

SSMG-Nav: Enhancing Lifelong Object Navigation with Semantic Skeleton Memory Graph

Haochen Niu, Lantao Zhang, Xingwu Ji, Rendong Ying, Peilin Liu[†] and Fei Wen

Abstract—Navigating to out-of-sight targets from human instructions in unfamiliar environments is a core capability for service robots. Despite substantial progress, most approaches underutilize reusable, persistent memory, constraining performance in lifelong settings. Many are additionally limited to single-modality inputs and employ myopic greedy policies, which often induce inefficient back-and-forth maneuvers (BFMs). To address such limitations, we introduce SSMG-Nav, a framework for object navigation built on a *Semantic Skeleton Memory Graph* (SSMG) that consolidates past observations into a spatially aligned, persistent memory anchored by topological keypoints (e.g., junctions, room centers). SSMG clusters nearby entities into subgraphs, unifying entity- and space-level semantics to yield a compact set of candidate destinations. To support multimodal targets (images, objects, and text), we integrate a vision-language model (VLM). For each subgraph, a multimodal prompt synthesized from memory guides the VLM to infer a target belief over destinations. A long-horizon planner then trades off this belief against traversability costs to produce a visit sequence that minimizes expected path length, thereby reducing backtracking. Extensive experiments on challenging lifelong benchmarks and standard ObjectNav benchmarks demonstrate that, compared to strong baselines, our method achieves higher success rates and greater path efficiency, validating the effectiveness of SSMG-Nav.

I. INTRODUCTION

Service robots are expected to navigate to target objects beyond their current field of view in previously unseen environments based on human instructions. Conventional methods typically rely on reinforcement learning (RL) or imitation learning (IL) [1], [2]. Despite notable progress, these methods exhibit two fundamental limitations: (i) they are data- and interaction-intensive, making training costly; and (ii) they often degrade under distribution shifts when deployment scenarios diverge from the training domain.

Human navigation suggests prioritizing high-likelihood locations (e.g., bedrooms or shelves for a stuffed toy), highlighting the value of commonsense spatial semantic reasoning. Meanwhile, large-scale pretrained models have shown strong generalization in commonsense and reasoning, offering a new paradigm for navigation. In recent years, numerous zero-shot object navigation methods that draw on vision-language models (VLMs) or large language models (LLMs) have emerged, demonstrating promising advantages by leveraging broad prior knowledge to link environmental semantic cues with target objects [3]–[6].

This work was supported by the National Natural Science Foundation of China (No. 62276166) and the STI 2030-Major Projects (No. 2022ZD0208700).

All authors are with the Brain-Inspired Application Technology Center (BATC), Shanghai Jiao Tong University, Shanghai 200240, China (email: {haochen_niu, swagger, jixingwu, rdying, liupeilin, wenfei}@sjtu.edu.cn).

Nevertheless, several gaps remain for real-world, user-centered navigation, especially in lifelong settings where agents are expected to leverage historical experience for continual improvement [7], [8]. First, the design of memory mechanisms remains underexplored: many methods make purely reactive decisions from instantaneous observations; others construct value maps [4], [9] or semantic maps [10], [11] that lack persistence and reusability. Recent efforts on graph memory are promising but most of them model “objects” [6], [12] or “waypoints” [13] as nodes, connect them via adjacency relations, and construct the memory graph in a text-based form, which limits the overall expressiveness. Second, most methods target single-modality goals and cannot accommodate multimodal instructions. Third, planners are commonly greedy, driving the agent to the most likely target locations with short-sighted decisions. This often leads to inefficient back-and-forth maneuvers (BFMs) [14].

These limitations motivate a navigation framework that maintains persistent memory and supports multimodal goals for efficient lifelong task execution. At its core is a Semantic Skeleton Memory Graph (SSMG) that provides reusable, spatially grounded semantics to support long-horizon planning. Unlike existing graph representations that center on objects [6] or waypoints [13] as nodes, SSMG aggregates objects around topological keypoints of the environment skeleton (e.g., corridor junctions, room centers). This yields a set of candidate destinations, each represented as a subgraph that naturally combines entity semantics and spatial semantics. Then we incorporate a VLM and, for each SSMG subgraph, derive a multimodal prompt to guide the VLM in evaluating belief that the target lies in that subgraph. Leveraging the VLM’s multimodal commonsense and reasoning, we can accommodate targets specified as images, objects, or text while fully exploiting the local scene context encoded by each subgraph. Moreover, we design a long-horizon planner that jointly considers belief and traversability cost to compute a visit sequence minimizing the expected travel distance. Using an efficient approximate solver to address this optimization, the planner effectively mitigates myopic behavior and improves navigation efficiency.

