

Expand Your SCOPE: Semantic Cognition over Potential-Based Exploration for Embodied Visual Navigation

Ningnan Wang, Weihuang Chen*, Liming Chen, Haoxuan Ji,
Zhongyu Guo, Xuchong Zhang, Hongbin Sun

State Key Laboratory of Human-Machine Hybrid Augmented Intelligence
National Engineering Research Center for Visual Information and Applications
Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University

Abstract

Embodied visual navigation remains a challenging task, as agents must explore unknown environments with limited knowledge. Existing zero-shot studies have shown that incorporating memory mechanisms to support goal-directed behavior can improve long-horizon planning performance. However, they overlook visual frontier boundaries, which fundamentally dictate future trajectories and observations, and fall short of inferring the relationship between partial visual observations and navigation goals. In this paper, we propose Semantic Cognition Over Potential-based Exploration (SCOPE), a zero-shot framework that explicitly leverages frontier information to drive potential-based exploration, enabling more informed and goal-relevant decisions. SCOPE estimates exploration potential with a Vision-Language Model and organizes it into a spatio-temporal potential graph, capturing boundary dynamics to support long-horizon planning. In addition, SCOPE incorporates a self-reconsideration mechanism that revisits and refines prior decisions, enhancing reliability and reducing overconfident errors. Experimental results on two diverse embodied navigation tasks show that SCOPE outperforms state-of-the-art baselines by 4.6% in accuracy. Further analysis demonstrates that its core components lead to improved calibration, stronger generalization, and higher decision quality.

Code — <https://github.com/mrwangyou/SCOPE>

Extended version — <https://arxiv.org/abs/2511.08935>

1 Introduction

Embodied visual navigation (EVN) is a fundamental task in embodied intelligence, requiring agents to autonomously plan paths toward designated goals in previously unseen environments. This capability holds significant promise for real-world applications such as smart home robotics, disaster response, and deep space exploration, and has attracted growing research interest in recent years (Janny et al. 2025; Ehsani et al. 2024; Wang et al. 2025a).

Despite its utility, embodied visual navigation remains a significant challenge. Specifically, agents must operate

in unfamiliar environments, continuously inferring associations between partial visual observations and navigation goals (Li et al. 2025). Furthermore, modern tasks increasingly involve multi-modal goal specifications, ranging from natural language commands to reference images or object categories, while demanding robust long-horizon planning (Khanna et al. 2024b). These requirements place a premium on methods that combine strong semantic reasoning capabilities with reliable planning over extended sequences, substantially increasing decision-making complexity.

Strong exploration ability is essential to meeting these demands. Agents must efficiently acquire new information from current observations while effectively organizing and leveraging previously gathered knowledge. Most existing zero-shot navigation methods enhance exploration by incorporating memory mechanisms to support goal-directed behavior. For instance, ConceptGraph (Gu et al. 2024) constructs a graph-structured, object-centric representation of 3D scenes by fusing outputs from 2D foundation models, offering a compact memory for spatial reasoning. Similarly, VLFM (Yokoyama et al. 2024) builds occupancy maps and employs a Vision-Language Model (VLM) to generate a language-grounded value map over frontiers. While such approaches have shown promise, they face two main limitations: 1) VLMs, though proficient in multi-modal reasoning, are not specifically trained on dense 3D inputs, limiting their understanding of complex spatial relationships; and 2) memory systems that focus solely on visited regions often fail to connect current observations to unvisited goal-relevant areas, leaving agents without clear exploration cues.

To address these limitations, recent methods have begun to incorporate information from unexplored regions. Nav-iFormer (Xie et al. 2025) enhances object-goal navigation by integrating spatial layout encoding, agent pose trajectories, and a passable frontier map using spatio-temporal attention. 3D-Mem (Yang et al. 2025), the state-of-the-art method, proposes a snapshot-based 3D memory architecture that stores multi-view representations of explored areas alongside frontier snapshots of unexplored regions. However, these approaches still fall short in fully exploiting the *potential* of visual frontiers. Far from being passive boundaries, frontiers act as critical bridges linking the agent’s current state to potentially goal-relevant areas, shaping both fu-

*Corresponding author (chenwh@xjtu.edu.cn).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ture trajectories and the observations an agent can acquire.

This paper aims to propose a zero-shot method that fully exploits potential for effective exploration. Specifically, we introduce Semantic Cognition Over Potential-based Exploration (SCOPE), a new approach that fully exploits the informational value of visual frontier boundaries to enhance long-horizon navigation. The key idea is to treat frontiers as semantic cognition cues that enhance scene understanding, memory construction, and decision reliability. SCOPE is composed of three key components. Firstly, a *frontier-level potential estimator* evaluates the semantic relevance of unexplored regions with respect to the navigation goals from raw visual inputs. Secondly, a *potential graph* constructs and dynamically updates a latent graph that propagates spatial-semantic utility as structured memory, guiding the VLM in decision-making. Finally, a *self-reconsideration mechanism* revisits and refines initial decisions, preventing premature commitments, reducing hallucinations in VLM reasoning, and improving environmental understanding and planning accuracy. These components work in synergy to foster more deliberate and informed navigation, leading to stronger performance across complex tasks.

Our main contributions are as follows:

1. We propose a zero-shot embodied navigation framework that explicitly prioritizes frontier information as a primary exploration cue, introducing a new paradigm for information-guided exploration.
2. We design a potential graph that efficiently integrates spatio-temporal information with frontier potential estimation, enabling structured environment representation and informed global planning.
3. We incorporate a self-reconsideration mechanism, encouraging iterative refinement and contextual alignment, leading to more reliable and deliberate actions.
4. We demonstrate statistically significant improvements over state-of-the-art baselines, achieving **4.6%** accuracy gains on two challenging embodied navigation benchmarks (GOAT-Bench and A-EQA).

