

# FAST-EQA: Efficient Embodied Question Answering with Global and Local Region Relevancy

Haochen Zhang<sup>1†\*</sup>Nirav Savaliya<sup>2\*</sup>Faizan Siddiqui<sup>2</sup>Enna Sachdeva<sup>2</sup><sup>1</sup>Carnegie Mellon University, USA<sup>2</sup>Honda Research Institute USA

haochen4@andrew.cmu.edu, {nsavaliya, faizan\_siddiqui, enna\_sachdeva}@honda-ri.com



Figure 1. In the illustrated scenario, our FAST-EQA agent first localizes relevant regions, such as the master bedroom and guest bedroom, and identifies the visual target: bedsheet. Guided by a semantic-aware global exploration strategy focused on relevant rooms, it navigates across these regions while maintaining and updating a target-specific memory based on visual relevance. Once sufficiently confident, the agent queries a large vision–language model (here, GPT-4o) to answer the question using the stored visual observations.

## Abstract

*Embodied Question Answering (EQA) combines visual scene understanding, goal-directed exploration, spatial and temporal reasoning under partial observability. A central challenge is to confine physical search to question-relevant subspaces while maintaining a compact, actionable memory of observations. Furthermore, for real-world deployment, fast inference time during exploration is crucial. We introduce FAST-EQA, a question-conditioned framework*

*that (i) identifies likely visual targets, (ii) scores global regions of interest to guide navigation, and (iii) employs Chain-of-Thought (CoT) reasoning over visual memory to answer confidently. FAST-EQA maintains a bounded scene memory that stores a fixed-capacity set of region–target hypotheses and updates them online, enabling robust handling of both single- and multi-target questions without unbounded growth. To expand coverage efficiently, a global exploration policy treats narrow openings and doors as high-value frontiers, complementing local target seeking with minimal computation. Together, these compo-*

\*Equal contribution. †Work done during internship at HRI.

*nents focus the agent’s attention, improves scene coverage, and improve answer reliability while running substantially faster than prior approaches. On HMEQA and EXPRESS-Bench, FAST-EQA achieves state-of-the-art performance, while performing competitively on OpenEQA and MT-HM3D. All source code will be publicly released. More information can be found on our project page: <https://astronirav.github.io/fasteqa>*

## 1. Introduction

In order to have generalizable robot assistants in human-centric environments such as homes, we seek the ability for robots to actively explore unknown environments and handle diverse natural language queries. This behavior can be characterized by the task of Embodied Question Answering (EQA), where a physical or simulated embodied agent is asked to answer a natural-language question via moving and perceiving in an unexplored 3D environment, rather than relying on a static image. The agent must combine natural language understanding, visual scene understanding, exploration, and spatial and temporal reasoning under partial observability to gather the evidence needed to answer. A core challenge is deciding where to explore at each step, narrowing the search to question-relevant regions while keeping a compact memory representation of what has been seen, especially in unseen environments. To add more complexity, practical tasks often have multiple targets, requiring embodied agent to navigate and gather information from multiple sources to answer the question. For instance, to answer the question “*Are the curtains in the bedroom the same color as those in the living room?*”, the agent must visit both the rooms, detect the curtains in each, store the observations in memory, and then reason over this evidence to provide the answer. Finally, in order to support physical robots in the real world, EQA approaches must follow deployment constraints with regards to power, memory, and inference time.

Recent advances in large language models (LLMs) and vision-language models (VLMs) have shown impressive capabilities in planning, reasoning, and question-answering (QA) [19, 20, 38] in the field of 2D vision. Thus, existing EQA systems increasingly lean on these models to choose exploration directions, quantify confidence, and inject pre-trained, web-scale knowledge that can anchor high-level goals. However, VLMs struggle with long-horizon embodied AI tasks due to shallow working memory, limited spatial reasoning, and weak integration of observations over time. Trained primarily on static image–text corpora, they lack the sequential perception–action grounding needed to maintain coherent representations of large, partially observed scenes or to selectively retain task-relevant information [27]. These issues become especially acute in multi-target tasks, where agents must jointly reason over spatial

and temporal evidence gathered from multiple locations.

To mitigate these limitations, many methods maintain structured memory in the form of object-centric metric-semantic scene graphs [15, 21, 29]. While effective for capturing symbolic relationships, explicit graph construction is computationally and memory intensive, slows real-time operation, and often collapses nuanced spatial detail into coarse edges. Alternatively, storing raw images or memory snapshots [33] preserves fine-grained layout but causes memory usage to grow without bound as exploration proceeds—an issue that is particularly acute in long-horizon and multi-target tasks. Thus, existing memory designs trade off either efficiency or scalability, limiting their practical utility.

Furthermore, exploration strategies present parallel challenges. Most existing approaches are frontier-based [25], which are effective in open spaces but blind to indoor regularities, inattentive to question-conditioned semantics, prone to repeated exploration when uncertain, and unable to target intuitive transitions such as doors, hallways, or rooms that are often key to exploration. Consequently, both memory and control components of existing EQA systems struggle to scale gracefully to complex, long-horizon tasks.