The main contributions are as follows.

- A novel Semantic Skeleton Memory Graph (SSMG) that consolidates historical observations into persistent, spatially grounded memory, unifying entity and spatial semantics to support informed decision-making.
- An SSMG-driven long-horizon planner that elicits VLM-based environmental reasoning via multimodal prompts and employs an efficient approximation to

balance belief and traversal cost for long-term visitation planning, improving efficiency.

- A complete lifelong multimodal object navigation framework, SSMG-Nav, which achieves state-of-the-art performance on public benchmarks.

II. RELATED WORK

A. Object Navigation

Existing approaches to object navigation can be broadly categorized into learning-based and zero-shot methods.

1) *Learning-based Object Navigation*: Learning-based methods have progressed rapidly, with reinforcement learning (RL) and imitation learning (IL) taking center stage. Building on large-scale human demonstrations, Habitat-Web [2] shows that IL can internalize the implicit, often intricate behaviors underlying object search. On the modular front, PONI [9] uses a potential-function network to decide “where to look?” and pairs it with an analytical local policy for action generation, while SGM [15] complements this line by predicting maps beyond the agent’s visible field via self-supervision. ZSON [16] brings zero-shot generality into the picture by training a policy with DD-PPO and leveraging CLIP [17] for image-goal navigation. More recently, MVVT [1] taps multi-view inputs to couple global image cues with local spatial context, further boosting policy learning. Despite these advances, learning-based methods typically require large-scale data and tend to degrade under distribution shift.

2) *Zero-shot Object Navigation*: In recent years, powered by advances in large models, zero-shot object navigation has seen a surge of new methods. Some studies leverage vision-language models (VLMs), such as CLIP [17], to align observations with target objects, incorporating advanced exploration or learning-based strategies for open-set object navigation [3]. However, these methods lack analysis of environmental semantics and memory, resulting in lower exploration efficiency.

PIVOT [18] and VLNav [19] cast navigation as video question answering, querying a VLM to select motion directions. Building on this idea, DyNavLM [20] augments the system with a graph memory architecture, representing objects as nodes and relational adjacencies as edges.

Methods such as PONI [9], VLFM [4] and WMNav [21] construct a value map to score frontier points on the map, thereby guiding the agent to prioritize exploration of regions with higher target relevance. However, such value maps are often not reusable because they are tightly coupled to the current task.

In addition, several approaches construct a semantic map. Methods like ESC [10], L3MVN [22], and OpenFMNav [11] aggregate object categories, locations, and relations into a semantic map and render candidate waypoints in the egocentric view as prompts; an LLM then chooses the next waypoint. Extending this paradigm, InstructNav [5] leverages dynamic chain-of-navigation (DCoN) and spatiotemporal chains [23] to explicitly structure spatial reasoning in natural language.

Some recent work focuses on graph representations of memory/map that can efficiently integrate different kinds of information. SG-NAV [12] constructed a hierarchical scene graph [24] and proposed a CoT prompting method for LLMs to reason about target locations. Similarly, SayNav [25] created an object-room-scene topological map using objects as nodes in a bottom-up way. MCGPT [13] samples viewpoints during movement as topological nodes and utilizes an offline constructed CoT database to retrieve the most suitable prompt templates for online analysis.