Further analysis reveals that our method delivers statistically significant gains, demonstrates better-calibrated decision confidence, and benefits from each proposed component. We release the complete implementation of SCOPE for transparency and reproducibility.

2 Related Work

EVN Tasks EVN tasks involve a wide range of task formulations. ObjectNav directs agents to locate an instance of a specified semantic category (Zhang et al. 2024; Khanna et al. 2024a; Yin et al. 2024; Zhang et al. 2025). ImageNav requires finding an object that visually matches a given reference image (Barsellotti et al. 2024; Li et al. 2025; Lei et al. 2024, 2025). Vision-Language Navigation (VLN) requires agents to ground semantic descriptions into spatial targets (Perincherry, Krantz, and Lee 2025; Wang et al. 2024b,a; An et al. 2024). Beyond these forms, Embodied Question Answering (EQA) requires agents to explore environments to collect information needed to answer natural-language questions (Ren et al. 2024a; Wang et al. 2025b;

Dang et al. 2025). Emerging multi-stage (Song et al. 2025) and multi-step (Yu et al. 2024) navigation settings further amplify complexity by demanding extended reasoning, persistent memory, and adaptive long-horizon planning, which are core capabilities forming the focus of this work.

Zero-shot EVN Methods Zero-shot EVN methods aim to enable agents to operate in unknown environments without task-specific fine-tuning, instead leveraging generalizable knowledge learned from large-scale pretraining. These approaches have recently gained significant attention and generally progress along two main directions. The first focuses on improving scene representation and memory organization to better support spatial reasoning and long-term planning (Liu, Wang, and Yang 2024; Yin et al. 2025; Zhao et al. 2024; Li et al. 2024; Lin et al. 2025b; Xie et al. 2025). The second emphasizes enhancing agent reasoning through advanced decision-making paradigms, such as chain-of-thought prompting (Lin et al. 2025a; Kong et al. 2024), diffusion-based decision policies (Ren et al. 2025; Zhang, Tang, and Yan 2024), or world models (Bar et al. 2025; Pan et al. 2025) for simulating and planning in dynamic environments. While these methods often excel at modeling visited regions, they lack structured and adaptive mechanisms for prioritizing which unexplored areas to visit next. In contrast, our approach treats frontiers as a primary semantic and structural cue for exploration, enabling more targeted, informed, and goal-aligned navigation.

Frontier-related EVN Methods Several recent works have explicitly incorporated frontier information into EVN policies. VLFM (Yokoyama et al. 2024) derives semantically valuable cues from occupancy maps to guide exploration toward unseen object targets. Cui et al. (2024) represent visible but unexplored frontiers with “ghost nodes” to improve exploration awareness in navigation agents. However, both methods rely mainly on simple geometric frontier features, limiting adaptability to diverse navigation contexts. 3D-Mem (Yang et al. 2025) stores snapshots of unexplored regions, enabling VLM-based reasoning over frontier areas. While this extends frontier usage beyond geometric estimation, the approach largely passes raw snapshots to VLMs and still depends on heuristic frontier selection, leaving much of the potential of frontiers untapped. In contrast, our method introduces a potential-based frontier exploration paradigm that explicitly estimates the semantic utility of each frontier, integrates it into a structured spatio-temporal potential graph, and refines decisions via self-reconsideration.

3 Methodology

3.1 Preliminary

EVN tasks require an agent to traverse an unknown environment to reach goals specified by point coordinates, object references, or natural language instructions. Formally, at each time step t , the agent selects an action a^t using a navigation policy π_θ , conditioned on the current goal q , the agent’s state s^t , and its accumulated knowledge \mathcal{K}^t :

$$a^t = \pi_\theta(q, s^t, \mathcal{K}^t), \quad (1)$$

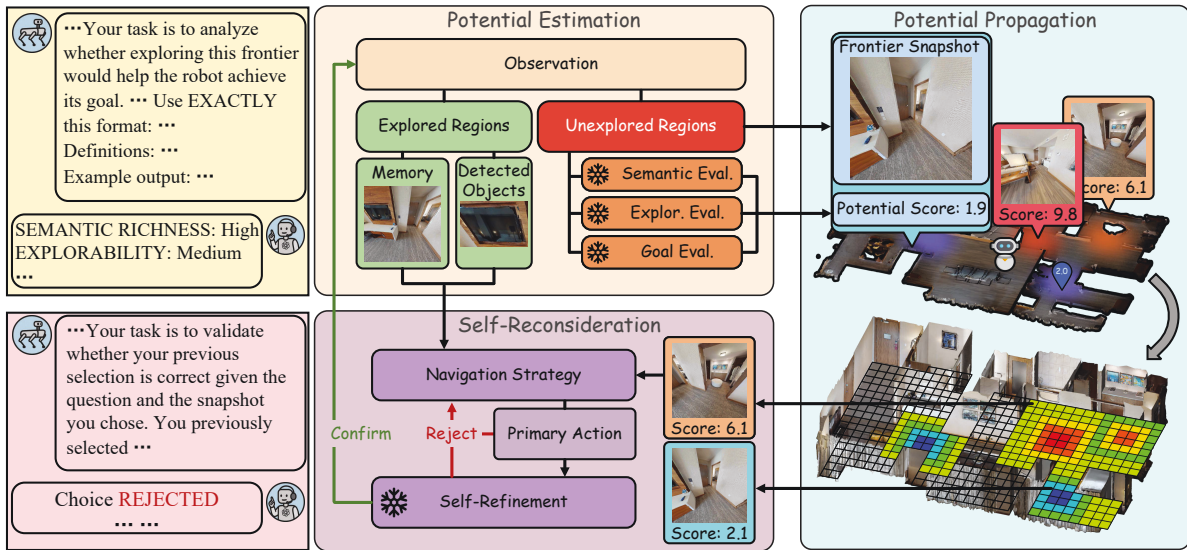


Figure 1: Overview of SCOPE. The agent predicts frontier utility via a VLM-based estimator and encodes it into a structured potential graph for spatiotemporal reasoning. Action decisions are guided by this graph and further reconsideration through a self-refinement module to avoid impulsive errors.