In the context of the above challenges, we present FAST-EQA - FAst, Semantics-aware, Target-driven Exploration for Embodied Question Answering (Figure 1), an embodied QA system that couples semantically-guided global and local exploration policies with bounded visual memory. Given a question, FAST-EQA first uses an LLM to extract candidate visual goals or targets, initializes a spatial memory and ranks regions of the environment by their likelihood of containing relevant evidence for the goals. Exploration then proceeds with two complementary policies: Global Relevance (GR) Exploration, which departs from standard frontier-based exploration (FBE) by prioritizing transitional waypoints such as hallways and doors to efficiently reach promising areas, and Local Relevance (LR) Exploration, which assess relevant local regions (e.g. rooms) on their informativeness for answering the question. The agent interleaves GR and LR steps dynamically until it can answer confidently. As the agent explores, it maintains an explicitly-bounded visual memory, maintaining a fixed budget of  $k$  visual snapshots per target to ensure efficiency and scalability. After taking LR steps, the system invokes the VLM’s chain-of-thought reasoning over these snapshots to produce the final answer.

We evaluate FAST-EQA comprehensively on both multiple-choice QA and open-answer datasets across HMEQA [25], EXPRESS-Bench [17], OpenEQA [22], and MT-HM3D [37], achieving superior performance on HMEQA and EXPRESS-Bench and competitive results on OpenEQA and MT-HM3D, underscoring the method’s versatility across question types. Additionally, we show that

our method has superior real-time inference capabilities, consistently producing navigation decisions faster during each step of exploration. Our contributions can be summarized as follows:

- We present FAST-EQA, a lightweight embodied QA framework that runs in near real-time with a compact memory footprint, making it well-suited for deployment on embodied agents.
- We propose semantically guided frontier-selection policy designed for indoor environments, which prioritizes narrow openings and doors as informative frontiers to transition between semantically different regions while directing exploration toward relevant visual targets and goals.
- To ensure efficient scaling to multi-target questions, we introduce a bounded memory that selectively retains target-specific visual snapshots, enabling lightweight operation.

## 2. Related Works

### 2.1. Embodied Question Answering

Embodied Question Answering (EQA) tasks agents with navigating environments to answer language queries. Early works such as EmbodiedQA [8] and IQA [12] established the problem of coupling active perception with language grounding. Hierarchical methods like Neural Modular Control [9] decompose navigation into semantic sub-goals, while Multi-Target EQA [35] extended reasoning to comparative queries across multiple objects. Subsequent advances incorporated richer perception: point cloud inputs improved photorealistic navigation [30], and VideoNavQA [2] highlighted reasoning challenges even when navigation is removed. Recent efforts emphasize efficiency and generalization. EfficientEQA [6] and HM-EQA [25] introduced semantic-guided exploration and early-stopping, while Fine-EQA [17] benchmarked fine grained exploration-aware answering. Foundation model integration has also emerged: OpenEQA reframed EQA in the context of VLMs, and graph- or memory-augmented approaches (e.g., GraphEQA [26], 3D-Mem [33], MemoryEQA [37]) addresses long-horizon reasoning. Collectively, these works show progress toward real-world EQA but reveal persistent gaps in spatial reasoning, structured memory for complex queries, and exploration time.

### 2.2. VLMs for Embodied Tasks

Large-scale vision-language models (VLMs) have been adapted to embodied contexts to improve reasoning. PaLM-E [10] demonstrated that grounding multimodal LLMs in sensory inputs enables general-purpose planning and perception, with transfer across VQA and robotics. SpatialVLM [5] addressed VLMs’ weakness in 3D reasoning by training on large-scale spatial QA, improving distance

and relation understanding. Complementary to this, Think, Act, and Ask [7] showed how LLMs can orchestrate exploration and dialog for personalized navigation, highlighting the potential of interactive reasoning in embodied tasks. While these works showcase the versatility of foundation models, their effectiveness often depends on coupling with structured spatial memory for reliable reasoning.

### 2.3. Memory Representations for Embodied Tasks

Effective memory remains central to Embodied QA. Metric- and semantic-based approaches like Semantic MapNet [3] and Goal-Oriented Semantic Exploration [4] build allocentric maps that accumulate semantic cues for navigation and object search. Semi-Parametric Topological Memory and VLMs further demonstrated that topological and vision-language fused maps support robust localization and zero-shot navigation. Extensions such as VLFM [34] leverage language-grounded frontier exploration, while ConceptGraphs [13] and GraphEQA [26] use 3D semantic scene graphs for relational reasoning. Memory-centric designs, e.g. MemoryEQA [37] and ReMEMBR [1], place structured, queryable memory at the core of decision-making, enabling long-horizon and temporal reasoning. These advances confirm that semantically enriched memory is key to bridging perception and reasoning, but also underscore challenges in scaling memory to open-vocabulary, complex environments.

## 3. Problem Formulation

We formulate the problem of EQA similar to [25, 37]. For an EQA task, an agent is initialized at a random location within a given single-floor environment and asked a natural language question. In any given episode, the agent will explore up to  $t \leq N$  timesteps where  $N$  is the maximum time steps allowed for the robot, which is scaled relative to the size of the scene. Each action step corresponds to a discrete change in the agent’s pose, either through translation or rotation, and planning of the next pose. An EQA scenario then consists of a sequence of RGB-D observations  $O = \{o_i\}_{i=0}^N$  obtained from exploration, a state sequence of agent poses during exploration,  $S = \{s_i\}_{i=0}^N$ , an input question  $Q$ , and some ground-truth answer  $y$ . In this work, we consider two types of questions: multiple-choice, where the set of candidate answers is guaranteed to contain the correct label  $y$ , and open-ended, where the answer is expressed as a free-form text sentence.

