

PanoNav: Mapless Zero-Shot Object Navigation with Panoramic Scene Parsing and Dynamic Memory

Qunchao Jin, Yilin Wu, Changhao Chen*

PEAK-Lab, The Hong Kong University of Science and Technology (Guangzhou)
qunchao.jin@outlook.com, ywu044@connect.hkust-gz.edu.cn, changhaochen@hkust-gz.edu.cn

Abstract

Zero-shot object navigation (ZSON) in unseen environments remains a challenging problem for household robots, requiring strong perceptual understanding and decision-making capabilities. While recent methods leverage metric maps and Large Language Models (LLMs), they often depend on depth sensors or prebuilt maps, limiting the spatial reasoning ability of Multimodal Large Language Models (MLLMs). Mapless ZSON approaches have emerged to address this, but they typically make short-sighted decisions, leading to local deadlocks due to a lack of historical context. We propose PanoNav, a fully RGB-only, mapless ZSON framework that integrates a Panoramic Scene Parsing module to unlock the spatial parsing potential of MLLMs from panoramic RGB inputs, and a Memory-guided Decision-Making mechanism enhanced by a Dynamic Bounded Memory Queue to incorporate exploration history and avoid local deadlocks. Experiments on the public navigation benchmark show that PanoNav significantly outperforms representative baselines in both SR and SPL metrics.

Introduction

Object Goal Navigation (ObjectNav) (Sun et al. 2024) is a fundamental capability for domestic robots to effectively fulfill their functions. This task requires a robot to identify a specific object described by a human, plan a path toward it, and actively navigate through the environment to reach the target. As a prerequisite to downstream tasks—such as robotic search and exploration, mobile manipulation (Billard and Kragic 2019) and human-robot interaction (Ajoudani et al. 2018)—reliable ObjectNav is critical for enabling autonomous robotic assistance.

Since the introduction of standardized definitions and evaluation metrics in the Habitat Challenge (Batra et al. 2020), ObjectNav has attracted increasing research interest. Various approaches have been proposed to tackle this problem, including end-to-end learning (Ramrakhya et al. 2023; Zhang et al. 2023), modular frameworks (Chen et al. 2023; Chang, Gupta, and Gupta 2020), and zero-shot generalization methods (Dorbala, Mullen, and Manocha 2023;

Cai et al. 2024), many of which have shown promising results. In particular, the open-vocabulary setting presents a greater challenge, as it requires the agent to recognize and navigate to objects beyond a predefined set of categories. Recent advances in vision-language models (VLMs) (Radford et al. 2021; Li et al. 2023) and large language models (LLMs) (Achiam et al. 2024; Liu et al. 2025; Bai et al. 2025) have opened new possibilities for solving this problem, leading to a surge of LLM-based approaches (Majumdar et al. 2022; Yokoyama et al. 2024; Zhou et al. 2023; Wu et al. 2024; Yu, Kasaei, and Cao 2023; Zhao et al. 2025; Nie et al. 2025).

Despite this progress, several limitations remain in existing methods. First, many approaches rely on imitation learning (Ramrakhya et al. 2022) or reinforcement learning (Wortsman et al. 2019), which are often confined to narrow domains or limited object categories—making them unsuitable for real-world open-vocabulary scenarios. Second, a significant number of prior works depend on precise depth images to construct 2.5D scene representations for spatial reasoning. However, in both physical robotic systems and virtual agents, RGB-only input is far more common. Dependence on depth information thus constrains the applicability and generalizability of these methods. Third, many recent methods (Yokoyama et al. 2024; Long et al. 2024; Nie et al. 2025) require prebuilt metric maps for navigation, which involve costly depth sensors and accurate robot localization. Such dependencies increase the system’s hardware burden and reduce robustness in dynamic or noisy environments. To address these challenges, a few recent works (Majumdar et al. 2022; Cai et al. 2024; Zhao et al. 2025) have proposed mapless open-vocabulary navigation frameworks. However, these methods often make decisions based only on the agent’s current observation, ignoring historical trajectory information. As a result, they are prone to local deadlocks—e.g., repeatedly revisiting already explored areas due to lack of memory.

To this end, we propose **PanoNav**, a **Panoramic Scene Parsing-based Navigation** framework for RGB-only, mapless, and open-vocabulary Object Goal Navigation. PanoNav features two core components: (1) Panoramic Scene Parsing, which leverages six multi-view RGB observations and an open-source multimodal large language model (MLLM) to build a panoramic spatial understanding without requiring

*Corresponding author.