B. Lifelong Multimodal Object Navigation

Recently, some researchers have begun to focus on object navigation tasks that better align with real-world scenarios. [6] pointed out the limitations of existing zero-shot methods in constructing reasoning frameworks on LLMs for specific tasks, noting that they cannot generalize across different types of targets. Therefore, they proposed a graph structure with objects as nodes for unified multimodal target representation, verifying targets through multi-stage subgraph matching based on LLM inferences, though lifelong task execution was not considered. SayNav [25] studied multi-object navigation tasks but can only present all targets at once and operates in a single modality. OpenIN [26] focuses on lifelong task execution but is also limited to a single modality. In contrast, GOAT [7] formalizes Lifelong Multimodal Object Navigation with sequential goals specified by object categories, images, or language descriptions, and GOAT-Bench [8] provides a recent standardized benchmark with a unified protocol and public baselines for fair comparison.

III. METHOD

A. Task Definition and System Overview

The lifelong multimodal object navigation task places an agent in a previously unseen scene and requires it to navigate sequentially to multiple goal objects, each specified by an object category (e.g., “towel”), a text description (e.g., “a bath towel hanging on a hook in the bathroom, relative to the shower stall and shower bench”), or a unique image. Each goal defines a subtask, which ends when the time budget is exhausted or the agent issues a “stop” action, after which it switches to the next subtask until completion. At each switch, the subtask is successful if the agent is within a distance threshold of the goal object; otherwise, it fails.

As shown in Fig. 1, our proposed SSMG-Nav is structured as follows. At each step, the agent receives RGB, depth, and pose from the simulator. Using these sensor inputs, we incrementally update geometric information (occupancy map and skeleton topology) and semantic information (value map and object map), thereby constructing the SSMG over time (Sec. III-B). The navigation policy then transitions between states (Sec. III-D). Conditioned on the current state, we generate multimodal prompts from SSMG subgraphs to guide a VLM in belief estimation (Sec. III-C.1). Finally, a long-horizon planner jointly optimizes traversal cost and belief to produce a long-term visitation route and executes its first waypoint (Sec. III-C.2).

