, 2025. 3

[17] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. LLaVA-OneVision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024. 3

[18] Mingfu Liang, Ying Wu, et al. Toa: Task-oriented active vqa. In *NeurIPS*, 2023. 2, 3

[19] Jijia Liu, Feng Gao, Bingwen Wei, Xinlei Chen, Qingmin Liao, Yi Wu, Chao Yu, and Yu Wang. What can rl bring to vla generalization? an empirical study. In *NeurIPS*, 2025. 3

[20] Li Liu, Diji Yang, Sijia Zhong, Kalyana Suma Sree Tholeti, Lei Ding, Yi Zhang, and Leilani Gilpin. Right this way: Can vlms guide us to see more to answer questions? In *NeurIPS*, 2024. 2, 3

[21] Yichen Liu, Yuhan Zhang, Chen Wang, and Hao Li. EXPRESS-Bench: Beyond the destination: A novel benchmark for exploration-aware embodied question answering. In *ICCV*, 2025. 2, 3, 6, 7, 8, 11, 12, 14, 16, 21

[22] Wufei Ma, Yu-Cheng Chou, Qihao Liu, Xingrui Wang, Celso de Melo, Jianwen Xie, and Alan Yuille. Spatialreasoner: Towards explicit and generalizable 3d spatial reasoning. In *NeurIPS*, 2025. 3, 6, 16

[23] Wufei Ma, Luoxin Ye, Celso M de Melo, Alan Yuille, and Jieneng Chen. Spatialllm: A compound 3d-informed design towards spatially-intelligent large multimodal models. In *CVPR*, pages 17249–17260, 2025. 3

[24] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, et al. OpenEQA: Embodied question answering in the era of foundation models. In *CVPR*, 2024. 3, 6, 8, 11

[25] OpenAI. Gpt-4 technical report, 2023. 6

[26] OpenAI. Introducing gpt-5, 2025. 6, 7, 8, 11, 14, 16

[27] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *NeurIPS*, pages 27730–27744. Curran Associates, Inc., 2022. 3

[28] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. In *NeurIPS*, 2022. 3

[29] Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, Manolis Savva, Yili Zhao, and Dhruv

Batra. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. 4, 6, 12, 13

[30] Arijit Ray, Jiafei Duan, Ellis Brown, Reuben Tan, Dina Bashkirova, Rose Hendrix, Kiana Ehsani, Aniruddha Kembhavi, Bryan A Plummer, Ranjay Krishna, et al. Sat: Dynamic spatial aptitude training for multimodal language models. In *COLM*, 2025. 11

[31] Allen Z. Ren, Jaden Clark, Anushri Dixit, Masha Itkina, Anirudha Majumdar, and Dorsa Sadigh. Explore until confident: Efficient exploration for embodied question answering. In *Robotics: Science and Systems*, 2024. 2, 3, 8

[32] Saumya Saxena, Blake Buchanan, Chris Paxton, Peiqi Liu, Bingqing Chen, Narunas Vaskevicius, Luigi Palmieri, Jonathan Francis, and Oliver Kroemer. Grapheqa: Using 3d semantic scene graphs for real-time embodied question answering. 2025. 2, 3

[33] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 3, 5

[34] Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao, Qianqian Zhang, et al. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*, 2025. 3

[35] Alex Su, Haozhe Wang, Weiming Ren, Fangzhen Lin, and Wenhu Chen. Pixel reasoner: Incentivizing pixel-space reasoning with curiosity-driven reinforcement learning. In *NeurIPS*, 2025. 2

[36] Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025. 3

[37] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-VL: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 3

[38] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In *ACL*, 2023. 3

[39] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *ICLR*, 2022. 3

[40] Junfei Wu, Jian Guan, Kaituo Feng, Qiang Liu, Shu Wu, Liang Wang, Wei Wu, and Tieniu Tan. Reinforcing spatial reasoning in vision-language models with interwoven thinking and visual drawing. *arXiv preprint arXiv:2506.09965*, 2025. 3, 6, 7

[41] Yuncong Yang, Jiageng Liu, Zheyuan Zhang, Siyuan Zhou, Reuben Tan, Jianwei Yang, Yilun Du, and Chuang Gan. Mindjourney: Test-time scaling with world models for spatial reasoning. In *NeurIPS*, 2025. 2

[42] Yuncong Yang, Han Yang, Jiachen Zhou, Peihao Chen, Hongxin Zhang, Yilun Du, and Chuang Gan. 3D-Mem: 3d scene memory for embodied exploration and reasoning. In *CVPR*, pages 17294–17303, 2025. 2, 3

[43] Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025. 3

[44] Jonas Zausinger, Lars Pennig, Anamarija Kozina, Sean Sdahl, Julian Sikora, Adrian Dendorfer, Timofey Kuznetsov, Mohamad Hagog, Nina Wiedemann, Kacper Chlodny, et al. Regress, don't guess – a regression-like loss on number tokens for language models. In *ICML*, 2025. 11, 12

[45] Simon Zhai, Hao Bai, Zipeng Lin, Jiayi Pan, Peter Tong, Yifei Zhou, Alane Suhr, Saining Xie, Yann LeCun, Yi Ma, et al. Fine-tuning large vision-language models as decision-making agents via reinforcement learning. In *NeurIPS*, 2024. 3

[46] Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. Group sequence policy optimization, 2025. 3

[47] Chenming Zhu, Tai Wang, Wenwei Zhang, Jiangmiao Pang, and Xihui Liu. Llava-3d: A simple yet effective pathway to empowering lmms with 3d-awareness. In *ICCV*, 2025. 3

Appendix

A.1. Extended Experiments

A.1.1. Comparison with More Training Variants

As discussed in Section 4.1, we further evaluate three internal variants for a comprehensive comparison of training strategies. The quantitative results are summarized in Table A3. Below, we detail each variant and its training setup:

- **SFT w/ NT Loss** [44]: this variant is trained with an auxiliary regression-aware loss, called Number Token Loss (NT Loss) [44], on action magnitudes instead of a cross-entropy loss. In our setting, this change brings no benefit over vanilla SFT training.
- **SFT (extended)**: to fairly compare against our two-stage training, we extend SFT training until convergence (from 7 to 20 epochs) so that the only difference lies in the training objective. While longer SFT training yields modest gains, SFT alone still underperforms our two-stage training (SFT+RL).
- **RL w/ r^{pose}**: this variant incorporates an additional distance-based reward inversely proportional to the discrepancy between the predicted and target camera poses (position and orientation), thereby injecting weak supervision of the ground-truth action into RL. Nevertheless, it still fails to learn meaningful actions without the SFT warm-up stage, highlighting that a supervised initialization followed by RL is more effective than relying on more sophisticated reward designs for RL alone.