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

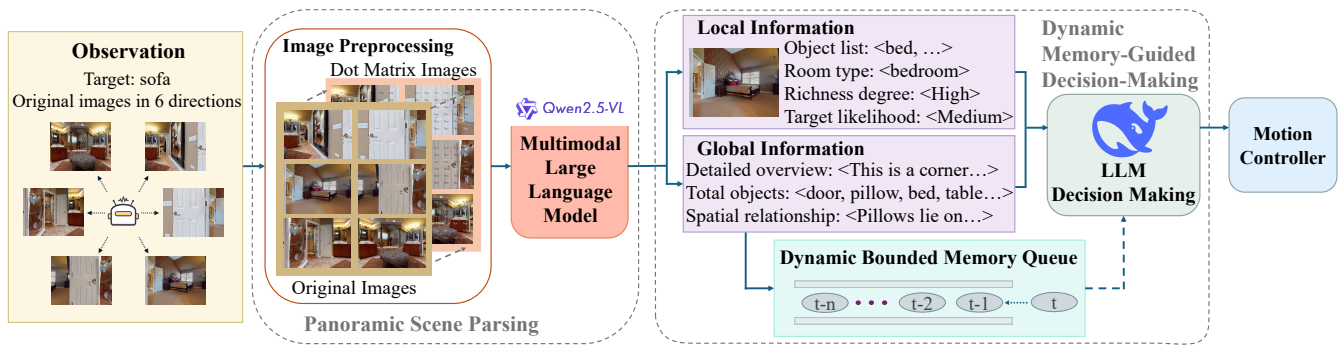


Figure 1: Overview of the proposed PanoNav framework. At each timestep, the robot captures six directional RGB images to form a panoramic view. Each image is preprocessed into a dot matrix map, and both the RGB and dot matrix images are fed into an MLLM for spatial parsing. The model outputs local directional descriptions and a global scene summary. Global summaries are stored in a Dynamic Bounded Memory Queue, and together with local and global information, are used by the LLM to make navigation decisions that guide the robot’s movement.

depth sensors or metric maps; and (2) Dynamic Memory-Guided Decision-Making, which integrates historical trajectory information via a Dynamic Bounded Memory Queue to improve exploration efficiency and avoid local deadlocks. To the best of our knowledge, PanoNav is the first RGB-only, mapless open-vocabulary ObjectNav framework that addresses the local deadlock issue. Experiments on the HM3D dataset demonstrate that PanoNav outperforms state-of-the-art baselines such as PixNav (Cai et al. 2024) and ZSON (Majumdar et al. 2022) under the same settings in terms of both Success Rate (SR) and Success weighted by Path Length (SPL), and even surpasses several map-based and closed-vocabulary methods.

In summary, our main contributions are as follows:

- We propose PanoNav, an RGB-only framework that leverages multi-view visual input and vision-language reasoning to perform panoramic scene parsing for mapless, open-vocabulary object navigation.
- We introduce a Dynamic Memory-guided Decision-Making mechanism that incorporates a Dynamic Bounded Memory Queue to track visited areas and avoid local deadlocks.
- Extensive experiments demonstrate the superiority of our approach, showing significant improvements over state-of-the-art baselines in the open-vocabulary ObjectNav setting.

Related Work

Open-Vocabulary Object Navigation

While object navigation has seen significant progress and achieve accepted performance, many existing methods (Yadav et al. 2023; Ramrakhya et al. 2022, 2023; Ye et al. 2021) remain limited to predefined object categories, failing to generalize to the open-vocabulary real-world scenarios. Therefore, open-vocabulary object navigation has gained growing attention in the research community. Gadre et al. (2023) suggested a baseline algorithm named CoW for language-driven zero-shot object navigation by using

open-vocabulary model. SSNet (Zhao et al. 2023) addresses the generalization gap in ObjectNav for new categories via a semantic similarity network. VLFM (Yokoyama et al. 2024) leverages vision-language models to generate semantic value maps for frontier-based exploration (Topiwala, Inani, and Kathpal 2018) without task-specific training.

Large-Language-Models for Embodied Navigation

In recent years, research on foundation models has flourished, with numerous studies applying LLMs to embodied navigation tasks, including vision-language navigation (VLN), demand-driven navigation (DDN), and object navigation. NavGPT (Zhou, Hong, and Wu 2024) is a navigation system that leverages LLMs for instruction decomposition and commonsense reasoning to achieve decision-making, ultimately enabling effective VLN. Qiao et al. (2025) present Open-Nav for zero-shot VLN in continuous environments using open-source LLMs, reducing costs and enhancing privacy protection. SG-Nav (Yin et al. 2024) advances ObjectNav by representing environments as dynamically updated 3D scene graphs, where LLMs leverage hierarchical node relationships to prioritize exploration frontiers. VoroNav (Yin et al. 2024) is the another method for ObjectNav, which leverages LLMs to reason about topological waypoints extracted from graphs. Long et al. (2024) attempted to integrate navigation tasks with diverse linguistic instructions, proposing the InstructNav, which makes great use of potential of LLMs to integrate different types of instruction into its Dynamic Chain-of-Navigation module. These LLM-based approaches often rely on the metric map for navigation and decision-making.

Mapless Object Navigation

Recently, some progress has been made in mapless ObjectNav within the field of embodied navigation. ZSON (Majumdar et al. 2022) achieves mapless zero-shot object navigation by leveraging the multimodal representation and alignment capabilities of CLIP (Radford et al. 2021), enabling generalization in open-world scenarios. Imagine-

Nav (Zhao et al. 2025) predicts the next viewpoint using a novel synthesis model, then employs a vision-language model to select the most promising direction. This approach can achieve mapless object navigation. PixelNav (Cai et al. 2024) represents a zero-shot navigation approach that operates solely on RGB images, leveraging panoramic image analysis and LLMs to identify exploration-worthy pixels. However, these mapless approaches typically make decisions based solely on the current situation without considering historical information, which can easily lead to local deadlocks.

Mapless Zero-Shot Object Navigation

Problem Formulation

In the Object Goal Navigation task, an agent must explore an unfamiliar indoor environment to locate an instance I of a specified object category C (e.g., bed, toilet, chair) based solely on a linguistic instruction. Each episode begins with the agent randomly placed in an unseen environment without any prior map or spatial knowledge.