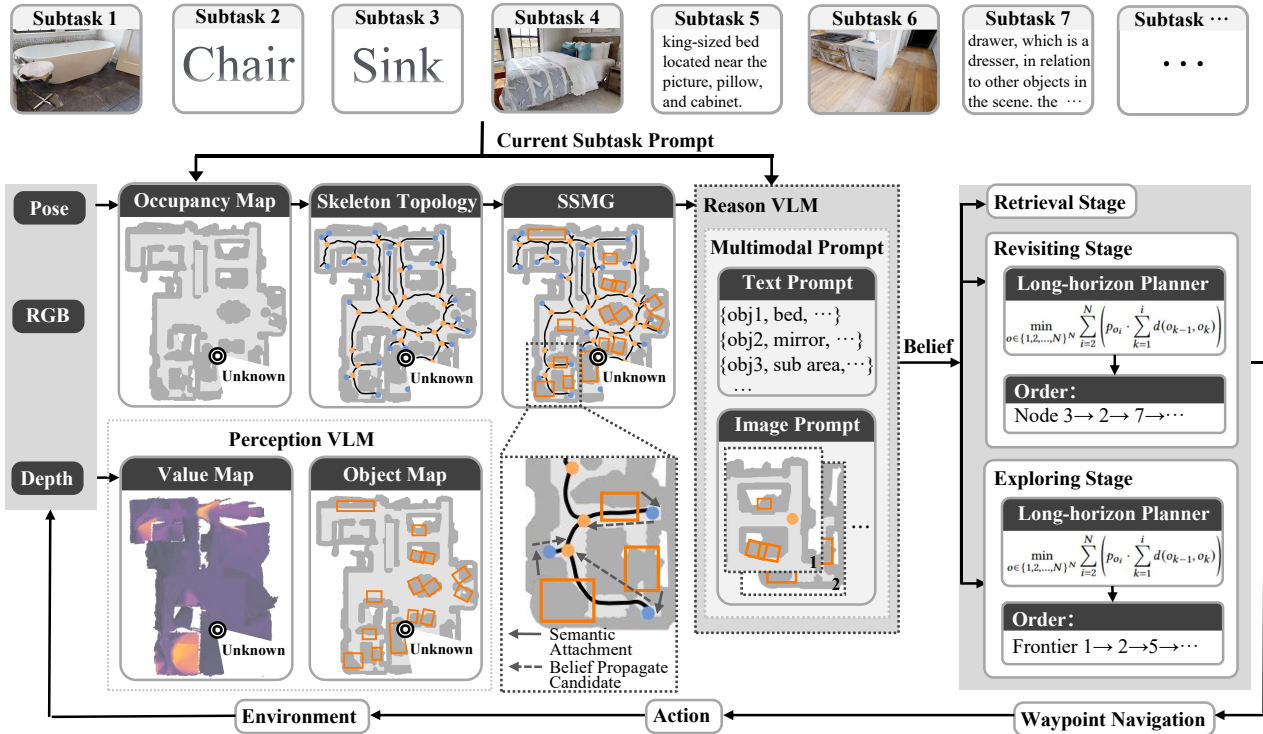


Fig. 1: An overview of the proposed SSMG-Nav.

B. Semantic Skeleton Memory Graph Construction

1) *Visual Perception*: At each time step t , the agent acquires sensor observations $obs_t = \{I_t, D_t, P_t\}$, comprising the RGB I_t , depth D_t , and pose P_t .

We first convert the depth D_t into a 3D point cloud and project it onto a bird’s-eye-view (BEV) plane. Using ray casting from the sensor origin, we construct and incrementally update a 2D occupancy grid M_{occ} . Each grid cell is labeled as *free*, *occupied*, or *unknown* to track exploration progress.

To analyze the semantic information in the observations, we employ BLIP-2 [27] to evaluate the relevance between the current observation and the target. Following [4], we adopt the text template “*Seems like there is a <target object> ahead.*” to perform text–image matching on the current RGB input I_t . Analogous to M_{occ} , we project the relevance scores onto the BEV plane to obtain a value map M_{value} , which encodes the correlation between the scene context and the target object across the explored area.

While the value map captures target-specific relevance, lifelong navigation requires a target-agnostic semantic representation of observations. We therefore use BLIP-2 again to generate a caption for the current image I_t , denoted as cap_t . To obtain structured semantics, we apply the open-set detector Grounding-DINO [28] to I_t and produce a set of detections

$$Det_t = \{(\mathbf{b}_i, s_i, c_i) \mid i = 1, \dots, N\}, \quad (1)$$

where \mathbf{b}_i is the bounding box, s_i the confidence score, c_i the predicted category, and N the number of detected objects. For each detection in Det_t , we apply MobileSAM [29]

within the corresponding bounding box to obtain an instance mask. By fusing the mask with the depth, we back-project pixels to 3D and construct an object-level point cloud map M_{obj} . For each object $j \in \mathcal{O}$, we store:

$$\mathcal{O}_j := \{\mathbf{PC}_j, c_j, \mathcal{V}_j\}, \quad \mathcal{V}_j := \{(\mathbf{I}_j^k, \mathbf{P}_j^k, cap_j^k, s_j^k)\}_{k=1}^K. \quad (2)$$

Here \mathbf{PC}_j is the point cloud, c_j the category, and \mathcal{V}_j collects up to K views: RGB \mathbf{I}_j^k , pose \mathbf{P}_j^k , caption cap_j^k , and score s_j^k . We set $K = 3$ in our experiments.

Over time, M_{occ} and M_{value} are incrementally updated via coverage. Updates to M_{obj} are gated by semantic category consistency and the alignment between 3D bounding boxes of the point cloud. For objects subject to fusion, historical entries \mathcal{V}_j are replaced based on detection confidence, thereby retaining higher-confidence observations that typically correspond to better viewpoints.

2) *Skeleton Topology Extraction*: Object-cluster graphs [12] may lose geometric fidelity, while trajectory-based node generation [30] can yield non-representative nodes. We therefore aim to extract a concise geometric backbone and enrich it with semantics later.

Following [14], we approximate the skeleton of M_{occ} via morphological thinning to obtain a one-pixel-wide medial axis. Each skeleton pixel is classified as endpoint \mathcal{N}_e (degree 1), connector \mathcal{N}_c (degree 2), or junction \mathcal{N}_j (degree ≥ 3). We then prune connectors and retain only \mathcal{N}_e and \mathcal{N}_j (special nodes), adding edges between them. This yields the graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where connectors are stored as edge attributes to support fast path planning.