A.1.2. More Experiments on EQA

We report additional results on EQA setups using the Open-EQA [24] dataset in Table A4. We use the same Fine-EQA [21] configuration as the base EQA pipeline as in Section 4.2, and then incorporate our AVS framework on top. As shown, plugging our AVS framework into this pipeline improves performance, yielding a 6.8%pt gain in average LLM-Match score compared to the Fine-EQA baseline.

We also illustrate the integration of our AVS framework into an existing EQA pipeline in Figure A6.

A.1.3. Cross-Dataset Generalization

To assess whether training on our relatively small-scale dataset harms the general spatial understanding of the original VLM, Qwen2.5-VL-7B [1], we additionally evaluate our models on the SAT [30] spatial reasoning benchmark, which comprises several downstream tasks across synthetic and real-image splits. As shown in Table A5, although our framework is not explicitly trained for this benchmark, it generally improves the backbone's performance, with the RL-only training strategy being the only variant that slightly underperforms. These results indicate that training on our synthetic dataset does not lead to noticeable catastrophic forgetting and can transfer reasonably well to broader spatial reasoning tasks.

A.1.4. Results with Multi-Turn Actions

We provide additional results with multi-turn action sequences in Table A6 and observe that executing multiple actions does not yield clear performance gains over the single-step setting.

We conduct these multi-turn experiments with two models. The first is our main model used in Section 4, which is trained with query views as input only and denoted as SFT+RL (Q-only). We also train a variant that takes both the query and target views as input, denoted as SFT+RL (Q+T). This SFT+RL (Q+T) model attains accuracy comparable to our original SFT+RL (Q-only) model and exhibits more stable performance between single- and multi-turn rollouts, yet its performance likewise does not improve in the multi-turn setting.

Overall, these results indicate that our continuous action design enables the agent to reach an informative viewpoint within a single turn.

A.2. Training and Experiment Setup Details

Training. We use Qwen2.5-VL-7B [1] as the backbone VLM, keeping the vision encoder frozen during training. We train the model with a batch size of 32 for SFT and 128 for RL, running 7 epochs of SFT followed by 5 additional epochs of GRPO-based reinforcement learning. Following prior work [44], we set the weight of the Number Token Loss to 0.3 for the SFT model trained with NT Loss (see Table A3). The total reward in GRPO is a weighted sum of a format reward and a verifier reward, with weights 0.3 and 1.0, respectively. We set the KL penalty coefficient β to 0.04. During GRPO training, we use a separate, frozen Qwen2.5-VL-7B [1] as the VLM verifier for computing verifiable reward r^{ver}, with a group size of 16. We use learning rates of 2×10^{-5} for SFT and 10^{-6} for GRPO.

EQA Baseline. In Table 1, for a comparison with an EQA framework, we employ the state-of-the-art Fine-EQA approach [21]. This framework includes both a frontier-based exploration strategy, which simply aims to visit unseen regions, and a goal-oriented exploration (GOE) strategy, which guides the agent toward semantically informative areas given the current observation and language query. We adopt the GOE strategy so that the framework can actively adjust its viewpoint based on contextual cues, aligning it with our Visually-Grounded Active View Selection (VG-AVS) task. Following Fine-EQA [21], we use `prism-dinosiglip+7b` introduced in Prismatic VLMs [14] as the VLM backbone and replace the older GPT-4 models with the GPT-5 family [26] for region prioritization and exploration termination decisions.