At each timestep t , the agent receives an egocentric panoramic observation \mathbf{V}_t , consisting of six RGB images $\{\mathbf{V}_t^i\}_{i=1}^6$, captured at 60-degree intervals. No additional sensor data (e.g., depth, GPS) is provided. The agent must then choose an action to navigate toward the target. The task is considered successful if the agent stops within a pre-defined distance threshold of the object. The action space includes six discrete commands: *Stop*, *MoveAhead*, *TurnLeft*, *TurnRight*, *LookUp*, *LookDown*. Specifically, *MoveAhead* advances the agent by 0.25 meters, while *TurnLeft* and *TurnRight* rotate the agent by 30 degrees.

PanoNav Framework

PanoNav is a mapless, RGB-only navigation framework for open-vocabulary ObjectNav. Its architecture is illustrated in Fig. 1. At each timestep, the agent captures a panoramic image set composed of six RGB views taken at 60-degree intervals. These images are first converted into corresponding dot matrix representations to enhance spatial structure cues. The panoramic RGB and dot matrix images are then fed into the Panoramic Scene Parsing module, which outputs: 1) local textual descriptions for each directional view, and 2) a global scene summary describing the agent’s current surroundings. These textual cues, along with a Dynamic Bounded Memory Queue that stores recent global summaries, are input into the Memory-Guided Decision-Making module. This module selects the most promising action based on both current and historical context. The chosen action is then executed by the motion controller.

Each component is detailed in the following sections.

Panoramic Scene Parsing

Recent years have witnessed increasing interest in using multimodal large language models (MLLMs) for scene understanding (Bai et al. 2025; Liu et al. 2023). While previous zero-shot object navigation methods have shown potential, they often suffer from limited perceptual scope or rely heavily on depth images to construct spatial 2.5D scene

representations. To address these challenges, we propose a Panoramic Scene Parsing module that leverages multiview RGB images and the reasoning capabilities of MLLMs to generate both fine-grained local directional parsing and a holistic global panoramic summary at each timestep.

Local Directional Parsing The goal of local directional parsing is to extract rich contextual information from each directional view, including object presence, spatial relationships among objects, room type, the likelihood of target object presence, and overall information richness. In particular, parsing spatial relationships plays a crucial role in enhancing the agent’s spatial understanding.

To fully exploit the MLLM’s capability in spatial relationship reasoning, we introduce dot matrix images as additional inputs. At each timestep t , the agent captures an egocentric panoramic observation composed of six RGB views, denoted as $\mathbf{V}_t = \mathbf{V}_t^1, \dots, \mathbf{V}_t^6$. For each view \mathbf{V}_t^i , we generate a corresponding dot matrix image \mathbf{M}_t^i using the Scaffold (SCA) (Lei et al. 2025) processing method:

$$\mathbf{M}_t^i = \text{SCA}(\mathbf{V}_t^i), \quad i \in \{1, \dots, 6\}. \quad (1)$$

As shown in Fig. 2, the raw RGB image \mathbf{V}_t^i captures geometric distance cues, while the paired dot matrix image \mathbf{M}_t^i enhances planar position understanding. Together, these complementary inputs allow the MLLM to reason over both spatial dimensions. We construct a spatial relation graph G_t^i in the latent space:

$$G_t^i = \mathcal{P}(\Psi(\mathbf{V}_t^i), \Phi(\mathbf{V}_t^i, \mathbf{M}_t^i)), \quad (2)$$

where $\Psi(\cdot)$ extracts geometric distance relations, $\Phi(\cdot)$ parses planar positional relationships, and $\mathcal{P}(\cdot)$ aggregates these two relational representations into a unified graph. This process yields a detailed and structured description of spatial relations for each direction, significantly enhancing the agent’s local spatial awareness.

Global Panoramic Summary Beyond local parsing, we perform a global analysis of the robot’s spatial context to derive higher-level semantic understanding. Unlike the directional views that focus on object-level detail, the global summary emphasizes holistic environmental perception. Specifically, this involves identifying what objects are generally present in the surrounding environment and determining the type of room or scene the agent currently occupies.

This global parsing provides a form of implicit self-location awareness, offering cues about the agent’s position within the broader environment (e.g., kitchen, hallway). The output of this panoramic summarization serves as a foundational input to our decision-making module, particularly in the construction of the dynamic memory state used to guide future actions.

Navigation Decision-Making with Dynamic Memory

In the absence of metric maps, previous mapless zero-shot object navigation methods often fail to incorporate historical exploration information, making them susceptible to local deadlocks—situations where the agent becomes trapped