3) *Semantic Attachment*: We augment the skeletal topology $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with object information from M_{obj} to obtain a complete SSMG: for each object $j \in \mathcal{O}$, we select its highest-confidence viewpoint $k_j^* := \arg \max_k s_j^k$ and set the reference pose $\mathbf{P}_j^* := \mathbf{P}_j^{k_j^*}$; projecting pose to BEV via $\pi(\cdot)$, we attach the object to its nearest node:

$$n_j^* := \arg \min_{n \in \mathcal{N}} \|\mathbf{x}_n - \pi(\mathbf{P}_j^*)\|, \mathcal{O}(n_j^*) \leftarrow \mathcal{O}(n_j^*) \cup \{j\}. \quad (3)$$

To improve coverage while remaining concise, each node n defines a data-driven radius from the farthest currently attached object,

$$r_n := \begin{cases} \max_{j \in \mathcal{O}(n)} \|\mathbf{x}_n - \pi(\mathbf{P}_j^*)\|, & \mathcal{O}(n) \neq \emptyset, \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

and adds all objects within this radius,

$$\mathcal{O}(n) \leftarrow \mathcal{O}(n) \cup \{j \in \mathcal{O} \mid \|\mathbf{x}_n - \pi(\mathbf{P}_j^*)\| \leq r_n\}, \quad (5)$$

which naturally allows objects to belong to multiple nodes (i.e., $j \in \mathcal{O}(n_1) \cap \mathcal{O}(n_2)$). To support downstream belief propagation, we augment the edge set by linking each endpoint to directly connected junction node(s) on the skeleton (i.e., one-hop connectivity on the skeleton graph):

$$\mathcal{E} \leftarrow \mathcal{E} \cup \{(n_e, n_j) \in \mathcal{N}_e \times \mathcal{N}_j \mid \text{skeleton_adj}(n_e, n_j)\}, \quad (6)$$

where $\text{skeleton_adj}(n_e, n_j)$ is true iff n_e and n_j are one-hop neighbors on the skeleton topology.

C. SSMG-driven Long-Horizon Planner

Zero-shot object navigation increasingly leverages VLMs/LLMs to translate perception into planning priors. However, prevailing systems select the single most promising frontier at each step [6], [10], [12], without sequencing multiple candidates over long horizons, which frequently results in BFMs and inefficient trajectories—a phenomenon also reported in autonomous exploration [14]. Moreover, limitations in spatial understanding [31] and long-horizon planning [32], [33] cast doubt on the reliability of directly feeding large-scale global maps and rich semantic information into LLMs in some form for long-horizon planning.

We therefore propose an SSMG-driven Long-Horizon Planner (LHP) that builds SSMG-based multimodal prompts for reasoning over nodes of SSMG, followed by an approximation algorithm for long-term path planning.

1) *VLM-Based Belief Evaluation*: Belief estimation is performed per special node. We first build multimodal prompts from endpoints \mathcal{N}_e . Each node stores nearby object information; for each object $j \in \mathcal{O}(n_i)$, we form a textual prompt $\{\text{object_id}_j, \text{object_class}_j, \text{object_caption}_j\}$. We then extract a subgraph as an image prompt: centered at node n_i , we annotate the BEV plane with each object’s point-cloud centroid and its object_id_j . Concatenating the textual and visual prompts yields a background-knowledge prompt. Finally, CoT reasoning guides the VLM (Qwen-VL-Plus

[34]) to analyze node n_i step by step and produce a formatted belief estimate, belief_i .

After evaluating all endpoints \mathcal{N}_e , we propagate endpoint cues bottom-up to \mathcal{N}_j . This focuses candidates on junctions, reducing the planning search space. For each junction, we augment its background-knowledge prompt with information from its associated endpoints and query the VLM again to output the junction belief.