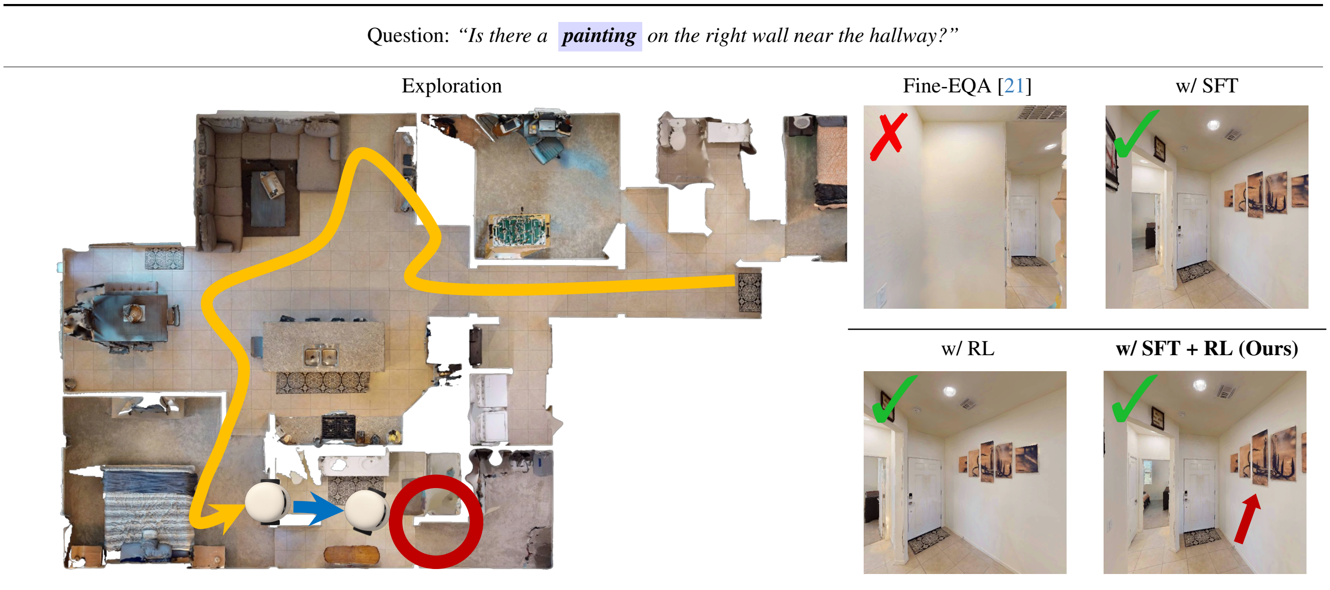

Figure A6. **Illustration of integrating our AVS framework into an EQA pipeline.** On the left, the yellow line represents the exploration path of the EQA pipeline, the blue arrow the final viewpoint refinement by our model, and the red circle the region of interest given the question. On the right, we show the views obtained by each method, where ✓ denotes a correct view with a high LLM-Match score ($= 5$) and ✗ an incorrect view with a low score (≤ 2). Given the question about the *painting*, our AVS framework effectively refines the viewpoint such that the final predicted view reveals the painting on the wall. This demonstrates that the training-free EQA pipeline lacks query-conditioned final viewpoint refinement, whereas our learning-based method can minimally adjust the viewpoint to obtain an answerable view.

Table A3. **Quantitative comparison on internal baselines.** We report VQA accuracy on AVS-ProcTHOR and LLM-Match scores on AVS-HM3D, normalized to a percentage scale. The best in each column is in **bold** and second best is <u>underlined</u>.

AVS Framework	**AVS-ProcTHOR**				**AVS-HM3D**					
	Existence	**Counting**	**State**	**Average**	**Existence**	**Counting**	**State**	**Attribute**	**Object**	**Average**
SFT	<u>91.28</u>	57.06	83.84	77.39	67.50	<u>70.77</u>	62.16	66.67	55.56	64.53
SFT (Extended)	<u>91.28</u>	61.52	<u>85.59</u>	<u>79.46</u>	76.67	67.69	68.92	73.33	55.56	68.43
SFT w/ NT Loss [44]	81.59	60.41	77.51	73.17	77.08	70.00	63.24	60.83	**63.89**	67.01
RL	86.82	<u>65.24</u>	83.41	78.49	**81.25**	70.00	<u>72.97</u>	<u>69.17</u>	60.00	<u>70.68</u>
RL w/ r^{pose}	86.05	61.52	83.41	76.99	<u>78.33</u>	66.92	72.43	**70.83**	58.89	69.48
SFT+RL (Ours)	**91.47**	**69.52**	**90.17**	**83.72**	74.58	**71.54**	**73.78**	**70.83**	<u>62.78</u>	**70.70**

Table A4. **Comparison on Open-EQA benchmark.** Each column reports normalized LLM-Match scores for different question types. Plugging our AVS framework into an exiting EQA pipeline [21] improves performance. Best in **bold**, second best <u>underlined</u>.

Method	**Object Recognition**	**Spatial Understanding**	**Object State Recognition**	**Attribute Recognition**	**Object Localization**	**Functional Reasoning**	**Average**
Fine-EQA [21]	40.00	45.26	51.11	<u>52.73</u>	**41.71**	**48.24**	<u>46.51</u>
w/ SFT	<u>44.80</u>	42.11	<u>60.00</u>	43.64	35.43	45.88	45.31
w/ RL	41.60	<u>46.32</u>	55.56	43.03	35.43	**48.24**	45.03
w/ SFT+RL (Ours)	**52.80**	**65.20**	**61.40**	**58.20**	<u>36.60</u>	<u>46.20</u>	**53.40**

A.3. More Details on Dataset Curation

We provide additional details on the data curation procedure briefly introduced in Sections 3.1 and 4.1. We first describe the AVS dataset and AVS-ProcTHOR, both curated with ProcTHOR [7] 3D scenes, and then present AVS-HM3D, which is constructed from real indoor scenes in Habitat-Matterport3D [29].

A.3.1. AVS Dataset and AVS-ProcTHOR

Both of our synthetic datasets are constructed using the same fully automatic curation pipeline on ProcTHOR [7] 3D scenes, as introduced in Section 3.1. The AVS dataset is

Table A5. **Results on SAT benchmark.** We report performance both for synthetic and real splits. For each row, the best is in **bold** and the second best is underlined.

Scene Type	Backbone	AVS Framework		
	Qwen2.5-VL-7B [1]	SFT	RL	SFT + RL (Ours)
Synthetic	59.11	**69.33**	57.31	**69.33**
Real	60.00	67.33	57.33	**77.33**

Table A6. **Quantitative comparison with multi-turn actions.** We report VQA accuracy on AVS-ProcTHOR. The best in each column is in **bold** and second best is underlined.

Training Variants	Action Steps	AVS-ProcTHOR			
		Existence	Counting	State	Average
SFT+RL (Q-only)	1	**91.47**	**69.52**	90.17	**83.72**
	2	87.60	61.52	87.55	78.89
SFT+RL (Q+T)	1	90.50	68.22	**90.39**	83.04
	2	90.50	68.22	90.17	82.96

used for training, whereas AVS-ProcTHOR is reserved for evaluation. They share identical scene configurations; the only difference lies in the question types.