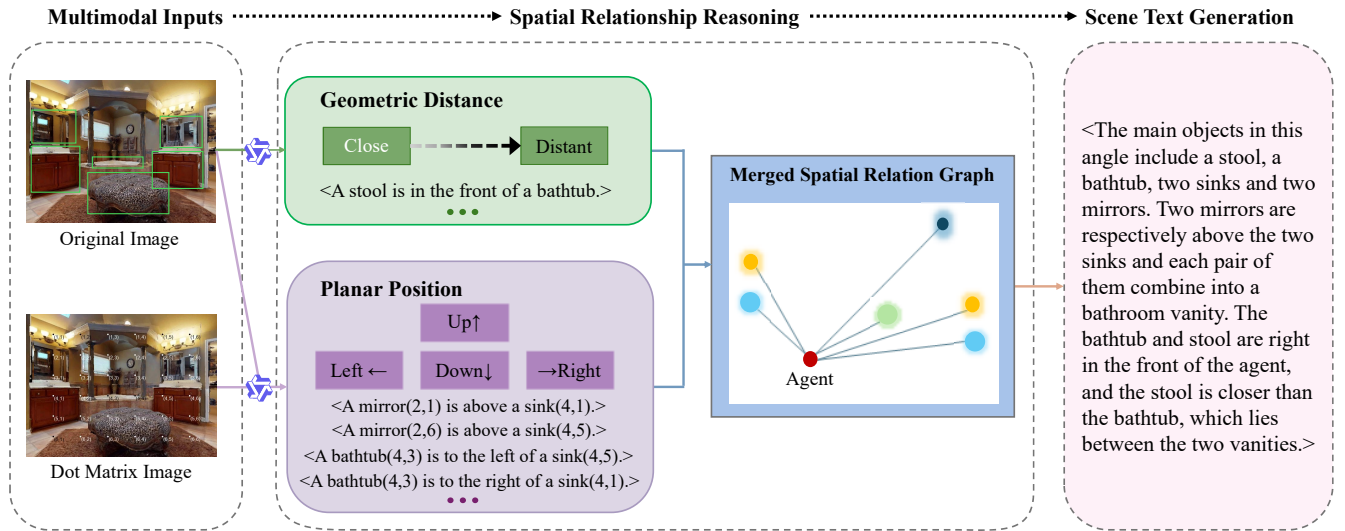


Figure 2: Spatial relationship parsing in Panoramic Scene Parsing. The MLLM processes both the original RGB image and its corresponding dot matrix image to extract geometric distance and planar positional relationships between objects, producing textual descriptions of the spatial scene.

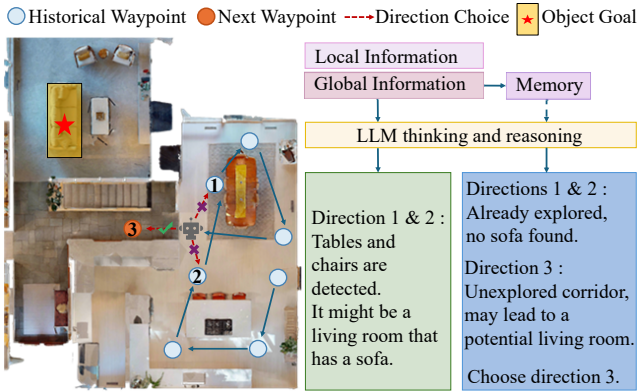


Figure 3: Workflow of the Dynamic Memory-Guided Decision mechanism. Without memory, the LLM may cause the robot to revisit previously explored areas. Incorporating a memory mechanism guides the LLM to improve exploration and avoid redundant navigation.

in a limited area, endlessly circling within a room. We observed that agents tend to remain in regions that semantically resemble likely target locations (e.g., a living room when searching for a sofa), even if the target is not actually present. This is primarily because LLMs tend to over-rely on object-room priors during decision-making, without considering the agent’s exploration history.

As shown in Fig. 3, although the agent has already searched the right-hand region, the LLM still favors this area as the likely target location due to the contextual presence of tables and chairs—objects commonly associated with living rooms. This bias causes the agent to overlook other unexplored yet potentially relevant regions, such as corridors.

To alleviate this challenge, we introduce a Dynamic Memory-guided Decision-Making mechanism featuring the Dynamic Bounded Memory Queue, which effectively incorporates the robot’s historical exploration data into the decision-making process. As depicted in Fig. 1, the panoramic observation is processed through our Local-Global Panoramic Parsing module, yielding a sequence of local descriptions information ld_t^i and a global summary information gs_t for the current timestep t .

We utilize the global summary description of each timestep to construct historical information, which is maintained as the Dynamic Bounded Memory Queue \mathcal{Q} with a maximum length of n . At the initial stage of each episode, the queue is empty ($\mathcal{Q}_0 = \{\}$). As the agent moves, new elements are continuously appended to the queue until it reaches full capacity. At this point, the queue’s full flag f_t is set to 1. At timestep t , the queue is represented as:

$$\mathcal{Q}_t = \{\text{gs}_{t-n}, \dots, \text{gs}_{t-2}, \text{gs}_{t-1}\}. \quad (3)$$

Thereafter, when a new waypoint is reached, the oldest waypoint description is dequeued while the newest one is added, maintaining a constant queue length.

The decision system receives local descriptions information and global summary information of the current waypoint as inputs. When the Dynamic Bounded Memory Queue is not full, only these two inputs are considered by the LLM for determining the promising direction. Once the queue reaches full capacity, the system additionally incorporates the memory queue as input for the decision-making process:

$$\begin{cases} \mathbf{r}_t = \mathcal{F}(\text{ld}_t, \text{gs}_t), & f_t = 0 \\ \mathbf{r}_t = \mathcal{F}(\text{ld}_t, \text{gs}_t, \mathcal{Q}_t), & f_t = 1 \end{cases}, \quad (4)$$

where $\mathcal{F}(\cdot)$ donates the decision LLM, and \mathbf{r}_t represents the decision result at timestep t , which includes both the deci-