2) *Long-Horizon Planning*: Our LHP treats all \mathcal{N}_j as candidate nodes and computes a visitation path that sequentially accesses them. The agent proceeds along this path until the target object is found. Given N candidate nodes $\{n_1, \dots, n_N\}$ with target probabilities $\{p_1, \dots, p_N\}$ and pairwise distances $d(i, j)$, we seek an order $\{n_{o_1}, \dots, n_{o_N}\}$ that minimizes the expected travel while searching for the target:

$$\min_{o \in \{1, 2, \dots, N\}^N} \sum_{i=2}^N \left(p_{o_i} \cdot \sum_{k=2}^i d(o_{k-1}, o_k) \right), \quad (7)$$

subject to:

$$\begin{aligned} \sum_{i=1}^N p_i &= 1, \quad p_i \geq 0, \\ |\{o_1, o_2, \dots, o_N\}| &= N. \end{aligned} \quad (8)$$

This problem is NP-hard. Inspired by the 2-opt local search heuristic [35], we propose an approximate solution. We first convert the estimated belief_i from Sec. III-C.1 into probabilities p_i via softmax with temperature τ :

$$p_i = \frac{\exp(\text{belief}_i / \tau)}{\sum_{j=1}^N \exp(\text{belief}_j / \tau)}, \quad \tau > 0. \quad (9)$$

We first obtain an initial path via a greedy heuristic that always visits the nearest unvisited node. For simplicity, let n_1 denote n_{o_1} , and denote the initial path as $\pi_0 = \{n_1, n_2, \dots, n_N\}$. Its expected travel cost $E(\pi_0)$ is computed using (7). Next, select a segment $\{n_p, n_{p+1}, \dots, n_q\}$, reverse it, and replace the original segment to form $\pi_{pq} = \{n_1, \dots, n_{p-1}, n_q, n_{q-1}, \dots, n_{p+1}, n_p, \dots, n_N\}$. Compute $E(\pi_{pq})$; if it is lower than the current expectation, accept π_{pq} . Iterate over all (p, q) pairs and repeat until reaching the maximum iterations or no further improvement.

D. Navigation Policy

Based on the aforementioned content, we develop a strategy with four stages: retrieval, revisiting, exploring, and waypoint navigation.

1) *Retrieval Stage*: When switching to a new subtask, we first query the SSMG for the specified target. If same-category instances are found, we construct prompts as in Sec. III-C.1. If an instance is verified as the target, we proceed to waypoint navigation. If the target is absent from the current SSMG, we assess exploration status on M_{occ} . If valid frontiers remain, we enter the exploring stage to search unexplored regions. If exploration is complete, we consider potential missed detections (e.g., due to viewpoint) and transition to revisiting.

TABLE I: Lifelong Multimodal Object Navigation Comparison.

Method	Zero-shot	GOAT val.seen			GOAT val.unseen		
		s-SR (\uparrow)	e-SR(\uparrow)	SPL(\uparrow)	s-SR (\uparrow)	e-SR(\uparrow)	SPL(\uparrow)
SenseAct-NN Skill Chain [8]	×	0.292	–	0.128	0.295	–	0.113
SenseAct-NN Monolithic [8]	×	0.168	–	0.094	0.123	–	0.068
SenseAct-NN Monolithic [†]	×	0.159	0.003	0.104	0.126	0.008	0.088
CoW [3]	✓	0.148	–	0.087	0.161	–	0.104
Modular GOAT [8]	✓	0.263	–	<u>0.175</u>	0.249	–	<u>0.172</u>
PIVOT [18]	✓	–	–	–	0.102	–	0.055
VLMNav [19]	✓	–	–	–	0.201	–	0.096
DyNaVLM [20]	✓	–	–	–	0.255	–	0.102
VLFM [†] [4]	✓	<u>0.308</u>	<u>0.025</u>	0.129	<u>0.406</u>	<u>0.028</u>	0.159
SSMG-Nav(Ours)	✓	0.412	0.056	0.318	0.465	0.086	0.341

2) *Revisiting Stage*: In this stage, we follow Sec. III-C to compute the access sequence π and visit nodes sequentially. At each junction, the agent performs a 360° spin to acquire panoramic observations. If the target is matched, we transition to waypoint navigation; otherwise, we check for adjacent end nodes, visit them (if any), and repeat the 360° scan. If still unmatched, proceed to the next junction in π and repeat.