The AVS training dataset consists exclusively of binary object–existence questions of the form "*Is there a target object on the supporting object?*", where the correct answer is always "*yes*". In every sample, the target object mentioned in the question is guaranteed to exist on the specified supporting object, so that selecting a viewpoint that clearly reveals the object consistently yields a positive reward from the frozen VLM verifier, while uninformative viewpoints yield zero reward. This design aligns the verifier's feedback directly with the quality of view selection during RL training.

In contrast, AVS-ProcTHOR employs more diverse question types for a more comprehensive evaluation: *Existence*, *Counting*, and *State*. Existence questions ask whether the target object is on the supporting object, counting questions ask how many instances of the target object are on that supporting object, and state questions query the state of the target object. For the *Existence* category, we additionally cast questions into a multiple-choice format to make random guessing less likely to succeed.

Building upon the notations introduced in Section 3.1, the overall curation pipeline consists of three stages:

Stage 1: Scene Modification by Question Type. Given a data sample, we first select a pair of a supporting object and a target object placed on it using ProcTHOR scene metadata, and modify the scene if needed to make it suitable for the specified question type.

The scene modification rules for each question type are as follows:

- ***Existence.*** No scene modification done.
- ***Counting.*** Place between two and five instances of the target object on or near the supporting object.
- ***State.*** Choose a target object with a controllable state (e.g., Faucet with on/off, Book with open/closed, Mug with filled/empty) and randomly assign one of its possible states.

Stage 2: Viewpoint Sampling. To obtain a target camera pose s_{tgt}, we leverage the built-in function in ProcTHOR, which returns a set of agent poses (positions and orientations) from which a specified object is both visible and interactable. Specifically, we call the `GetInteractablePoses()` built-in function on the supporting object to retrieve a list of nearby camera poses that can effectively observe it. We then sort these candidate poses by their distance to the supporting object, select the top 10 closest poses. We then randomly sample one pose from this subset and use it as the target camera pose s_{tgt}, provided that it satisfies the rule described in Equation 1.

Once s_{tgt} is determined, we generate candidate query poses s_{qry} by applying actions that move the agent slightly away from the target view. Specifically, we randomly sample 10 actions, parameterized as in Section 3.2, each consisting of a heading rotation in $[-45°, 45°]$, a forward translation in $[50, 150]$ (centimeter), and a view rotation in $[-45°, -15°] \cup [15°, 45°]$, where the forward translation is applied in the opposite direction so that the agent moves backward from the target viewpoint. Among the resulting candidates, we randomly select one query pose that satisfies the rule described in Equation 2.

In Equations 1 and 2, we set $\epsilon^{\text{sup}}_{\text{vis}} = 5{,}000$, $\epsilon^{\text{obj}}_{\text{vis}} = 10{,}000$, and $\epsilon^{\text{obj}}_{\text{inv}} = 30$, with a 90° field of view, and the image resolution is 512×512. We use the default camera height of 90 centimeters and the default camera elevation angle of 15° downward in ProcTHOR [7].

Stage 3: Question–Answer Generation and Filtering. Given the selected target/supporting objects and their relation in the scene, we instantiate rule-based question templates for each question type.

For every data sample tuple, we retain only the samples for which the VLM verifier answers the question correctly given the target view o_{tgt}.

A.3.2. AVS-HM3D

As discussed in Section 4.1, we construct an additional benchmark, AVS-HM3D, on real indoor environments from Habitat-Matterport3D [29] to assess the generalization of our method beyond the synthetic ProcTHOR scenes used for training. We reuse triplets of (question, answer, ground-truth view) from the validation split of the Fine-EQA dataset [21], where the ground-truth view is the human-annotated frame from which the question–answer pair is derived. Among its seven question types, we discard those that rely on external world knowledge or holistic layout priors and retain five locally grounded types: *Attribute*, *Counting*, *Existence*, *Object*, and *State*. Given the ground-truth view, we then sample five candidate query views by applying small actions that move the agent away from the ground-truth view, following the same procedure as in the ProcTHOR setup. Due to the absence of scene metadata (e.g., object visibility and object relations) in real scenes, we manually select for each triplet, the query view that offers partial visual evidence—sufficient to guide the agent toward the ground-truth view, yet insufficient to directly answer the question. Using a GPT-5 model [26], we filter out low-quality items, such as cases where the question is not solvable from the ground-truth view or the ground-truth view has poor visual quality. The resulting curated benchmark comprises 208 samples.

A.4. Prompts

We present our input prompts in Figure A7, including the action prompt that enforces the VLM to predict action parameters and the format prompt that specifies the desired output structure. In the action prompt, `VQA_question` is a placeholder that is replaced by the actual question about the scene. When training purely with RL, we adopt a "think-then-act" format, whereas during SFT the model is only supervised to directly predict the action parameters without any reasoning traces. Directly switching an SFT-trained model to the think-then-act format often causes a degenerate behavior where the model outputs the action first and then hallucinates a post-hoc "thinking" process. To mitigate this mismatch, we redesign the RL output format so that the model first makes an initial action guess, then reasons to refine it, and finally outputs a revised action, which encourages a smooth transition from direct prediction to genuine reason-before-acting behavior.