Method	Navigation Setting			SR \uparrow	SPL \uparrow
	Modality	Scope	Map		
FBE	RGB-D, GPS+Compass	Close-Set	map-based	33.7	15.3
SemExp	RGB-D, GPS+Compass	Close-Set	map-based	37.9	18.8
Habitat-Web	RGB-D, GPS+Compass	Close-Set	mapless	41.5	16.0
OVRL	RGB-D, GPS+Compass	Close-Set	mapless	62.0	26.8
VLFM	RGB-D, GPS+Compass	Open-Set	map-based	52.2	30.4
ESC	RGB-D, GPS+Compass	Open-Set	map-based	39.2	22.3
VoroNav	RGB-D, GPS+Compass	Open-Set	map-based	42.0	26.0
L3MVN	RGB-D, GPS+Compass	Open-Set	map-based	50.4	23.1
ImagineNav	RGB-D	Open-Set	mapless	53.0	23.8
ZSON	RGB Only	Open-Set	mapless	25.5	12.6
PixNav	RGB Only	Open-Set	mapless	37.9	20.5
PanoNav (Ours)	RGB Only	Open-Set	mapless	43.5	23.7

Table 1: Comparison of Success Rate (SR) and Success-weighted Path Length (SPL) across state-of-the-art methods under various navigation settings. ZSON, PixNav, and our method are the only mapless, open-vocabulary approaches using RGB-only inputs. Our method outperforms both and even surpasses several methods with stronger assumptions.

sion direction and a flag indicating whether the target has been found.

Once a decision is made, we employ PixNav (Cai et al. 2024) as the motion controller to carry out the navigation action. In summary, the Dynamic Bounded Memory Queue empowers the decision system to reason over historical exploration contexts, reducing the risk of local deadlocks and enabling more efficient object search in mapless open-vocabulary settings.

Experiment

Benchmark and Evaluation Metrics

We evaluate the effectiveness of our proposed method using the Habitat simulator on the ObjectNav benchmark, specifically the Habitat-Matterport 3D Research Dataset (HM3D) (Ramakrishnan et al. 2021), a widely used benchmark for object goal navigation. HM3D provides high-fidelity reconstructions of 20 complete buildings. For our experiments, we randomly selected 200 episodes from its validation set. In our framework, the Local-Global Panoramic Parsing module utilizes Qwen-2.5-VL (Bai et al. 2025) as the Multimodal Large Language Model (MLLM), while the Memory-Guided Decision-Making System is powered by DeepSeek-V3 (Liu et al. 2025) as the Large Language Model (LLM).

We evaluate navigation performance using two standard metrics:

- Success Rate (SR): the percentage of episodes in which the agent successfully navigates to the target object within a defined threshold.
- Success-weighted Path Length (SPL): a measure of navigation efficiency, computed as:

$$SPL = \frac{1}{N} \sum_{i=1}^N s_i \left(\frac{\ell_i}{\max(\rho_i, \ell_i)} \right) \quad (5)$$

where ρ_i and ℓ_i denote the agent path length and ground true path length in episode i , respectively. s_i is a binary success indicator, which means the $SPL_i = 0$ when the certain episode i fails.

	SR \uparrow	SPL \uparrow	DTS (f) \downarrow	ER \uparrow
without Memory	12.0	4.9	6.7	32.0
with Memory	48.0	19.2	4.7	82.0

Table 2: Results of deadlock avoidance test. The SR, SPL, DTS (f), and ER correspond to Success Rate, Success-weighted Path Length, Distance to Success (in fail cases), and Escape Rate, respectively.

Performance on Mapless Zero-Shot ObjectNav

We selected multiple object navigation baselines with diverse navigation setting for comparative analysis. As shown in Table 1, these baselines were systematically categorized along three dimensions in the navigation framework: (1) Modality: the information required for the navigation method; (2) Scope: whether this method supports open-vocabulary navigation; (3) Map: whether the approach relies on building a metric map.

For close-set object navigation approaches, we selected four baseline methods, including two map-based approaches (FBE (Topiwala, Inani, and Kathpal 2018) and SemExp (Chaplot et al. 2020)) that utilize RGB-D information to construct metric maps for guiding agent exploration. Additionally, we incorporated two mapless methods: Habitat-Web (Ramrakhya et al. 2022), which employs large-scale imitation learning by replicating human task execution strategies, and OVRL (Yadav et al. 2023), which obtains a powerful visual encoder through extensive pre-training.

We also introduced several open-vocabulary ObjectNav methods as baselines. Among them, VLFM (Yokoyama