where $\mathcal{K}^t = \{I^1, I^2, \dots, I^t\}$ denotes the sequence of historical observations up to time t . Upon executing a^t , the agent transitions to a new state s^{t+1} and receives new sensory input. If the goal is deemed satisfied, the episode terminates, the task is considered finished, and a goal checker would examine whether the goal is achieved. Otherwise, the agent continues exploring to collect additional information.

At each step, the agent captures a panoramic set of ego-centric observations:

$$I^t = \{I_1^t, I_2^t, \dots\}, \quad (2)$$

where each I_i^t is an RGB-D image from a discrete directional view. From these observations, the agent constructs an *explored memory* \mathcal{E}^t , which encodes both the raw visual snapshot and detected semantic entities:

$$\mathcal{E}^t = \{I^t, O^t\}, \quad (3)$$

with O^t representing object-level annotations or region semantics extracted from I^t .

In parallel, the agent identifies *frontier regions*—boundaries between explored and unobserved space. These frontiers are represented as spatial clusters $\mathcal{F}^t = \{\mathcal{F}_1^t, \mathcal{F}_2^t, \dots\}$ at the periphery of explored areas, each associated with a representative visual snapshot.

We define the structured environmental knowledge $\tilde{\mathcal{K}}^t$ as the union of all previously explored content and current frontier candidates:

$$\tilde{\mathcal{K}}^t = \underbrace{\{\mathcal{E}^1, \mathcal{E}^2, \dots, \mathcal{E}^t\}}_{\text{Explored regions}} \cup \underbrace{\mathcal{F}^t}_{\text{Frontier regions}}. \quad (4)$$

Therefore, the environment is partitioned into two complementary components: explored regions, which contain visited areas with available visual and semantic information, and frontier regions, which represent unobserved boundary

areas serving as gateways to new information. Our framework explicitly models and exploits frontier regions as informative cues for decision-making, enabling goal-directed and efficient exploration in unknown areas.

3.2 Framework

The overall framework of SCOPE is illustrated in Fig. 1. The agent first captures observations encompassing both explored and frontier regions. To fully leverage information from frontier observations, we introduce a frontier-level potential estimator built upon open-domain VLMs, which integrates visual features from frontier snapshots with task-specific contextual cues to predict the semantic utility of each unexplored region. These predicted utilities are then incorporated into a potential graph that encodes, propagates, and updates utility estimates across spatial and temporal dimensions, enabling the agent to prioritize high-utility frontiers, revisit semantically valuable areas, and avoid low-utility regions to improve efficiency. The enriched frontier snapshots, now carrying propagated utility estimates, are fed back into the agent’s decision-making loop to guide more informed action selection. To further ensure planning stability and mitigate erroneous or impulsive choices, we incorporate a self-reconsideration mechanism that performs internal consistency checks on candidate actions before execution, allowing the agent to correct flawed reasoning and refine its strategy based on accumulated knowledge.

3.3 Frontier-level Potential Estimation

To assess the utility of each unexplored frontier, we propose a frontier-level potential estimator grounded in pretrained VLMs. This estimator jointly leverages visual features from frontier snapshots and goal-specific prompts to infer semantic utility, effectively transforming the VLM into a general-

izable semantic oracle across scenes and tasks.

Specifically, the potential of an unexplored region can be assessed from three complementary perspectives. **Semantic Richness:** Regions with high semantic density would efficiently earn more information and are more likely to include goal-relevant content. **Explorability:** A valuable frontier should not only be informative itself but also serve as a gateway to other unseen areas. **Goal Relevance:** Semantically related objects often co-occur in the same place. A frontier showing partial evidence of such a semantic cluster increases the likelihood that the goal target lies nearby.

Formally, for each frontier snapshot $\mathcal{F}_i^t \in \mathcal{F}^t$, we define:

$$\mathbf{p}_i^t = [p_{i,sem}^t, p_{i,explore}^t, p_{i,goal}^t] = f_{VLM}(\mathcal{F}_i^t, q), \quad (5)$$

where f_{VLM} is a prompting-based interface to a VLM, and q is the task goal. The model also provides an aggregated scalar score:

$$P_i^t = \text{Aggregate}(\mathbf{p}_i^t), \quad (6)$$

indicating the overall utility of each frontier snapshot.

This component transforms the VLM into a goal-aware semantic oracle that generalizes across tasks and scenes without additional training. It grounds frontier evaluation in rich pretrained priors and enables interpretable, context-aware decision-making.

3.4 Structured Memory over Potential Graph

To enable long-horizon planning and mitigate aimless exploration, we introduce the potential graph, a spatial memory structure that encodes the estimated utility of different regions across the navigable environment. This graph captures past frontier observations, propagates their potential across space, and enables strategic revisitation or avoidance.

We discretize the environment into a 2D grid $\mathcal{G} = \{v_{m,n}\}_{m,n}$, where each cell $v_{m,n}$ corresponds to a spatial location and maintains the properties including a potential score $P_{m,n}$, a visit count $n_{m,n}$, and semantic attributes $p_{m,n}^{sem}$, $p_{m,n}^{explore}$, $p_{m,n}^{goal}$.