Furthermore, in our method we define the relevant room(s) for exploration for a given query as a set  $R$ . The visual target(s) that needs to be observed are denoted as  $T = \{T_m\}_{m=0}^M$  where  $M$  is the total number of visual targets, and the current region the agent sees at timestep  $t$  is denoted as  $R_t$ . For example, for the query “*Is the bedsheet color in the master bedroom the same*

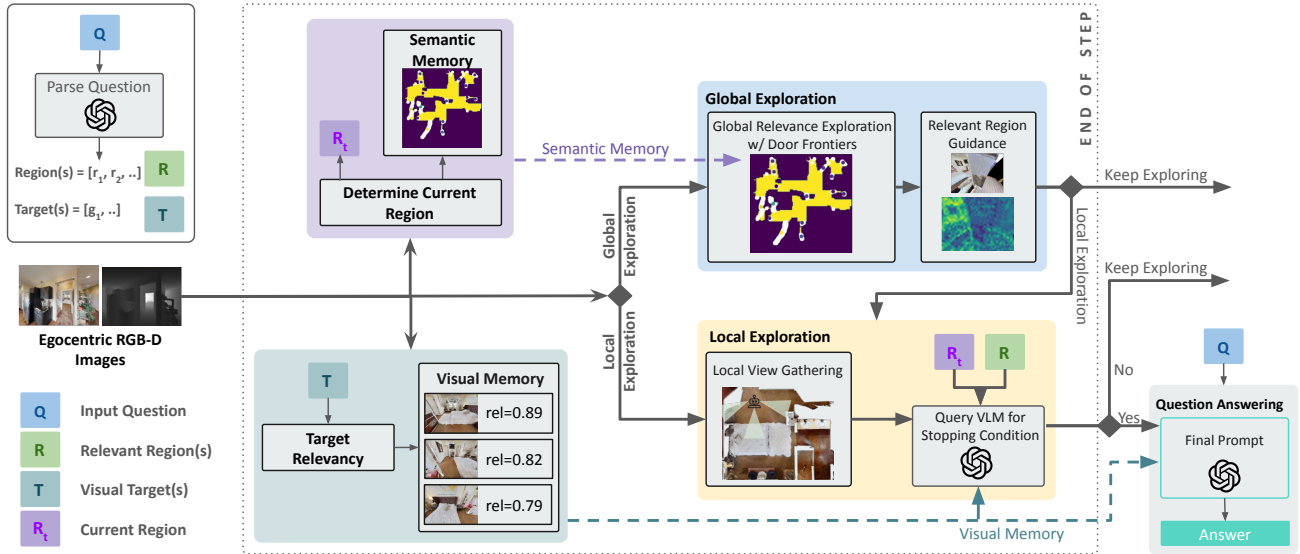


Figure 2. FAST-EQA processes the question (Q) by extracting relevant regions (R) and visual targets (T). At each step, it localizes its current region ( $R_t$ ) and updates a semantic memory. For each target  $T_m$ , it maintains a dedicated memory  $\psi_m$  that is refined using a visual relevance score. The agent employs a semantic frontier-guided global exploration strategy, leveraging narrow passages (e.g., doors and hallways) to effectively search for relevant semantic regions. When a relevant region is reached, it switches to local exploration to refine the target-specific memory. Once the stopping criterion is satisfied, the agent queries a VLM to generate the final answer. Dotted lines indicate module inputs while solid lines indicate procedural direction of the system.

as in the guest bedroom”, the relevant regions would be  $R = \{\text{'master bedroom'}, \text{'guest bedroom'}\}$  and the visual targets would be  $T = \{\text{bedsheet}_{\text{master}}, \text{bedsheet}_{\text{guest}}\}$ . The scenario concludes upon satisfaction of the stopping criterion, which, in our setup, corresponds to prompting GPT-4o with stopping prompt and the top- $k$  relevant images. All the prompts are described in the Supplementary Material.

## 4. Methodology

The FAST-EQA framework consists of global and local relevancy-guided exploration, visual memory retrieval, and chain-of-thought reasoning. The system diagram for how the framework operates in a given timestep  $t$ , is shown in Figure 2.

At the core of EQA methods is the exploration policy used. Existing methods for EQA have often relied on exploration policies that ask a VLM to select the next best direction from an annotated set of points on the current 2D observation to semantically weight frontiers [25, 37]. In most cases however, scenes have multiple rooms and the agent is often not be in a region relevant to the question. Thus, selecting the next direction in this way can be both inefficient and unintuitive. Given a certain query, humans already have some prior belief of which regions may be relevant and can quickly explore each of these regions. E.g. if asked about a bed in an unknown scene,

the most natural strategy is to look for bedrooms, eliminating rooms that are not a bedroom and going through hallways to reach where a bedroom may be. Inspired by this, our method first leverages an LLM to parse the input query and select which room types may be relevant for exploration. For explicit queries such as “Is the kitchen tap on?”, the relevant room would be {‘kitchen’}. For implicit queries such as “Where did I leave my blanket?”, the LLM is asked to guess which rooms may be relevant, returning {‘bedroom’, ‘living room’}. This extracted information acts as an information filter for the subsequent exploration, narrowing the search space and focusing the exploration on observing the correct global region(s) first. From its initial location, we also spin the agent 360 degrees to capture a panoramic view of the scene and form an initial occupancy map. Each subsequent exploration step follows either Global Relevance Exploration or Local Relevance Exploration, which are treated as distinct step types.