A.5. More Qualitative Comparisons

We provide more qualitative comparison results of Visually-Grounded Active View Selection in Figure A8.

```
Action Prompt:
You are an embodied agent navigating a 3D scene from an egocentric camera. Given the
current image and a question about the scene, predict the optimal NEXT action parameters
to reach a better viewpoint.
Action parameters (return integers only):
1) Heading rotation (deg) in (-180, 180]: Azimuth yaw about your vertical axis BEFORE
moving. Positive = clockwise/right, negative = counterclockwise/left, 0 = no rotation.
2) Forward distance (cm) >= 0: Move forward in the NEW facing direction after the
rotation. 0 = no move.
3) View rotation (deg) in (-180, 180]: Final azimuth adjustment AFTER moving, relative
to your post-move heading. Same sign convention as rotation.

Goal: Choose a heading rotation angle, moving forward distance, final-viewing rotation
angle that maximizes visibility of task-relevant objects and minimizes occlusion.
Example: Rotating -90 degrees, moving forward 50 cm, then rotating 90 degrees is
equivalent to translating 50 cm to your left while keeping the original heading.
Question: "{VQA_question}"
DO NOT answer the question; ONLY predict the next action parameters.

RL Format Prompt:
First output the reasoning process in <think> </think> tags. Then, output the final
predictions in <head> </head>, <fwd> </fwd>, <view> </view> tags in order.
The text between <head> and </head> must be the angle in degrees (-180, 180], <fwd>
and </fwd> must be the nonnegative forward distance, and <view> and </view> must be the
final viewing angle in degrees (-180, 180].
Each must be exactly one integer number (no units, no extra text).
In the reasoning process, explicitly reason about (1) how much to rotate to determine
the moving direction, (2) how far to move forward to approach, (3) how much to further
adjust your azimuth angle from your moving direction for the best view.

SFT-then-RL Format Prompt:
First, output your initial guess for the action parameter values.
Then, think carefully to refine your initial guess for each action parameter.
After output initial guess for the action parameters, output your reasoning process
within <think> </think> tags, and then provide the final guess within <head> </head>,
<fwd> </fwd>, and <view> </view> tags, respectively.
The text between <head> and </head> must be the rotation angle in degrees in the range
(-180, 180]; the text between <fwd> and </fwd> must be the nonnegative forward distance;
and the text between <view> and </view> must be the final viewing angle in degrees in
the range (-180, 180].
Each must be exactly one integer (no units, no extra text). In the reasoning process,
explicitly reason about (1) how much to rotate to determine the moving direction, (2)
how far to move forward to approach, (3) how much to further adjust your azimuth angle
from your moving direction for the best view.

For example:
<head> INITIAL_GUESS </head> <fwd> INITIAL_GUESS </fwd> <view> INITIAL_GUESS </view>
<think> REASONING PROCESS </think>
<head> FINAL_GUESS </head> <fwd> FINAL_GUESS </fwd> <view> FINAL_GUESS </view>
```

Figure A7. **Input System Prompts.** The action prompt provides the task context and action parameterization explanation, while the format prompt specifies the required reasoning trace and output tag structure.