When a frontier \mathcal{F}_i is observed at position $\mathbf{p}_{m,n}$ with estimated potentials $s_i = (p_i^{sem}, p_i^{explore}, p_i^{goal}, P_i)$, its information is propagated to nearby graph nodes within a fixed radius R . For each affected cell $v_{m,n}$ with spatial coordinate $\mathbf{p}_{m,n}$, we perform the following weighted update:

$$P_{m,n}^t \leftarrow (1 - \alpha_{m,n}^t) \cdot P_{m,n}^{t-1} + \alpha_{m,n}^t \cdot P_i, \quad (7)$$

$$\alpha_{m,n}^t = \max\left(0, 1 - \frac{\|\mathbf{p}_{m,n} - \mathbf{p}_{\mathcal{F}_i}\|}{R}\right), \quad (8)$$

where $\alpha_{m,n}^t \in [0, 1]$ determines the influence weight based on the spatial distance between the frontier and the cell. A similar weighted update is applied to semantic components p^{sem} , $p^{explore}$, p^{goal} .

To encourage wide coverage and avoid local cycles, we compute a frontier exploration value for each node by combining its spatial potential and semantic relevance, while penalizing frequent revisits. Specifically, a weighted sum of the node’s potential score, along with its estimated semantic richness, explorability, and goal relevance

assessments, is computed using a weight vector $\mathbf{w} = [\omega_{pot}, \omega_{sem}, \omega_{explore}, \omega_{goal}]$. This semantic score is then scaled by an inverse visitation factor:

$$E_{m,n} = (\omega_{pot} \cdot P_{m,n} + \omega_{sem} \cdot p_{m,n}^{sem} + \omega_{explore} \cdot p_{m,n}^{explore} + \omega_{goal} \cdot p_{m,n}^{goal}) \cdot \frac{1}{1 + \gamma \cdot n_{m,n}}, \quad (9)$$

where $n_{m,n}$ is the visit count of node $v_{m,n}$ and $\gamma > 0$ is a decay coefficient controlling the strength of revisitation penalty. This scoring encourages selection of novel yet promising areas. During navigation, the agent is prompted with the estimated potential score and consults the active frontier scores to plan its next move, facilitating strategic long-term exploration and reducing inefficient behaviors.

3.5 Self-reconsideration for Robust Action Validation

To avoid premature and inaccurate decisions, we incorporate a self-refinement technique that validates candidate actions before execution. This technique assesses whether a selected memory snapshot truly contains the goal-relevant content.

Refinement Trigger Given a goal query q and structured memory $\tilde{\mathcal{K}}^t$, the agent first produces a primary action:

$$a_t^{(0)} = \pi_\theta(q, \tilde{\mathcal{K}}^t), \quad (10)$$

which corresponds to either selecting a memory snapshot and object pair or proposing a frontier snapshot to explore. If this action selects a memory snapshot-object pair, refinement is triggered:

$$\delta(a_t^{(0)}) = \begin{cases} 1 & \text{if } a_t^{(0)} = \text{memory}(i, o), \\ 0 & \text{if } a_t^{(0)} = \text{frontier}(j). \end{cases} \quad (11)$$

Validation Process If refinement is triggered, a secondary decision is made by invoking a validation model g_ϕ , using a VLM, to assess whether the selected snapshot actually satisfies the task goal:

$$r_t^{(0)} = g_\phi(q, \mathcal{E}, o), \quad (12)$$

where $r_t \in \{\text{CONFIRM}, \text{REJECT}\}$. The input to g_ϕ includes the task goal q , the selected snapshot $\mathcal{E} = \langle I_{\mathcal{E}}, O_{\mathcal{E}} \rangle$, and the object $o \in O$ that was chosen as the goal candidate.

If $r_t^{(0)} = \text{CONFIRM}$, the original action $a_t^{(0)}$ is accepted as the final action $a_t = a_t^{(0)}$. If $r_t^{(0)} = \text{REJECT}$, the agent discards the selection and re-consults the main policy:

$$a_t^{(1)} = \pi_\theta(q, \tilde{\mathcal{K}}^t, a_t^{(0)}), \quad (13)$$

possibly repeating this process with updated guidance until either a valid action is confirmed or a retry limit is reached.

Final Action Selection Overall, the agent’s action at time t is determined as:

$$a_t = \text{SR}(a_t^{(0)}), \quad (14)$$

$$\text{SR}(a_t^{(i)}) = \begin{cases} a_t^{(i)} & \text{if } \delta(a_t^{(i)}) = 0 \text{ or } r_t^{(i)} = \text{CONFIRM}, \\ a_t^{(i+1)} & \text{if } r_t^{(i)} = \text{REJECT}. \end{cases} \quad (15)$$

This forms a corrective feedback loop that allows the agent to reassess decisions and avoid overconfident mistakes, improving robustness in ambiguous scenarios.

4 Experiments

4.1 Experiment Settings

Benchmarks To enable a broad evaluation across key embodied reasoning challenges, we leverage two complementary benchmarks covering distinct tasks: GOAT-Bench (Khanna et al. 2024b), which targets goal-conditioned navigation grounded in diverse language instructions, and A-EQA (Majumdar et al. 2024), which emphasizes embodied question answering that demands multi-step spatial and semantic reasoning. These benchmarks reflect distinct task formulations—spatial navigation and semantic reasoning—making them ideal to test the generality and robustness of our method.

Metrics We adopt four established metrics to comprehensively evaluate our agent’s performance across navigation and embodied question answering tasks. For GOAT-Bench, we follow prior work in using Success Rate (SR) and Success weighted by Path Length (SPL). SR measures the percentage of subtasks in which the agent stops within 1 meter of the correct goal object, while SPL accounts for both success and trajectory efficiency by penalizing unnecessarily long paths. For A-EQA, we evaluate answer quality using Correctness, based on the LLM-Match score—a large language model-based metric that aligns closely with human judgment. To assess information-seeking behavior, we also report Efficiency, defined as the ratio between the shortest possible path to the relevant information and the actual path the agent traverses before answering. Together, these four metrics provide a holistic view of agent performance in terms of task success, path optimality, answer quality, and exploration efficiency.