3) *Exploring Stage*: In this stage, we also conduct long-horizon planning to improve exploration efficiency. Specifically, we extract the boundaries between known and unknown regions as frontiers. These frontiers are then clustered to obtain several frontier points $\{f_1, f_2, \dots, f_n\}$. We then solve the path using the method described in Sec. III-C.2. Considering the incompleteness of SSMG information and the uncertainty of unknown areas during the exploration process, we simply use a value map M_{value} as the belief for each frontier point, similar to the approach in [4]. After computing the access path, we navigate to the first frontier; as the environment updates during motion, we replan at each step and navigate to the new target via waypoint navigation.

4) *Waypoint Navigation Stage*: In the navigation stage, we perform A* path planning based on the target points and generate lookahead points at specific distances. We then adopt a point navigation scheme, utilizing the model trained with the Variable Experience Rollout (VER) reinforcement learning algorithm on the HM3D dataset [36] as a local planner, as described in [4]. This planner uses depth and the relative position and orientation of the target points as inputs, producing the corresponding actions.

IV. EXPERIMENTAL RESULTS

A. Experimental Setup

1) *Benchmarks*: We evaluate on four high-resolution Habitat simulator-based benchmarks [37] spanning two tasks. First, we consider the challenging lifelong multimodal object navigation task introduced by GOAT-Bench [8]. We use the *GOAT val.seen* and *GOAT val.unseen* splits; each split contains 36 indoor scenes, each scene includes 10 episodes, and each episode consists of 5–10 subtasks. Each subtask

TABLE II: Object Navigation Comparison.

Method	Zero-shot	HM3D		MP3D	
		SR (\uparrow)	SPL(\uparrow)	SR (\uparrow)	SPL(\uparrow)
PONI [9]	×	–	–	0.318	0.121
Habitat-Web [2]	×	0.415	0.160	0.316	0.085
ZSON [16]	×	0.255	0.126	0.153	0.048
SGM [15]	×	0.602	0.308	0.377	0.147
ESC [10]	✓	0.392	0.223	0.287	0.142
L3MVN [22]	✓	0.504	0.231	0.349	0.145
OpenFMNav [11]	✓	0.549	0.244	0.372	0.157
SG-Nav [12]	✓	0.540	0.249	0.402	0.160
VLFM [4]	✓	0.525	0.304	0.364	<u>0.175</u>
VLFM [†]	✓	0.523	0.303	0.339	0.165
InstructNav [5]	✓	0.510	0.187	0.420	0.161
UniGoal [6]	✓	0.545	0.251	0.410	0.164
WMNav [21]	✓	0.581	<u>0.312</u>	0.454	0.172
Ours	✓	<u>0.595</u>	0.337	<u>0.443</u>	0.191

specifies a target object via one of three modalities: object class, a reference image, or a natural-language description.

To further assess our method’s generality, we also evaluate on the classic object navigation task using two standard datasets, *HM3D* [36] and *MP3D* [38]. The first split comprises 2,000 episodes across 20 scenes with 6 goal categories, and the second split contains 2,195 episodes across 11 scenes with 21 goal categories. This task can be regarded as a special case of the lifelong multimodal setting—restricted to the single modality of object class—with exactly one subtask per episode.

2) *Metrics*: In object navigation tasks, the *Success Rate (SR)* and *Success Rate Weighted by Inverse Path Length (SPL)* are commonly used as evaluation metrics.

Given the particularities of lifelong multimodal object navigation, we refine the overall SR into two granular metrics: *subtask Success Rate (s-SR)* and *episode Success Rate (e-SR)*. Concretely, s-SR is computed by first measuring the success of each subtask within an episode and then averaging these subtask-level success rates across all episodes. We deem an episode successful only if all of its subtasks are completed

successfully; correspondingly, e-SR reports the proportion of episodes that satisfy this criterion, i.e., the fraction of fully completed episodes.

3) *Baselines*: For lifelong multimodal object navigation, We compare SSMG-Nav against two RL methods, SenseAct-NN Skill Chain and Monolithic [8], and six zero-shot methods. Among them, CoW [3], PIVOT [18], and VLNav [19] are memory-free (stateless) methods. Modular GOAT [8] constructs an explicit semantic map, while DyNavLM [20] incorporates a graph-based memory architecture. To establish a strong value-map baseline, we also adapt VLFM [4] to the benchmark, carefully reconfiguring hyperparameters to ensure a fair comparison.