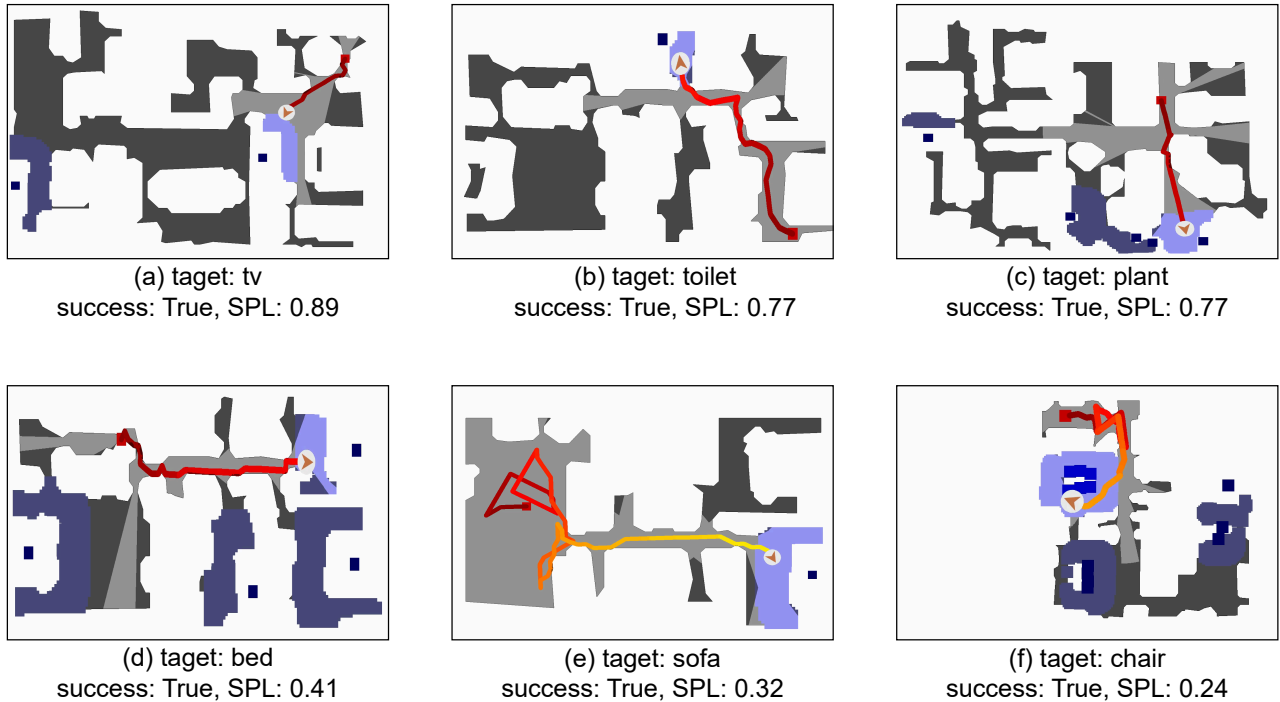


Figure 4: The top-down view of the navigation trajectories includes six representative examples with different object types. Cases (a)-(d) demonstrate several smooth navigation scenarios, while cases (e) and (f) present more complex situations where the robot initially circled within a local space before successfully escaping and ultimately reaching the target.




Figure 5: Visualization Results of the Deadlock Avoidance Test. (a) Memory-less decision (b) Memory-guided decision.

et al. 2024), ESC (Zhou et al. 2023), VoroNav (Wu et al. 2024), and L3MVN (Yu, Kasaei, and Cao 2023) are map-based approaches that rely on metric maps, employing different models to extract richer semantic features to guide the LLM decision. ImagineNav (Zhao et al. 2025) represents a mapless method that imagines future observation images and selects the most promising future direction for navigation. However, it incorporates spatial knowledge by using

depth information during training of its imagination model. ZSON (Majumdar et al. 2022) and PixNav (Cai et al. 2024) are the two methods most similar to our approach in navigation settings, as they both require only RGB data and can achieve mapless open-vocabulary object navigation.

Table 1 presents the comparative results of our proposed PanoNav against previous methods on the HM3D dataset, and our method achieves performance scores of 43.5% in SR and 23.7% in SPL. Compared with the baseline PixNav, which has the same navigation settings as ours, PanoNav demonstrates a 14.76% improvement in SR and a 15.61% improvement in SPL, respectively. Notably, the SR of our method even surpasses several close-set methods (FBE, SemExp, Habitat-Wen) and map-based methods (FBE, SemExp, ESC, VoroNav).

Fig. 4 presents the navigation trajectory results for locating different types of objects. As shown in Fig. 4 (a-d), the agent demonstrates the capability to rapidly and accurately localize targets and navigate to them in some scenarios. In more challenging environments with potential deceptive (Fig. 4 (e-f)), the agent exhibits the ability to recognize if it’s trapped in a deadlock situation and actively attempts to escape the current area to continue searching for the target.

Deadlock Avoidance Test

To further validate the effectiveness of our proposed Dynamic Bounded Memory Queue, we designed an interesting Deadlock Avoidance Test. Specifically, we selected several

Panorama Views	Decoupled Parsing and Decision	Memory Guidance	Success Rate \uparrow	SPL \uparrow
\times	\checkmark	\checkmark	19.5	9.97
\checkmark	\times	\times	35.0	20.47
\checkmark	\checkmark	\times	38.5	22.57
\checkmark	\checkmark	\checkmark	43.5	23.73

Table 3: Results of Ablation Study.

highly deceptive episodes where the agent’s starting positions were either: (1) locate where the target object would most likely appear (but is actually absent), such as “starting in a living room looking for a sofa”, or (2) spatial boundary areas, such as “starting in a hallway looking for a toilet”.

We conducted 10 repeated experiments on our selected 5 episodes, evaluating performance differences between using the memory queue and not using it. We evaluated by using four metrics. SR and SPL were the same metrics in the comparative experiments. Distance to Success (DTS) refers to the distance between the robot’s stopping position and the object location. We selected failed navigation cases to calculate the mean DTS (f). Escape Rate (ER) means the percentage of the robot escaping from a local space; we defined an escape as unsuccessful if the robot stops near its starting point.

As shown in Table 2, both the SR and SPL metrics demonstrate that agents equipped with memory guidance exhibit significantly higher chances to find the object, achieving success rate up to four times greater than the memory-less variant. Through comparative analysis of the DTS (f), we observed that even in unsuccessful episode, memory-guided navigation enables agents to approach object more closely. Furthermore, memory-guided method achieves an impressive escape rate of 82.0%, substantially outperforming the memory-less version (32.0%). The superior performance in both DTS (f) and ER metrics shows that the dynamic memory possesses stronger exploratory capabilities while avoiding deadlock to limited regions.