Baselines To ensure a fair and comprehensive comparison, we benchmark our method against 3D-Mem (Yang et al. 2025), the state-of-the-art framework which introduces a spatial memory module to support long-term, goal-directed reasoning. By persistently storing visual and semantic information throughout exploration, 3D-Mem enables effective retrieval and reasoning over past observations, making it a strong representative baseline for open-ended embodied navigation. In addition to 3D-Mem, we evaluate against several recent and competitive methods. For example, Explore-EQA (Ren et al. 2024b) integrates semantic-guided frontier exploration with a calibrated stopping mechanism based on conformal prediction. TANGO (Ziliotto et al. 2025) is a training-free framework that assembles pre-trained vision and language models into agents capable of solving open-world tasks, demonstrating strong generalization without task-specific training. MTU3D (Zhu et al. 2025) proposes a unified framework that bridges visual grounding and exploration by jointly optimizing scene understanding and navigation

GOAT-Bench				
Method	FR	FI	SR	SPL
SenseAct-NN Skill Chain	–	–	29.5	11.3
SenseAct-NN Monolithic	–	–	12.3	6.8
CoW (2023)	–	–	16.1	10.4
DyNaVLM (2025)	–	–	25.5	10.2
VLMnav (2024)	–	–	16.3	6.6
Modular GOAT (2024b)	–	–	24.9	17.2
Explore-EQA (ICRA’2024b)	✓	–	55.0	37.9
CG + Frontier (ICRA’2024)	✓	–	61.5	41.3
TANGO (CVPR’2025)	–	–	32.1	16.5
MTU3D (ICCV’2025)	✓	–	47.2	27.7
3D-Mem (CVPR’2025)	✓	–	69.1	48.9
SCOPE (Ours)	✓	✓	73.7	53.5

Table 1: Performance comparison on **GOAT-Bench**. **Bold** indicates the best performance.

gation policies within a memory-augmented transformer architecture. Together, these baselines represent diverse and advanced strategies for perception, memory, and reasoning, providing strong reference points for evaluating the effectiveness of our proposed approach.

4.2 Results

We evaluate SCOPE on vision-language navigation benchmark GOAT-Bench, and the result is presented in Table 1 and Fig. 2. FR denotes methods that incorporate frontier representations, while FI refers to methods that fully leverage frontier information for decision-making. SCOPE achieves the highest SR on GOAT-Bench, improving over the previous most competitive method, 3D-Mem, by +4.6% (73.7 vs. 69.1). Despite being more accurate, SCOPE also improves efficiency. It achieves a higher SPL by +4.6% (53.5 vs. 48.9).

We also evaluate SCOPE on A-EQA, an embodied question answering benchmark, and the result is shown in Table 2. On A-EQA, SCOPE boosts answer accuracy by +6.5% over 3D-Mem (59.1 vs. 52.6), demonstrating better




Figure 2: Performance across modalities on GOAT-Bench.

Method	A-EQA			
	FR	FI	Corr.	Eff.
LLaVA-1.5 Frame Captions	–	–	38.1	7.0
CG Scene-Graph Captions	–	–	34.4	6.5
SVM Scene-Graph Captions	–	–	34.2	6.4
Multi-Frame	–	–	41.8	7.5
Explore-EQA (ICRA’2024b)	✓	–	46.9	23.4
CG + Frontier (ICRA’2024)	✓	–	47.2	33.3
TANGO (CVPR’2025)	–	–	37.2	/
MTU3D (ICCV’2025)	✓	–	51.1	42.6
3D-Mem (CVPR’2025)	✓	–	52.6	42.0
SCOPE (Ours)	✓	✓	59.1	41.0

Table 2: Performance comparison on **A-EQA**. **Bold** indicates the best performance. “/” denotes unreported results.

semantic comprehension and task alignment. Compared to MTU3D (the most efficient prior method on A-EQA with efficiency 42.6), SCOPE slightly trades off efficiency (41.0 vs. 42.6) for a notable gain in answer accuracy (+8.0%, 59.1 vs. 51.1), suggesting better generalizability and reliability. These results underscore the utility of our utilization of frontier information in enhancing long-term decision quality.

5 Analysis

To further evaluate the effectiveness of SCOPE and the contributions of our principle of fully utilizing frontier information, we conduct a series of analyses beyond standard evaluation. We evaluate the significance and calibration of SCOPE’s performance, and further isolate the contributions of each core component through targeted studies on potential estimation, potential graph propagation, and self-refinement.

5.1 Statistical Significance Analysis

To rigorously evaluate the performance improvements of SCOPE over the strongest baseline, we conduct a statistical significance analysis against the state-of-the-art method 3D-Mem. Both methods are independently run five times under identical random seed settings, and we report their mean success rates along with standard deviations to capture variability across runs.

The results are shown in Fig. 3, we apply an unpaired two-tailed t-test to assess the statistical significance of these improvements. The resulting p-value on GOAT-Bench is 0.046, indicating that the improvement is statistically significant at the 5% level. On A-EQA, the p-value is 0.1365, suggesting a positive trend in favor of SCOPE. These results indicate that SCOPE not only achieves higher average performance but also exhibits lower variance across runs. This highlights the robustness and consistency of its improvements over 3D-Mem. Together with the calibration analysis, the findings support SCOPE as a more stable and effective agent for embodied navigation tasks.