### 4.1. Global Relevance Exploration

For agent’s global exploration, we leverage a lightweight 3D voxel representation with fixed height only for occupancy mapping and exploration tracking. Voxels are updated in real-time through Truncated Signed Distance Function (TSDF) fusion [36] as the agent explores, tracking both occupied voxels and explored voxels within view using the input depth. Leveraging the occupancy map obtained, we implement a detector that identifies navigable narrow open-

ings in the scene, which likely indicate doorways or hallways that are informative transitions between regions that are semantically different. To determine narrow openings that are likely doorways and hallways, we take a 2D slice of the 3D voxelized occupancy grid at the height of the agent and use the scene structure to determine candidate voxels  $\mathcal{V}$  by checking the occupancy of each voxel in the slice and the voxels within  $\Delta x$ , and  $\Delta y$  directions nearby. Here  $\phi(v)$  is the occupancy of voxel  $v$ .

$$\mathcal{V} = \left\{ v \mid (\phi(v) = 0 \wedge \phi(v \pm \Delta x) = 1 \vee \phi(v \pm \Delta y) = 1) \right\} \quad (1)$$

We then cluster these candidate voxels using the DBSCAN algorithm [11] and use the cluster centroids as our global exploration frontiers,  $\mathcal{F} = \{f_1, f_2, \dots, f_J\}$  for varying number of frontiers,  $J$ . We maintain these frontiers in a priority queue, prioritizing frontiers that are a) formed from larger clusters and b) closer to unexplored regions. With this global exploration method, our agent is able to efficiently transition between semantically different regions in an environment to maximize coverage of the scene and the likelihood of traveling to question-relevant regions.

At every step of exploration, we query a VLM with the current observation  $o_t$  to output what room(s) are seen from the agent’s current perspective. For the VLM, we leverage Prismatic-VLM [18], which is a 7B parameter model that can be queried locally, to account for the high computational cost of querying this VLM at every step. If a relevant room is detected from the agent’s current position but the agent is outside the room, we leverage language-aligned feature guidance to first guide the agent inside the relevant room. As shown in Figure 3, the 2D image is featurized using AM-RADIO with a SigLIP adaptor [24] and queried with the relevant room name to obtain language-aligned features to obtain a contour segment for the target room to move into. The centroid of this contour is then projected into 3D space using depth information, producing a direction vector that points the agent toward the room, enabling efficient and semantically grounded entry.

## 4.2. Local Relevance Exploration

Once the agent is inside a relevant room or region, the primary goal of the agent is to gather local views in search of the visual target. We spin the agent around 360 degrees to gather a panoramic viewpoint of the room. The stop condition is queried at each timestep  $t$  with the visual memory and current observation if the agent is still facing towards the region, determined by checking that the detected region at the current timestep  $t$  is from the set of relevant regions,  $R_t \in R$ . Each time the stop condition is queried, the agent can choose whether to stop exploration—if it can already answer the question—or to continue exploring. If the

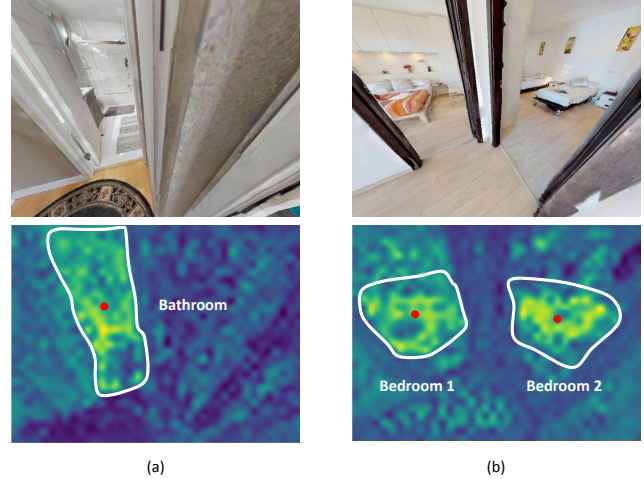


Figure 3. FAST-EQA leverages language-aligned features from AM-RADIO together with a SigLIP adaptor, to direct the agent toward the target regions. For queries such as (a) Bathroom and (b) Bedroom, the predicted heatmaps are thresholded to produce white contour segments, while the red dot indicates the contour centroid to step towards. This visualization illustrates how semantic grounding enables precise localization of task-relevant areas in the environment to guide exploration from global to local.

agent chooses to stop, relevant visual memory is retrieved for question-answering as detailed below. We use GPT-4o [16] as the VLM.