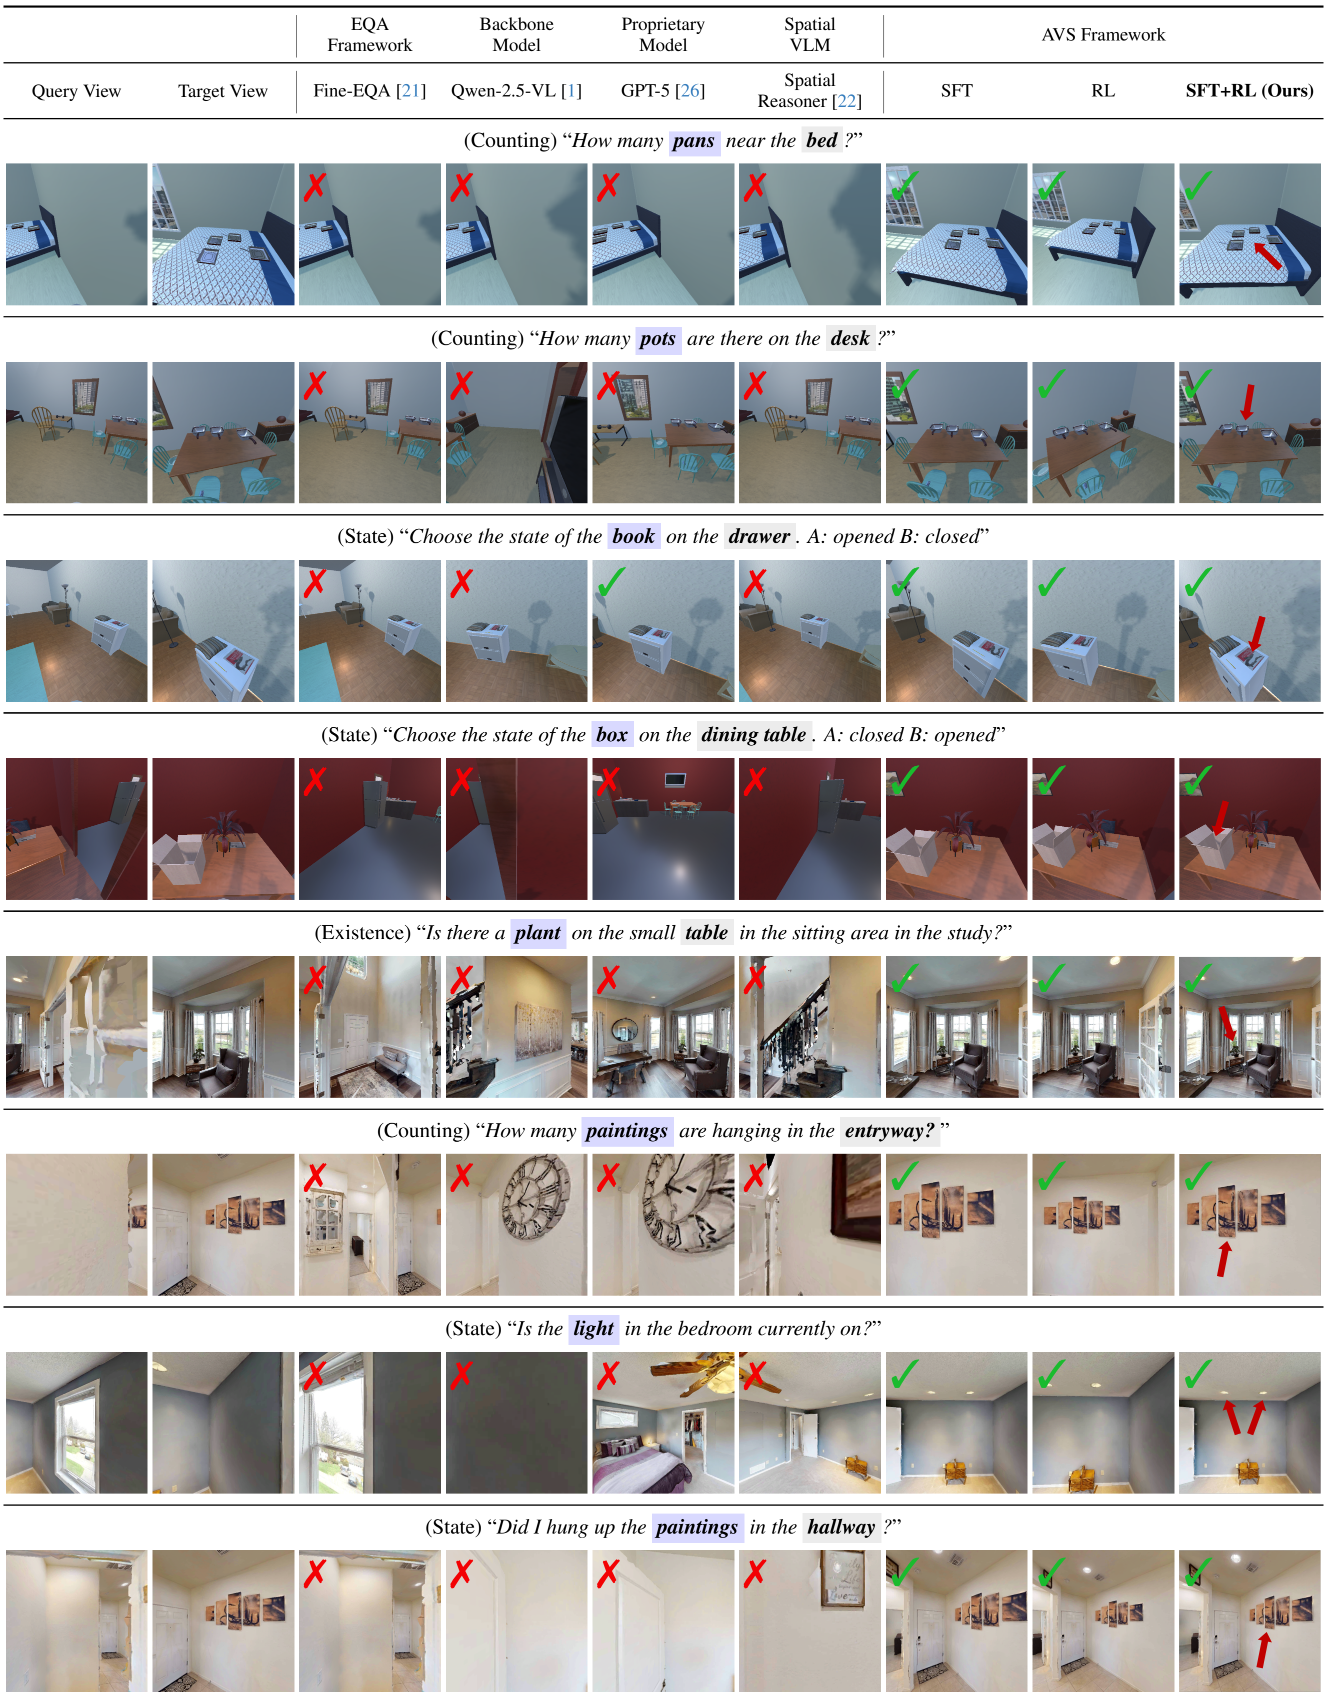

Figure A8. **More qualitative results on AVS-ProcTHOR (top four rows) and AVS-HM3D (bottom four rows).** Blue and gray mark the object of interest and surrounding cue objects, respectively. ✓ denotes correct answers (LLM-Match $=5$), ✗ incorrect ones (LLM-Match ≤ 2).

A.6. Qualitative Examples of AVS Framework Reasoning

In the following, we present qualitative examples illustrating how our AVS framework reasons from the query view to predict the desired action parameters.

Table A7. **Qualitative example illustrating the model's reasoning process and executed actions in AVS-ProcTHOR.** Given the input view (left), the model identifies partially observable visual cues and infers the appropriate action needed to complete the missing context. The predicted action sequence (<head>, <fwd>, <view>) is then executed, producing the output view (right), from which the verifier successfully answers.

<table>
<tr><th>Input View</th><th>Output View</th><th>Input View</th><th>Output View</th></tr>
<tr><td></td><td></td><td></td><td></td></tr>
<tr><td colspan="2">Choose an object that is present on the bed:
A. bottle B. teddy bear C. pen D. towel</td><td colspan="2">Choose an object that is present on the dresser:
A. vase B. potato C. apple D. faucet</td></tr>
<tr><td colspan="2"><think> The task is to find an object on the bed. To do this, I need to exit the room by opening the door and then proceed to the bedroom. The door is slightly ajar, so I can see part of the adjacent room. The bedroom is located to the right of the current room, which means I should turn right. After exiting the room, I will be able to see the bed and search for the object there. <think>
<head>73</head> <fwd>145</fwd> <view>-49</view></td><td colspan="2"><think> The task is to find an object on the dresser. To do this, I need to move towards the dresser. The dresser is located to the left of the door, so I should turn left and walk forward. Once there, I can look around to see if any of the objects on the dresser are visible. <think>
<head>-64</head> <fwd>79</fwd> <view>21</view></td></tr>
<tr><td></td><td></td><td></td><td></td></tr>
<tr><td colspan="2">How many bottles near the TV stand?</td><td colspan="2">How many mugs are there on the TV stand?</td></tr>
<tr><td colspan="2"><think> I need to move towards the right side of the room where the TV stand is located. By rotating 76 degrees to the right, I can face the direction of the TV stand. Then, by moving forward 125 units, I will be closer to the TV stand. Finally I'll turn -44 degrees to the left to face the TV stand directly. <think>
<head>76</head> <fwd>125</fwd> <view>-44</view></td><td colspan="2"><think> The TV stand is located to the upper left in the image. To see it, I need to turn my head to the left. After that, I should move forward until I can see the TV stand clearly. Once I'm close enough, I'll turn my head slightly to the right to focus on the mugs. <think>
<head>-54</head> <fwd>145</fwd> <view>36</view></td></tr>
</table>

Input View Output View | Input View Output View


Choose a faucet's state on the sink.
A: turned off. B: turned on.