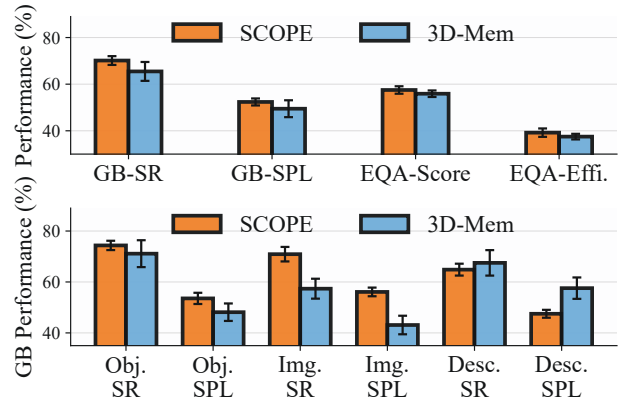


Figure 3: Performance comparison between SCOPE and 3D-Mem. **Top**: Results on the GOAT-Bench and A-EQA benchmarks, covering goal-based navigation (GB) and embodied question answering (EQA) tasks. **Bottom**: Detailed breakdown of GOAT-Bench SR and SPL across object-, image-, and description-goal settings. SCOPE achieves higher average performance and lower variance than 3D-Mem.

5.2 Model Calibration

Reliable navigation in embodied agents requires not only accurate decisions but also well-calibrated confidence scores that reflect the agent’s uncertainty—crucial for downstream decision-making and real-world deployment. Calibration measures how closely a model’s predicted confidence aligns with the true likelihood of correctness (Guo et al. 2017). A well-calibrated agent enables more trustworthy predictions and better risk assessment in practical deployments, which is especially important for an embodied agent.

To evaluate the reliability of SCOPE, we perform a calibration analysis on GOAT-Bench and A-EQA. We compute Expected Calibration Error (ECE) to quantify the discrepancy between predicted confidence and actual accuracy. As shown in Fig. 4, SCOPE achieves superior calibration performance. On GOAT-Bench, it reduces ECE from 11.62 (of the most competitive baseline 3D-Mem) to 3.83, and on A-EQA, from 11.55 to 8.12. These results indicate that SCOPE not only improves task performance but also produces more trustworthy estimates, equipping the agent with more reliable self-assessment for acting under uncertainty.

5.3 Utility Analysis

To further examine the utility of our utilization of the frontier, we design analysis experiments for each of our proposed components and evaluate its effectiveness.

Utility of Potential Estimator As for the proposed potential estimator, we design an analysis experiment that alters the form of the frontier input provided to the agent while keeping the estimated potential scores fixed. Specifically, we discard the original frontier images and instead provide the agent with a brief textual description of each frontier region, while still supplying the estimated potential score. This setup isolates the role of the potential scores as an inde-

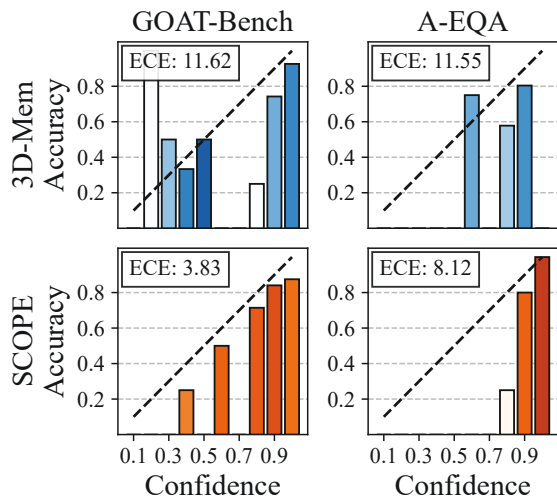


Figure 4: Calibration of 3D-Mem and SCOPE. “ECE” represents the estimated calibration error ($\times 100$), with lower values indicating better calibration. The dashed line denotes perfect calibration, and the bar colors become darker as they approach ideal calibration.

pendent guiding signal: if the agent can still navigate competently without visual frontier input, it indicates that the estimator encodes high-level, semantically meaningful, and goal-aware information beyond what is directly observable from raw pixels.

The results are shown in Fig. 5. Remarkably, even without frontier images, the agent maintains strong performance—achieving results on par with the full SCOPE model and significantly outperforming vanilla utilization of frontier images. This demonstrates that our potential estimation captures robust and task-relevant frontier representations that generalize beyond the visual modality. The potential scores alone equip the agent with a sufficiently rich understanding of the environment, enabling it to make informed decisions regardless of how the frontier is represented. These findings underscore the versatility and strength of our estimator in supporting high-level reasoning and navigation.

Utility of Potential Spread To assess the impact of the potential graph, we perform an ablation study where the agent only has access to the latest estimated potential scores of candidate frontiers, without the structured spatial-semantic propagation maintained by the potential graph. Unlike the full SCOPE model that accumulates and diffuses utility information over time and space, this variant makes decisions based solely on instantaneous frontier potentials, ignoring historical and contextual information.

As shown in Fig. 5, removing the potential graph leads to a clear performance drop: SR decreases from about 70.1% to 66.9%, and SPL falls from 52.4% to 47.0%. This result highlights the importance of structured spatial memory for long-term planning, allowing the agent to recall and leverage prior knowledge to revisit promising areas and avoid inefficient behaviors. The potential graph thus plays a crucial role




Figure 5: Ablation study evaluating the contribution of SCOPE components. **SCOPE w/o F. Img.** removes the frontier image input to the agent while retaining it for the potential estimator. **SCOPE w/o PG.** disables the potential graph module, exposing the agent only to raw estimated potential scores without spatial propagation.

in improving navigation efficiency and robustness beyond what raw frontier estimation alone can achieve.