## 4.3. Relevant Memory Retrieval

Throughout exploration, we maintain a lightweight visual memory,  $\psi$ , that selectively retains the  $k$  most relevant visual observations seen so far based on the visual relevance score,  $rel_{vis}$  of the observation (4). This memory is stored per visual target,  $T_m$ .

$$\psi_{T_m} = TopK_{o_t \in O}(rel_{vis}(o_t, Q, T_m)) \quad (2)$$

The selective retention of only the most relevant visual observations results in a bounded memory that scales directly to the number of targets  $M$ , extracted from the question. Upon either querying of the stop condition or final question-answering, the stored visual observations are retrieved to assist the VLM in making a decision. To rank the relevance of each observation, we compute the visual relevance score  $rel_{vis}(o_t, Q, T_m)$ , by a weighted combination of a relevance score obtained from the CLIP [23] embedding model  $rel_{CLIP}(o_t, T_m)$ , and the relevance score given by querying a generative VLM [18]  $rel_{VLM}(o_t, Q)$ , where the weight is denoted  $\lambda$ .

$$rel_{vis}(o_t, Q, T_m) = \lambda \cdot rel_{CLIP}(o_t, T_m) + (1 - \lambda) \cdot rel_{VLM}(o_t, Q) \quad (3)$$



Figure 4. FAST-EQA employs a bounded memory system that allocates a dedicated visual memory for each target, retaining the  $k$  most relevant images (here,  $k = 3$ ). The overall memory footprint scales only with the number of targets and remains constant over time, even in long-horizon tasks.

The CLIP relevance score  $rel_{CLIP}$ , calculated per visual target, is obtained by computing the cosine embedding similarity between the 2D RGB observation  $o_t$  and the text of the target object,  $T_m$  extracted from the input question, where  $f$  is the CLIP encoder.

$$rel_{CLIP}(o_t, T_m) = sim(o_t, T_m) = \frac{f_{\text{text}}(T_m) \cdot f_{\text{img}}(o_t)}{\|f_{\text{text}}(T_m)\| \|f_{\text{img}}(o_t)\|} \quad (4)$$

The relevancy score from a generative VLM,  $rel_{VLM}$ , is question-dependent and computed by prompting the model  $f$  with the input question  $Q$  and the agent’s current observation  $o_t$ . Following [25], we use a prompt  $P$  that asks the model to judge whether the observation contains sufficient evidence to answer the question, and we take the probability assigned to the response “yes” as the relevancy score. For  $\hat{f}$ , we employ Prismatic-VLM as our generative backbone.

$$rel_{VLM}(o_t, Q) = p(\text{“yes”} | \hat{f}(P(Q), o_t)) \quad (5)$$

Overall, the CLIP relevance score can be viewed as a focused quantification of whether the observation is related to the visual target, while the VLM relevance score is a holistic scoring of the observation given the entire context of the question. This combination scoring allows the agent to retrieve observations that align with both the focused target goal and the question-answering goal. We tune the weight  $\lambda$  through a hyperparameter search (Supplementary Sec. 3).

#### 4.4. Chain-of-Thought Question Answering

In recent years, Chain-of-Thought (CoT) reasoning [28] has proven effective in natural language and vision–language tasks [39], but remains underexplored in embodied AI [14]. Multi-target question scenarios are one example that require reasoning over diverse observations, which naturally aligns with the compositional reasoning abilities of CoT. This motivates the use of CoT prompting for the final QA module.

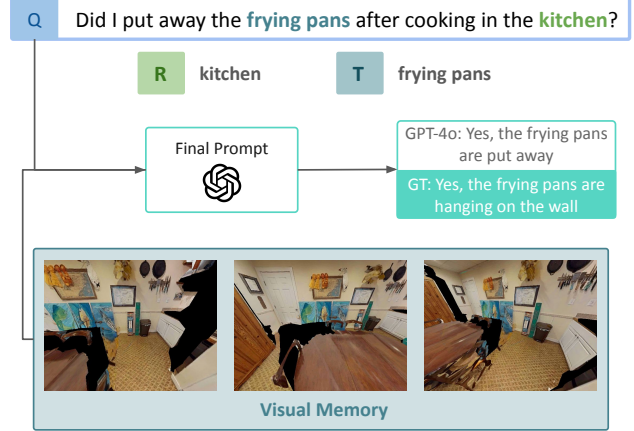


Figure 5. An example from EXPRESS-Bench illustrating how FAST-EQA identifies the relevant region  $R$  and target  $T$  from the question  $Q$ . It then explores the scene, and once the stopping condition is met, correctly generates the final answer from the retrieved visual memory.

Given the input question  $Q$  and the set of  $M \times k$  observations retained in memory, we prompt a multi-frame VLM at the end of exploration to generate an answer either as a letter for multiple-choice QA, or an open-ended text response (Figure 5) and asking the VLM to explicitly think step by step. This not only improves reasoning quality but also leads to more interpretable answers referring explicitly to the given visual observations, which existing baselines often lack. For the answering model, we use GPT-4o.

## 5. Evaluation and Results

We conduct experiments across diverse EQA benchmarks, keeping the definition of an exploration step consistent. Following the parameter set in [25], the agent is allowed to travel up to a maximum of 3 meters per translational step.

### 5.1. Comparison Across Benchmarks

To evaluate our approach, we run FAST-EQA on four benchmarks: HM-EQA [25], MT-HM3D [35], EXPRESS-Bench [17], and the released A-EQA 184-split of OpenEQA [22], which covers diverse question formats and question types. HM-EQA and MT-HM3D contain multiple-choice questions, with MT-HM3D focusing on questions comparing multiple targets. EXPRESS-Bench and OpenEQA consist of open-ended questions, with EXPRESS-Bench being the largest scale benchmark, containing over 2000 questions. To the best of our knowledge, we conduct the most comprehensive evaluation on diverse benchmarks compared to existing methods. Furthermore, we run three trials on each dataset, with temperature  $\tau = 0$ , to obtain error bounds for our results, which takes into consideration any stochasticity in our method. We note that existing methods do not

Table 1. Performance comparison across EQA benchmarks against SOTA baseline methods. \* indicates that the result is from a reproduced experiment reported by others. † indicates results are on full A-EQA split.