Fig. 5 presents two visualized navigation cases. The top row (a) shows trajectories without memory guidance, while the bottom row (b) demonstrates trajectories with memory guidance. Comparative analysis reveals that agents lacking dynamic memory exhibit repetitive looping behavior, as the LLMs may consider previously visited locations as the optimal choices.

Ablation Study

To evaluate the contribution of each component in our proposed framework, we conduct a series of ablation experiments.

Panorama View vs Limited View Our method leverages six-view panoramic observations for comprehensive scene understanding. However, collecting six views introduces additional observation steps. To assess the trade-off, we compare this setup with a limited three-view (forward-facing) configuration. As shown in the first row of Table 3, reducing to three views significantly degrades performance, with

the success rate dropping to 19.5% and SPL to 9.97%. The narrower field of view limits spatial awareness, causing the agent to miss critical navigational cues or make suboptimal decisions in complex environments.

Decoupled Parsing and Decision vs One-Step Decision

PanoNav adopts a decoupled parsing and decision strategy: an MLLM first parses RGB inputs into structured textual descriptions, which are then used by an LLM for decision-making. We compare this with a one-step approach where the MLLM directly outputs navigation decisions without intermediate descriptions. As shown in the second row of Table 3, the one-step method achieves moderate success but underperforms compared to the decoupled pipeline. This suggests that explicitly separating perception and reasoning helps reduce cognitive overload and encourages step-by-step processing. Current MLLMs may struggle with end-to-end reasoning when faced with rich, unstructured visual input.

Memory Guidance vs No Memory Guidance

We further evaluate the impact of the Dynamic Bounded Memory Queue by removing it from the framework. As shown in the third row of Table 3, the absence of memory guidance significantly weakens navigation performance. Without memory, the agent is more prone to local deadlocks and inefficient exploration. In contrast, our memory-guided approach enables better long-horizon planning and reduces revisiting previously explored areas, which is also supported by results in the Deadlock Avoidance Test.

Conclusion

In this work, we propose a mapless open-vocabulary object navigation approach PanoNav, that requires only RGB image inputs for navigation. Facing the limitations of most current ObjectNav methods that require mapping or depth images to parse spatial structures, we propose Panoramic Scene Parsing, which unlocks the spatial relationship parsing capability of MLLMs and provides more comprehensive reference information for decision-making systems. Furthermore, to mitigate the local deadlock issue commonly encountered in mapless methods, we propose a Dynamic Memory-guided Decision-Making mechanism that enables escape from local regions demonstrates superior exploration performance. Extensive experimental results have showed the effectiveness of both our overall framework and its individual components. In future work, we will explore leveraging multimodal information to construct memory queues for more robust mapless object navigation.