Utility of Re-consideration To evaluate the utility of the proposed self-refinement technique, we conduct both quantitative and qualitative analyses. We begin by aggregating the outcomes of refinement decisions across all episodes. When the refinement module confirms an answer, it is correct in 80.6% of cases. Conversely, when the module considers an answer needing reconsideration, yet the answer is ultimately executed due to reaching the retry limit, it turns out to be incorrect in 34.3% of cases. These results demonstrate the effectiveness of self-refinement in both reinforcing correct judgments and mitigating overconfident errors.

6 Conclusion

We presented SCOPE, a new framework for embodied visual navigation tasks that explicitly leverages visual boundaries as a central cue for informed exploration and long-term planning. By estimating frontier potential and organizing it within a structured potential graph, SCOPE enables agents to reason more effectively about where to explore next. The addition of a self-reconsideration mechanism further enhances decision reliability by revisiting and refining prior actions. Empirical results confirm that SCOPE outperforms state-of-the-art baselines and improves calibration, generalization, and decision quality. These findings highlight the importance of structured, frontier-aware reasoning in advancing goal-driven embodied navigation.

Acknowledgments

This research was supported by the National Natural Science Foundation of China (No. 92464302, No.62503381, No. U24A20291).

References

- An, D.; Wang, H.; Wang, W.; Wang, Z.; Huang, Y.; He, K.; and Wang, L. 2024. Etpnav: Evolving topological planning for vision-language navigation in continuous environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Bar, A.; Zhou, G.; Tran, D.; Darrell, T.; and LeCun, Y. 2025. Navigation world models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15791–15801.
- Barsellotti, L.; Bigazzi, R.; Cornia, M.; Baraldi, L.; and Cucchiara, R. 2024. Personalized instance-based navigation toward user-specific objects in realistic environments. *Advances in Neural Information Processing Systems (NeurIPS)*, 37: 11228–11250.
- Cui, X.; Liu, Q.; Liu, Z.; and Wang, H. 2024. Frontier-enhanced topological memory with improved exploration awareness for embodied visual navigation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 296–313. Springer.
- Dang, R.; Yuan, Y.; Zhang, W.; Xin, Y.; Zhang, B.; Li, L.; Wang, L.; Zeng, Q.; Li, X.; and Bing, L. 2025. Ecbench: Can multi-modal foundation models understand the egocentric world? a holistic embodied cognition benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 24593–24602.
- Ehsani, K.; Gupta, T.; Hendrix, R.; Salvador, J.; Weihs, L.; Zeng, K.-H.; Singh, K. P.; Kim, Y.; Han, W.; Herrasti, A.; et al. 2024. Spoc: Imitating shortest paths in simulation enables effective navigation and manipulation in the real world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16238–16250.
- Gadre, S. Y.; Wortsman, M.; Ilharco, G.; Schmidt, L.; and Song, S. 2023. Cows on pasture: Baselines and benchmarks for language-driven zero-shot object navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 23171–23181.
- Goetting, D.; Singh, H. G.; and Loquercio, A. 2024. End-to-End Navigation with VLMs: Transforming Spatial Reasoning into Question-Answering. In *Workshop on Language and Robot Learning: Language as an Interface*.
- Gu, Q.; Kuwajerwala, A.; Morin, S.; Jatavallabhula, K. M.; Sen, B.; Agarwal, A.; Rivera, C.; Paul, W.; Ellis, K.; Chellappa, R.; et al. 2024. Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 5021–5028. IEEE.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1321–1330. PMLR.
- Janny, S.; Poirier, H.; Antsfeld, L.; Bono, G.; Monaci, G.; Chidlovskii, B.; Giuliarri, F.; Del Bue, A.; and Wolf, C. 2025. Reasoning in visual navigation of end-to-end trained agents: a dynamical systems approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12111–12121.
- Ji, Z.; Lin, H.; and Gao, Y. 2025. DyNaVLM: Zero-Shot Vision-Language Navigation System with Dynamic Viewpoints and Self-Refining Graph Memory. *arXiv preprint arXiv:2506.15096*.
- Khanna, M.; Mao, Y.; Jiang, H.; Haresh, S.; Shacklett, B.; Batra, D.; Clegg, A.; Undersander, E.; Chang, A. X.; and Savva, M. 2024a. Habitat synthetic scenes dataset (hssd-200): An analysis of 3d scene scale and realism tradeoffs for objectgoal navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16384–16393.
- Khanna, M.; Ramrakhya, R.; Chhablani, G.; Yenamandra, S.; Gervet, T.; Chang, M.; Kira, Z.; Chplot, D. S.; Batra, D.; and Mottaghi, R. 2024b. Goat-bench: A benchmark for multi-modal lifelong navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16373–16383.
- Kong, X.; Chen, J.; Wang, W.; Su, H.; Hu, X.; Yang, Y.; and Liu, S. 2024. Controllable navigation instruction generation with chain of thought prompting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 37–54. Springer.
- Lei, X.; Wang, M.; Zhou, W.; and Li, H. 2025. Gaussnav: Gaussian splatting for visual navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Lei, X.; Wang, M.; Zhou, W.; Li, L.; and Li, H. 2024. Instance-aware exploration-verification-exploitation for instance imagegoal navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16329–16339.
- Li, H.; Wang, Z.; Yang, X.; Yang, Y.; Mei, S.; and Zhang, Z. 2024. Memonav: Working memory model for visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17913–17922.
- Li, P.; Wu, K.; Fu, J.; and Zhou, S. 2025. REGNav: Room Expert Guided Image-Goal Navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, 4860–4868.
- Lin, B.; Nie, Y.; Wei, Z.; Chen, J.; Ma, S.; Han, J.; Xu, H.; Chang, X.; and Liang, X. 2025a. Navcot: Boosting llm-based vision-and-language navigation via learning disentangled reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Lin, X.; Lin, T.; Huang, L.; Xie, H.; and Su, Z. 2025b. Bip3d: Bridging 2d images and 3d perception for embodied intelligence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9007–9016.
- Liu, R.; Wang, W.; and Yang, Y. 2024. Volumetric environment representation for vision-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16317–16328.
- Majumdar, A.; Ajay, A.; Zhang, X.; Putta, P.; Yenamandra, S.; Henaff, M.; Silwal, S.; Mcvay, P.; Maksymets, O.;