	HM-EQA		MT-HM3D		EXPRESS-Bench		A-EQA (184)	
	SR	Steps ( $\downarrow$ )	SR	Steps ( $\downarrow$ )	LLM Score	$E_{path}$ ( $\uparrow$ )	LLM-Match	$E_{path}$ ( $\uparrow$ )
GPT-4V (OpenEQA)[22]	-	-	-	-	-	-	41.8	7.5
Explore-EQA [25]	58.4	0.52	36.2*	0.64	-	-	46.9*	23.4
Graph-EQA [26]	63.5	<b>0.20</b>	45.63*	0.45	-	-	30.1* <sup>†</sup>	-
Memory-EQA [37]	63.4	0.40	<b>55.1</b>	<b>0.41</b>	-	-	36.8 <sup>†</sup>	-
Fine-EQA [17]	56.0	0.54	-	-	63.95	25.58	43.3 <sup>†</sup>	29.2
3D-Mem [33]	50.40	0.63	-	-	-	-	<b>52.6</b>	<b>42.0</b>
FAST-EQA (Ours)	<b>69.2</b> $\pm 0.7$	0.65 $\pm 0.01$	50.5 $\pm 0.3$	0.52 $\pm 0.01$	<b>68.7</b> $\pm 0.5$	<b>29.25</b> $\pm 0.55$	49.0 $\pm 1.7$	27.70 $\pm 1.70$

report error bounds, despite the usage of VLMs and LLMs often resulting in such variance.

For the multiple-choice datasets we report success rate (SR) and normalized steps (Steps) following [25, 37]. For the open-vocabulary questions, we report the LLM-Score following [17] for EXPRESS-Bench, LLM-Match ( $C$ ) following [22] for OpenEQA, and the corresponding path efficiency metric calculated with success weighted-by-path-length (SPL) using geodesic distance,  $E_{path}$ . For more details on metrics, refer to supplementary material.

**Results.** The performance of FAST-EQA against a number of SOTA baselines is shown in Table 1. We observe that our method surpasses existing methods on question-answering accuracy for both HM-EQA and EXPRESS-Bench, highlighting its ability to handle both multiple-choice and open-ended questions. In particular, we achieve a 9% improvement over the next-best method on HM-EQA and a 7% improvement over the SOTA method on EXPRESS-Bench, marking high performance on a larger-scale dataset. On the A-EQA subset of Open-EQA, we achieve results comparable to SOTA. While multi-target questions still pose a challenge, FAST-EQA is also able to achieve competitive performance against existing methods with relatively low variance while having a lightweight, bounded memory.

We further note that our method prioritizes question-answering accuracy against the step length tradeoff, allowing for more thorough exploration and gathering of diverse views before generating an answer. Simultaneously, the inference time speed-up we achieve maintains overall efficiency.

## 5.2. Inference Time Experiments

For EQA systems, the ability to operate in real-time is essential for eventual deployment on physical robots. Existing EQA methods however, do not report these evaluations. To address this gap, we benchmark and compare the infer-

	Average Step Time (s)
Fine-EQA	2.94 $\pm$ 0.14
Explore-EQA	4.90 $\pm$ 0.26
3D-Mem	14.53 $\pm$ 0.25
Memory-EQA	15.89 $\pm$ 0.65
FAST-EQA (Ours)	<b>2.54 <math>\pm</math> 0.05</b>

Table 2. Average step time, in seconds, of FAST-EQA compared to other methods on a subset of questions from HM-EQA

ence time of our method with existing open-source methods. Inference time experiments are run for three trials on a fixed subset of 100 questions from the HM-EQA dataset. All experiments are run on a single NVIDIA H100 Tensor Core GPU with 80GB RAM. We report the average time (in seconds) for each exploration step, which includes the time taken for querying large models.

Table 2 shows the inference time results. Our method achieves significantly faster inference times than existing EQA methods, resulting in a 13.6% speedup over the next fastest method. Furthermore, FAST-EQA achieves consistently fast run times, exhibiting markedly lower variance across trials than competing methods.

## 5.3. Ablation Studies

We further conduct a series of ablation studies to isolate the contributions of individual design choices in our framework. We run these experiments on a fixed sampled subset of 100 questions from HM-EQA, and summarize the results in Table 3. Each ablation removes one design choice at a time and compares against the full FAST-EQA model.

**Without Doorways/Openings as Frontiers.** To evaluate the role of structural scene analysis in exploration, we replace our doorway- and opening-guided strategy with stan-

	SR	Steps ( $\downarrow$ )
FAST-EQA w/o Doorway Frontiers	67.0	0.64
FAST-EQA w/o CLIP Scoring	69.0	0.63
FAST-EQA w/o VLM Scoring	72.0	0.63
FAST-EQA w/o CoT Reasoning	72.0	0.60
FAST-EQA (Ours)	<b>76.0</b>	0.65

Table 3. Ablation of various components of our EQA system on a fixed subset of HM-EQA

standard frontier-based exploration (FBE) [31], where frontiers are defined purely as boundaries between explored and unexplored space.