<think> The faucet is on the left side of the sink, so I need to move towards it. The angle needed to get there is -57 degrees, which means I should turn left. After reaching the sink, I'll need to adjust my view slightly to the right by 26 degrees to see the faucet clearly. <think>
<head>-57</head> <fwd>134</fwd> <view>26</view>


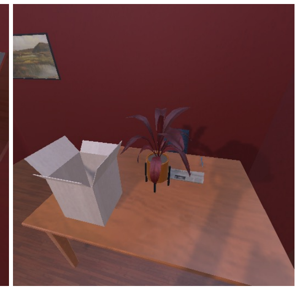

Choose a box's state on the dining table.
A: closed. B: opened.

<think> The table is on the left side of the image. To see the table, I need to turn to the left. Then, I can move forward to get closer to the table. Finally, I will turn slightly to the right to face the table. <think>
<head>-67</head> <fwd>55</fwd> <view>24</view>


Choose an object that is present on the dresser:
A. mirror B. CD C. pen D. coffee machine

<think> The goal is to find an object on the dresser. The dresser is located to the right of the current position, so I need to move towards it. The rotation value of 64 indicates a right turn, and the forward distance of 125 moves me closer to the dresser. The final look value of -38 suggests a slight left turn to align with the dresser. <think>
<head>64</head> <fwd>125</fwd> <view>-38</view>


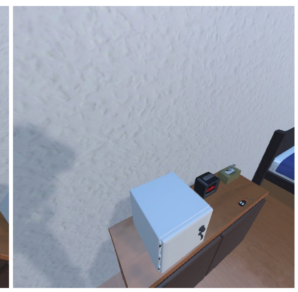

Choose an object that is present on the dresser:
A. fork B. tissue box C. cloth D. wine bottle

<think> The task is to find an object on the dresser. The dresser is located to the right of the room, which can be seen by looking towards the right. To reach it, I need to turn 76 degrees to the right, walk 90 units, and then look -34 degrees to the left to face the dresser. <think>
<head>76</head> <fwd>90</fwd> <view>-34</view>


Choose a paper towel roll's state on the toilet.
A: used up. B: unused.

<think> The toilet is on the left side of the image, and it is necessary to move towards it to see its state. By rotating -64 degrees, we can face the toilet, and by moving forward 79 units, we can get closer to it. Finally, by looking at 35 degrees, we can focus on the toilet's state. <think>
<head>-64</head> <fwd>79</fwd> <view>35</view>


Choose a desk lamp's state on the dining table.
A: turned off. B: turned on.

<think> I need to move towards the table on the left. By rotating -47 degrees, I can face the direction of the table. Then, by moving forward 145 units, I will get closer to the table. Finally, by adjusting my view with 33 degrees, I will have a clear view of the table. <think>
<head>-47</head> <fwd>145</fwd> <view>33</view>

Table A8. **Qualitative example illustrating the model's reasoning process and executed actions in AVS-HM3D.** Given the input view (left), the model identifies partially observable visual cues and infers the appropriate action needed to complete the missing context. The predicted action sequence (<head>, <fwd>, <view>) is then executed, producing the output view (right), from which the verifier successfully answers.

Input View	Output View	Input View	Output View
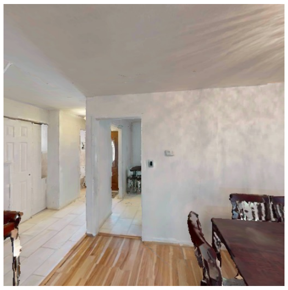		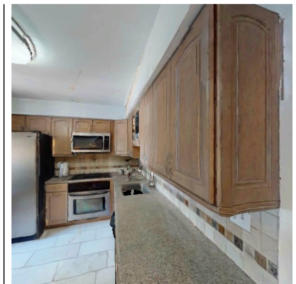	
Is there a vase on the dining table for the flowers?		Did I leave the refrigerator door open in the kitchen?	
<think> The table on the right side of the image is the dining table. The table is not fully visible, but the top right corner of the image shows part of it. To see the table, I need to turn to the right and move forward. <think> <head>94</head> <fwd>75</fwd> <view>-51</view>		<think> The refrigerator is on the left side of the image. To see it, I need to turn left and move forward. <think> <head>-64</head> <fwd>127</fwd> <view>21</view>	
		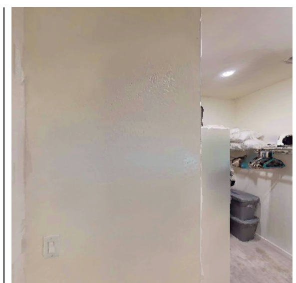	
Can you describe the hue of the sizable pot situated adjacent to the entrance?		Are there any suit jackets hanging in the closet?	
<think> The sizable pot is located at the entrance, which is on the right side. To see it, I need to turn to the right and move towards the door. <think> <head>64</head> <fwd>79</fwd> <view>-35</view>		<think> The closet is on the right side, so I need to turn right and walk forward to get there. <think> <head>46</head> <fwd>127</fwd> <view>-17</view>	
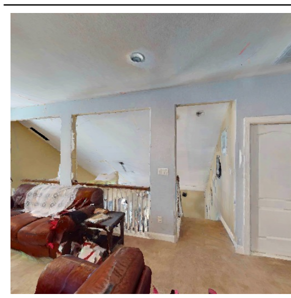	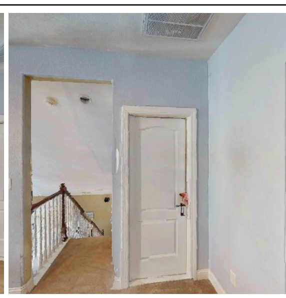	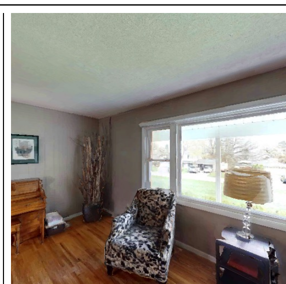	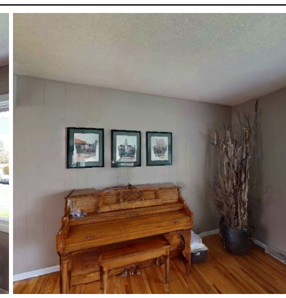
How many decorative items are hanging from the door handle next to the stairs?		How many framed pictures are hanging on the wall near the piano?	
<think> The door on the right side of the image is the one next to the stairs. I need to move towards it to see if there are any decorative items hanging from the door handle. <think> <head>64</head> <fwd>125</fwd> <view>-42</view>		<think> The framed picture is on the wall behind the piano. I need to move towards the piano to see it clearly. <think> <head>-64</head> <fwd>145</fwd> <view>33</view>	