Acknowledgments

This work is funded by the Young Elite Scientist Sponsorship Program by CAST (No. YESS20220181) and the National Natural Science Foundation of China (NFSC) under the Grant Number 62573370.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2024. GPT-4 Technical Report. arXiv:2303.08774.
- Ajoudani, A.; Zanchettin, A. M.; Ivaldi, S.; Albu-Schäffer, A.; Kosuge, K.; and Khatib, O. 2018. Progress and prospects of the human–robot collaboration. *Autonomous Robots*, 42(5): 957–975.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2.5-VL Technical Report. arXiv:2502.13923.
- Batra, D.; Gokaslan, A.; Kembhavi, A.; Maksymets, O.; Mottaghi, R.; Savva, M.; Toshev, A.; and Wijmans, E. 2020. ObjectNav Revisited: On Evaluation of Embodied Agents Navigating to Objects. arXiv:2006.13171.
- Billard, A.; and Kragic, D. 2019. Trends and challenges in robot manipulation. *Science*, 364(6446): eaat8414.
- Cai, W.; Huang, S.; Cheng, G.; Long, Y.; Gao, P.; Sun, C.; and Dong, H. 2024. Bridging Zero-shot Object Navigation and Foundation Models through Pixel-Guided Navigation Skill. In *International Conference on Robotics and Automation (ICRA)*, 5228–5234.
- Chang, M.; Gupta, A.; and Gupta, S. 2020. Semantic Visual Navigation by Watching YouTube Videos. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 4283–4294.
- Chaplot, D. S.; Gandhi, D. P.; Gupta, A.; and Salakhutdinov, R. R. 2020. Object Goal Navigation using Goal-Oriented Semantic Exploration. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 4247–4258.
- Chen, J.; Li, G.; Kumar, S.; Ghanem, B.; and Yu, F. 2023. How To Not Train Your Dragon: Training-free Embodied Object Goal Navigation with Semantic Frontiers. In *Robotics: Science and Systems (RSS)*.
- Dorbala, V. S.; Mullen, J. F.; and Manocha, D. 2023. Can an Embodied Agent Find Your “Cat-shaped Mug”? LLM-Based Zero-Shot Object Navigation. *IEEE Robotics and Automation Letters (RA-L)*, 9(5): 4083–4090.
- 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.
- Lei, X.; Yang, Z.; Chen, X.; Li, P.; and Liu, Y. 2025. Scaffolding Coordinates to Promote Vision-Language Coordination in Large Multi-Modal Models. In *Proceedings of the 31st International Conference on Computational Linguistics (COLING)*, 2886–2903.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. In *International Conference on Machine Learning (ICML)*, 19730–19742.
- Liu, A.; Feng, B.; Xue, B.; Wang, B.; Wu, B.; Lu, C.; Zhao, C.; Deng, C.; Zhang, C.; Ruan, C.; et al. 2025. DeepSeek-V3 Technical Report. arXiv:2412.19437.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in Neural Information Processing Systems (NeurIPS)*, 36: 34892–34916.
- Long, Y.; Cai, W.; Wang, H.; Zhan, G.; and Dong, H. 2024. InstructNav: Zero-shot System for Generic Instruction Navigation in Unexplored Environment. In *8th Annual Conference on Robot Learning (CoRL)*.
- Majumdar, A.; Aggarwal, G.; Devnani, B.; Hoffman, J.; and Batra, D. 2022. ZSON: Zero-Shot Object-Goal Navigation using Multimodal Goal Embeddings. *Advances in Neural Information Processing Systems (NeurIPS)*, 35: 32340–32352.
- Nie, D.; Guo, X.; Duan, Y.; Zhang, R.; and Chen, L. 2025. WMNav: Integrating Vision-Language Models into World Models for Object Goal Navigation. arXiv:2503.02247.
- Qiao, Y.; Lyu, W.; Wang, H.; Wang, Z.; Li, Z.; Zhang, Y.; Tan, M.; and Wu, Q. 2025. Open-Nav: Exploring Zero-Shot Vision-and-Language Navigation in Continuous Environment with Open-Source LLMs. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning (ICML)*, 8748–8763.
- Ramakrishnan, S. K.; Gokaslan, A.; Wijmans, E.; Maksymets, O.; Clegg, A.; Turner, J.; Undersander, E.; Galuba, W.; Westbury, A.; Chang, A. X.; Savva, M.; Zhao, Y.; and Batra, D. 2021. Habitat-Matterport 3D Dataset (HM3D): 1000 Large-scale 3D Environments for Embodied AI. arXiv:2109.08238.
- Ramrakhya, R.; Batra, D.; Wijmans, E.; and Das, A. 2023. PIRLNav: Pretraining with Imitation and RL Finetuning for OBJECTNAV. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17896–17906.
- Ramrakhya, R.; Undersander, E.; Batra, D.; and Das, A. 2022. Habitat-Web: Learning Embodied Object-Search Strategies from Human Demonstrations at Scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5173–5183.
- Sun, J.; Wu, J.; Ji, Z.; and Lai, Y.-K. 2024. A Survey of Object Goal Navigation. *IEEE Transactions on Automation Science and Engineering (TASE)*, 22: 2292–2308.
- Topiwala, A.; Inani, P.; and Kathpal, A. 2018. Frontier Based Exploration for Autonomous Robot. arXiv:1806.03581.

Wortsman, M.; Ehsani, K.; Rastegari, M.; Farhadi, A.; and Mottaghi, R. 2019. Learning to Learn How to Learn: Self-Adaptive Visual Navigation using Meta-Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6750–6759.

Wu, P.; Mu, Y.; Wu, B.; Hou, Y.; Ma, J.; Zhang, S.; and Liu, C. 2024. VoroNav: Voronoi-based Zero-Shot Object Navigation with Large Language Model. In *International Conference on Machine Learning (ICML)*, 53757–53775.

Yadav, K.; Ramrakhya, R.; Majumdar, A.; Berges, V.-P.; Kuhar, S.; Batra, D.; Baevski, A.; and Maksymets, O. 2023. Offline Visual Representation Learning for Embodied Navigation. In *Workshop on Reincarnating Reinforcement Learning at ICLR 2023*.

Ye, J.; Batra, D.; Das, A.; and Wijmans, E. 2021. Auxiliary Tasks and Exploration Enable Objectgoal Navigation. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, 16117–16126.

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.

Yokoyama, N.; Ha, S.; Batra, D.; Wang, J.; and Bucher, B. 2024. VLFM: Vision-Language Frontier Maps for Zero-shot Semantic Navigation. In *International Conference on Robotics and Automation (ICRA)*, 42–48.

Yu, B.; Kasaei, H.; and Cao, M. 2023. L3MVN: Leveraging Large Language Models for Visual Target Navigation. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3554–3560.

Zhang, S.; Song, X.; Li, W.; Bai, Y.; Yu, X.; and Jiang, S. 2023. Layout-based Causal Inference for Object Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10792–10802.

Zhao, Q.; Zhang, L.; He, B.; Qiao, H.; and Liu, Z. 2023. Zero-Shot Object Goal Visual Navigation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2025–2031.

Zhao, X.; Cai, W.; Tang, L.; and Wang, T. 2025. ImagineNav: Prompting Vision-Language Models as Embodied Navigator through Scene Imagination. In *The Thirteenth International Conference on Learning Representations (ICLR)*.

Zhou, G.; Hong, Y.; and Wu, Q. 2024. NavGPT: Explicit Reasoning in Vision-and-Language Navigation with Large Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, 7641–7649.

Zhou, K.-Q.; Zheng, K.; Pryor, C.; Shen, Y.; Jin, H.; Getoor, L.; and Wang, X. E. 2023. ESC: Exploration with Soft Commonsense Constraints for Zero-shot Object Navigation. In *International Conference on Machine Learning (ICML)*, 42829–42842.