**Without VLM Relevancy.** To quantify the effect of incorporating VLM confidence into scoring observation relevancy, we experiment with scoring the relevance of image observations solely based on the CLIP similarity to the visual target by setting  $\lambda = 1$  in Equation 3.

**Without CLIP Relevancy.** Similarly, we ablate the inclusion of CLIP similarity scores when scoring observation relevancy. We score the relevance of image observations only based on querying the VLM for whether the image is relevant by setting  $\lambda = 0$  in Equation 3.

**Without Chain-of-Thought Prompting.** We ablate the effect of using chain-of-thought prompting for the question-answering module by modifying the prompt to remove explicit instructions to “think step by step”.

**Results.** As shown in Table 3, exclusion of doorway-guided frontier exploration, VLM relevancy scoring, CLIP scoring, or chain-of-thought prompting all result in a performance drop for overall question-answering success, highlighting the importance of each component. Furthermore the semantic frontier-guided exploration leveraging doorways and openings has the largest effect for improving the performance, as changing to standard FBE results in a 9% drop in success rate. This demonstrates that semantic exploration guided by informative structural information is particularly critical for enabling efficient and accurate embodied QA. More broadly, these ablations validate our main contribution: combining semantics-aware exploration together with multimodal relevancy scoring and reasoning mechanisms yields substantial gains for accurate question-answering without compromising computation time.

## 6. Conclusion

We introduce FAST-EQA, an active exploration QA framework that couples semantically-guided global and local exploration with a bounded memory that scales only with targets, making it well-suited for single as well as multi-target tasks. Our exploration strategy is focused on intuitively

maximizing exploration coverage by leveraging scene morphology to transition between semantic regions, allowing for thorough search and verification of targets. The framework supports both multiple-choice and open-ended question settings, as well as single- and multi-target questions. Across major EQA benchmarks, FAST-EQA achieves state-of-the-art results on HMEQA and EXPRESS-Bench, while remaining competitive on OpenEQA and MT-HM3D. Coupled with this performance, our method delivers fast inference with low memory usage, enabling real-time deployment for embodied agents.

## 6.1. Limitations and Future Work

While FAST-EQA demonstrates strong performance across multiple EQA benchmarks, several limitations remain. First, the system’s effectiveness is bounded by the spatial reasoning capabilities of the current VLMs used. As highlighted by [32], even state-of-the-art MLLMs struggle with fine-grained spatial understanding, which in turn constrains the quality of reasoning in complex environments. Second, we observe that the VLM (e.g., GPT-4o) used for reasoning exhibits variance across runs, sometimes producing inconsistent answers. In certain cases, this variability can cause the agent to terminate prematurely, limiting robustness and reliability in long-horizon tasks.

Looking ahead, an exciting direction is the development of latent visual memory representations that can efficiently capture and compress both short-term and long-term scene information. Such memory mechanisms could enable more consistent reasoning, reduce reliance on repetitive VLM queries, and improve the agent’s ability to handle extended exploration tasks. We believe integrating structured spatial memory with language-driven reasoning offers a promising pathway toward more generalizable and reliable embodied agents.

## References

- [1] Abrar Anwar, John Welsh, Joydeep Biswas, Soha Pouya, and Yan Chang. Rememr: Building and reasoning over long-horizon spatio-temporal memory for robot navigation. In *arXiv preprint arXiv:2409.13682*, 2024. 3
- [2] Ctălina Cangea, Eugene Belilovsky, Pietro Li, and Aaron Courville. Videonavqa: Bridging the gap between visual and embodied question answering. In *Proceedings of the 30th British Machine Vision Conference (BMVC)*, 2019. Spotlight presentation at Visually Grounded Interaction Learning (ViGIL) Workshop, NeurIPS 2019. 3
- [3] Vincent Cartillier, Zhile Ren, Neha Jain, Stefan Lee, Irfan Essa, and Dhruv Batra. Semantic mapnet: Building allocentric semantic maps and representations from egocentric views. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 964–972, 2021. 3
- [4] Devendra Singh Chaplot, Dhiraj Gandhi, Abhinav Gupta, and Ruslan Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *Advances in Neural*