Input View	Output View	Input View	Output View

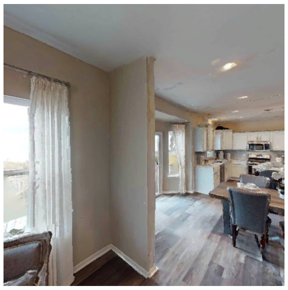
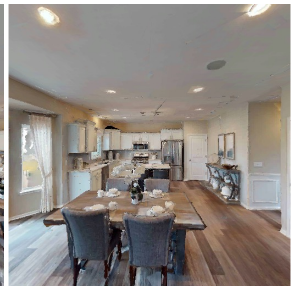

If I have six quests for dinner, how many chairs are available at the dining table?

<think> The dining table is located towards the right side of the room, and there are chairs visible on both sides of it. To see the entire table and the chairs, I need to move towards the right. <think>
<head>67</head> <fwd>125</fwd> <view>-29</view>


What is the small lamp next to the bed used for?

<think> The lamp is on the left side of the bed. I need to move towards it by turning left and then moving forward. After reaching the lamp, I will turn right slightly to face it directly. <think>
<head>-64</head> <fwd>125</fwd> <view>33</view>


What is the color of the lamp shade next to the plant?

<think> The plant is located on the far right side of the room, which is currently out of view. To see it, I need to turn right and move forward. This will allow me to see the plant and its lamp shade. <think>
<head>67</head> <fwd>125</fwd> <view>-24</view>


Is the dining table set for dinner?

<think> The dining table is located in the area behind the sofa, which can be seen by looking towards the left. <think>
<head>-54</head> <fwd>167</fwd> <view>23</view>

Table A9. **Qualitative example illustrating the model's reasoning process and executed actions in Fine-EQA [21].** Fine-EQA first performs its own exploration and provides an observation (left). Based on it, our model identifies visual cues that are insufficient for answering and reasons that an additional action is required to complete the missing context. The model then predicts an action sequence (<head>, <fwd>, <view>), executes it, and obtains the updated view (right), from which the verifier successfully answers the question.

Final View of Fine-EQA [21]	Refined View by Ours	Final View of Fine-EQA [21]	Refined View by Ours
		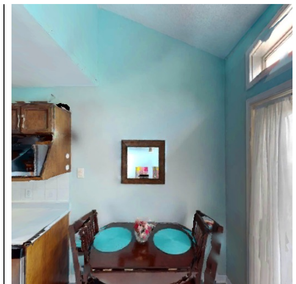	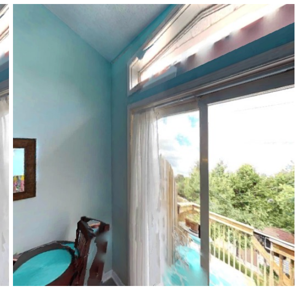
How many chairs are there at the bar counter in the kitchen?		Are the curtains in the living room closed?	
<think> The kitchen area is located on the left side of the hallway. To get a clear view of the kitchen, I need to turn left and walk forward. This will allow me to see the bar counter and the chairs around it. <think> <head>-64</head> <fwd>145</fwd> <view>33</view>		<think> The curtains are on the right side of the image, so I need to turn right and move forward to see them. <think> <head>76</head> <fwd>103</fwd> <view>-34</view>	
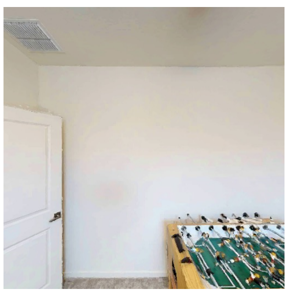		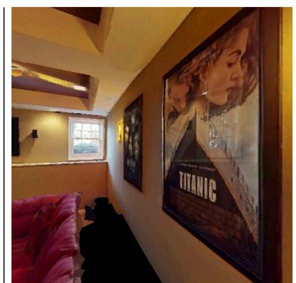	
Can you check if the welcome mat is still in front of the door in the hallway?		Did I leave the television on in the living room?	
<think> The door is on the left side of the image, so I need to turn left to face it. Then, I can move forward to get closer to the door. Finally, I'll slightly turn right to center the door in my view. <think> <head>-74</head> <fwd>90</fwd> <view>32</view>		<think> The television is likely in the living room, which is to the left. To get there, I need to tum left and walk forward. <think> <head>-76</head> <fwd>125</fwd> <view>42</view>	