- Arnaud, S.; et al. 2024. Openeqa: Embodied question answering in the era of foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16488–16498.
- Pan, Y.; Xu, Y.; Liu, Z.; and Wang, H. 2025. Planning from imagination: Episodic simulation and episodic memory for vision-and-language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, 6345–6353.
- Perincherry, A.; Krantz, J.; and Lee, S. 2025. Do Visual Imaginations Improve Vision-and-Language Navigation Agents? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3846–3855.
- Ren, A. Z.; Clark, J.; Dixit, A.; Itkina, M.; Majumdar, A.; and Sadigh, D. 2024a. Explore until Confident: Efficient Exploration for Embodied Question Answering. In *Proceedings of Robotics: Science and Systems XX (RSS 2024)*, Delft, The Netherlands, July 15-19, 2024.
- Ren, A. Z.; Clark, J.; Dixit, A.; Itkina, M.; Majumdar, A.; and Sadigh, D. 2024b. Explore Until Confident: Efficient Exploration for Embodied Question Answering. In *Proceedings of the First Workshop on Vision-Language Models for Navigation and Manipulation (VLMNM) at ICRA*. ArXiv:2403.15941.
- Ren, H.; Zeng, Y.; Bi, Z.; Wan, Z.; Huang, J.; and Cheng, H. 2025. Prior does matter: Visual navigation via denoising diffusion bridge models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12100–12110.
- Song, X.; Chen, W.; Liu, Y.; Chen, W.; Li, G.; and Lin, L. 2025. Towards long-horizon vision-language navigation: Platform, benchmark and method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12078–12088.
- Wang, F.; Lyu, S.; Zhou, P.; Duan, A.; Guo, G.; and Navarro-Alarcon, D. 2025a. Instruction-Augmented Long-Horizon Planning: Embedding Grounding Mechanisms in Embodied Mobile Manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, 14690–14698.
- Wang, L.; He, Z.; Dang, R.; Shen, M.; Liu, C.; and Chen, Q. 2024a. Vision-and-language navigation via causal learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13139–13150.
- Wang, W.; Duan, C.; Peng, Z.; Liu, Y.; and Zhou, B. 2025b. Embodied scene understanding for vision language models via metavqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 22453–22464.
- Wang, Z.; Li, X.; Yang, J.; Liu, Y.; Hu, J.; Jiang, M.; and Jiang, S. 2024b. Lookahead exploration with neural radiance representation for continuous vision-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13753–13762.
- Xie, W.; Jiang, H.; Zhu, Y.; Qian, J.; and Xie, J. 2025. Navigator: A Spatio-Temporal Context-Aware Transformer for Object Navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, 14708–14716.
- Yang, Y.; Yang, H.; Zhou, J.; Chen, P.; Zhang, H.; Du, Y.; and Gan, C. 2025. 3D-mem: 3D scene memory for embodied exploration and reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17294–17303.
- Yin, H.; Xu, X.; Wu, Z.; Zhou, J.; and Lu, J. 2024. Sg-nav: Online 3d scene graph prompting for llm-based zero-shot object navigation. *Advances in Neural Information Processing Systems (NeurIPS)*, 37: 5285–5307.
- Yin, H.; Xu, X.; Zhao, L.; Wang, Z.; Zhou, J.; and Lu, J. 2025. Unigoal: Towards universal zero-shot goal-oriented navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19057–19066.
- Yokoyama, N.; Ha, S.; Batra, D.; Wang, J.; and Bucher, B. 2024. Vlfm: Vision-language frontier maps for zero-shot semantic navigation. In *Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA)*, 42–48. IEEE.
- Yu, X.; Zhang, S.; Song, X.; Qin, X.; and Jiang, S. 2024. Trajectory diffusion for objectgoal navigation. *Advances in Neural Information Processing Systems (NeurIPS)*, 37: 110388–110411.
- Zhang, G.; Tang, H.; and Yan, Y. 2024. Versatile navigation under partial observability via value-guided diffusion policy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17943–17951.
- Zhang, S.; Song, X.; Yu, X.; Bai, Y.; Guo, X.; Li, W.; and Jiang, S. 2025. HOZ++: Versatile Hierarchical Object-to-Zone Graph for Object Navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Zhang, S.; Yu, X.; Song, X.; Wang, X.; and Jiang, S. 2024. Imagine before go: Self-supervised generative map for object goal navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16414–16425.
- Zhao, G.; Li, G.; Chen, W.; and Yu, Y. 2024. OVER-NAV: Elevating Iterative Vision-and-Language Navigation with Open-Vocabulary Detection and StructurEd Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16296–16306.
- Zhu, Z.; Wang, X.; Li, Y.; Zhang, Z.; Ma, X.; Chen, Y.; Jia, B.; Liang, W.; Yu, Q.; Deng, Z.; et al. 2025. Move to Understand a 3D Scene: Bridging Visual Grounding and Exploration for Efficient and Versatile Embodied Navigation. *arXiv preprint arXiv:2507.04047*.
- Ziliotto, F.; Campari, T.; Serafini, L.; and Ballan, L. 2025. TANGO: training-free embodied AI agents for open-world tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 24603–24613.