- Information Processing Systems (NeurIPS)*, 2020. Winner of the CVPR 2020 Habitat ObjectNav Challenge. 3
- [5] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14455–14465, Vancouver, BC, Canada, 2024. 3
- [6] Kai Cheng, Zhengyuan Li, Xingpeng Sun, Byung-Cheol Min, Amrit Singh Bedi, and Aniket Bera. Efficienteqa: An efficient approach for open vocabulary embodied question answering. In *Proceedings of the IEEE/CVF International Conference on Robotics and Automation (ICRA)*, 2025. Accepted for publication. 3
- [7] Yinpei Dai, Run Peng, Sikai Li, and Joyce Chai. Think, act, and ask: Open-world interactive personalized robot navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, page TBD, Singapore, 2024. 3
- [8] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–10, Salt Lake City, UT, USA, 2018. 3
- [9] Abhishek Das, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Neural modular control for embodied question answering. In *Proceedings of The 2nd Conference on Robot Learning (CoRL)*, pages 53–62, Zürich, Switzerland, 2018. 3
- [10] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. Palm-e: An embodied multimodal language model. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pages 8469–8488, Honolulu, HI, USA, 2023. 3
- [11] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, pages 226–231, 1996. 5
- [12] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4089–4098, Salt Lake City, UT, USA, 2018. 3
- [13] Qiao Gu, Alihusein Kuwajerwala, Sacha Morin, Krishna Murthy Jatavallabhula, Bipasha Sen, Aditya Agarwal, Corban Rivera, William Paul, Kirsty Ellis, Rama Chellappa, Chuang Gan, Celso Miguel de Melo, Joshua B. Tenenbaum, Antonio Torralba, Florian Shkurti, and Liam Paull. Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5021–5028, Karlsruhe, Germany, 2024. 3
- [14] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Inner monologue: Embodied reasoning through planning with language models. In *Conference on Robot Learning (CoRL)*, 2022. 6
- [15] Nathan Hughes, Yun Chang, and Luca Carlone. Hydra: A real-time spatial perception system for 3d scene graph construction and optimization. In *Proceedings of Robotics: Science and Systems (RSS)*, New York City, NY, USA, 2022. 2
- [16] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 5
- [17] Kaixuan Jiang, Yang Liu, Weixing Chen, Jingzhou Luo, Ziliang Chen, Ling Pan, Guanbin Li, and Liang Lin. Beyond the destination: A novel benchmark for exploration-aware embodied question answering. In *arXiv preprint arXiv:2503.11117*, 2025. 2, 3, 6, 7
- [18] 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 *Forty-first International Conference on Machine Learning*, 2024. 5
- [19] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 2
- [20] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. 2
- [21] Joel Loo, Zhanxin Wu, and David Hsu. Open scene graphs for open-world object-goal navigation. *arXiv preprint arXiv:2508.04678*, 2025. Also to appear in \*The International Journal of Robotics Research\*. 2
- [22] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul McVay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent-Pierre Berges, Shiqi Zhang, Pulkit Agrawal, Dhruv Batra, Yonatan Bisk, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Alexander Sax, and Aravind Rajeswaran. Openeqa: Embodied question answering in the era of foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 6, 7
- [23] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmlR, 2021. 5
- [24] Mike Ranzinger, Greg Heinrich, Jan Kautz, and Pavlo Molchanov. Am-radio: Agglomerative vision foundation model reduce all domains into one. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12490–12500, 2024. 5

- [25] Allen Z. Ren, Jaden Clark, Anushri Dixit, Masha Itkina, Anirudha Majumdar, and Dorsa Sadigh. Explore until confident: Efficient exploration for embodied question answering. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024. 2, 3, 4, 6, 7
- [26] Saumya Saxena, Blake Buchanan, Chris Paxton, Bingqing Chen, Narunas Vaskevicius, Luigi Palmieri, Jonathan Francis, and Oliver Kroemer. Grapheqa: Using 3d semantic scene graphs for real-time embodied question answering. *arXiv*, 2024. Accepted for publication. 3, 7
- [27] Weizhen Wang, Chenda Duan, Zhenghao Peng, Yuxin Liu, and Bolei Zhou. Embodied scene understanding for vision language models via metavqa. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22453–22464, 2025. 2
- [28] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 6
- [29] Abdelrhman Werby, Chenguang Huang, Martin Büchner, Abhinav Valada, and Wolfram Burgard. Hierarchical open-vocabulary 3d scene graphs for language-grounded robot navigation. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024. 2
- [30] Erik Wijmans, Samyak Datta, Oleksandr Maksymets, Abhishek Das, Georgia Gkioxari, Stefan Lee, Irfan Essa, Devi Parikh, and Dhruv Batra. Embodied question answering in photorealistic environments with point cloud perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6659–6668, Long Beach, CA, USA, 2019. 3
- [31] Brian Yamauchi. A frontier-based approach for autonomous exploration. In *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, pages 146–151, Monterey, CA, USA, 1997. 8
- [32] Jihan Yang, Shusheng Yang, Anjali W. Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in space: How multimodal large language models see, remember, and recall spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10632–10643, 2025. 8
- [33] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17294–17303, 2025. 2, 3, 7
- [34] Naoki Yokoyama, Sehoon Ha, Dhruv Batra, Jiuguang Wang, and Bernadette Bucher. Vlfm: Vision-language frontier maps for zero-shot semantic navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 42–48, Yokohama, Japan, 2024. 3
- [35] Licheng Yu, Xinlei Chen, Georgia Gkioxari, Mohit Bansal, Tamara L. Berg, and Dhruv Batra. Multi-target embodied question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4990–4999, 2019. 3, 6
- [36] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In *CVPR*, 2017. 4
- [37] Mingliang Zhai, Zhi Gao, Yuwei Wu, and Yunde Jia. Memory-centric embodied question answer. In *arXiv preprint arXiv:2505.13948*, 2025. 2, 3, 4, 7
- [38] Jiazhao Zhang, Kunyu Wang, Rongtao Xu, Gengze Zhou, Yicong Hong, Xiaomeng Fang, Qi Wu, Zhizheng Zhang, and He Wang. Navid: Video-based vlm plans the next step for vision-and-language navigation. *arXiv preprint arXiv:2402.15852*, 2024. 2
- [39] Zhuosheng Zhang, Aston Zhang Xu, Zhiyuan Li, Rui Wang, and Zhifang Liu. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*, 2023. 6