We evaluate four learning-based methods and eight zero-shot methods for the object-goal navigation task. Among them, ESC [10], L3MVN [22], and OpenFMNav [11] are semantic map-based; VLFM [4], InstructNav [5], and WMNav [21] are value map-based; and SG-Nav [12] and UniGoal [6] are scene graph-based.

4) *Configuration Details*: The agent in Habitat is modeled as a cylinder with a radius of 0.17 m and a height of 1.44 m, following [8]. An RGB camera is mounted at a height of 1.31 m, with a resolution of 640×480 and a horizontal field of view of 79° . A depth camera with the same intrinsics and extrinsics is also mounted, with a valid range of 0.5–5.0 m. The action space is {forward, turn left, turn right, stop}; each forward action moves 0.25 m, and each rotation is 30° . A task succeeds when the agent issues stop and is within 0.5 m of any candidate target viewpoint. Each task is capped at 500 time steps.

B. Benchmark Results

1) *Performance on lifelong multimodal object navigation task*: The quantitative results are shown in Table I, where boldface and underlining indicate the best and second-best results, respectively. The methods marked with \dagger (SenseAct-NN Monolithic and VLFM) represent our reimplementations, which include supplementary e-SR metrics to enhance analytical comprehensiveness. Observing the results, VLFM achieves a strong SR, evidencing its exploration capability; however, its SPL remains low due to the lack of persistent memory. In contrast, Modular GOAT benefits from its memory module and attains a substantially higher SPL than VLFM, albeit with a lower SR. Our method attains the best performance across all metrics and clearly surpasses the baselines, with particularly large gains in SPL. These results indicate that SSMG enables effective abstraction and retention of historical observations, allowing the agent to exploit past information during lifelong execution and leverage LLMs’ general knowledge and reasoning for long-horizon decision-making, thereby achieving both high success and high efficiency. On the other hand, the overall e-SR remains relatively low for all methods, underscoring that lifelong multimodal object navigation is still highly challenging and merits further investigation.

We further performed a statistical analysis of the performance of multimodal tasks, as shown in Fig. 2. Re-

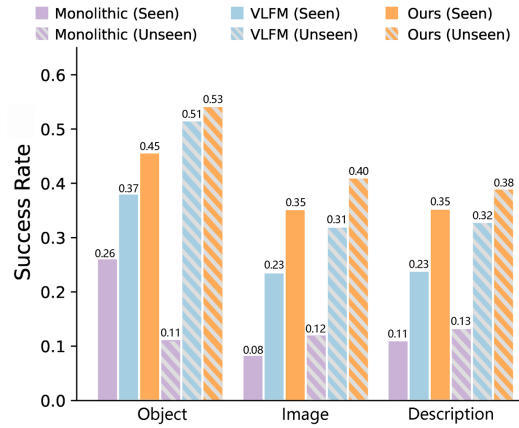


Fig. 2: Performance over different modalities.




Fig. 3: Performance over time.

sults demonstrate the superiority of our approach across all modalities. Notably, tasks utilizing the object modality generally underperform compared to those with image and description modalities, likely because instance-level object navigation requires a deeper understanding and analysis of target objects.

To investigate the characteristics of different approaches on lifelong tasks, we compute performance statistics as a function of subtask index; the results are shown in Fig. 3. We first observe that our method consistently outperforms the baselines across almost all subtask orders, reinforcing the conclusions drawn from the previous figures and tables. Moreover, both our method and the Monolithic baseline exhibit performance gains over subtasks (i.e., over time), relative to their initial states, with especially notable improvements in SPL on the *GOAT val_unseen* split for our method. As discussed in Sec. IV-A.3, the RL-based Monolithic approach implicitly encodes map information, whereas our approach explicitly abstracts historical observations into SSMG to assist decision making; thus, performance improvement over time is an expected and consistent finding. In contrast, VLFM outperforms Monolithic overall with general knowledge of VLMs, but it does not exhibit clear improvement over time, further underscoring the importance of memory in lifelong tasks, whether implemented explicitly or implicitly.

2) *Performance on object navigation task*: The quantitative results in Table II show that our method consistently outperforms in SPL while attaining comparable SR to strong

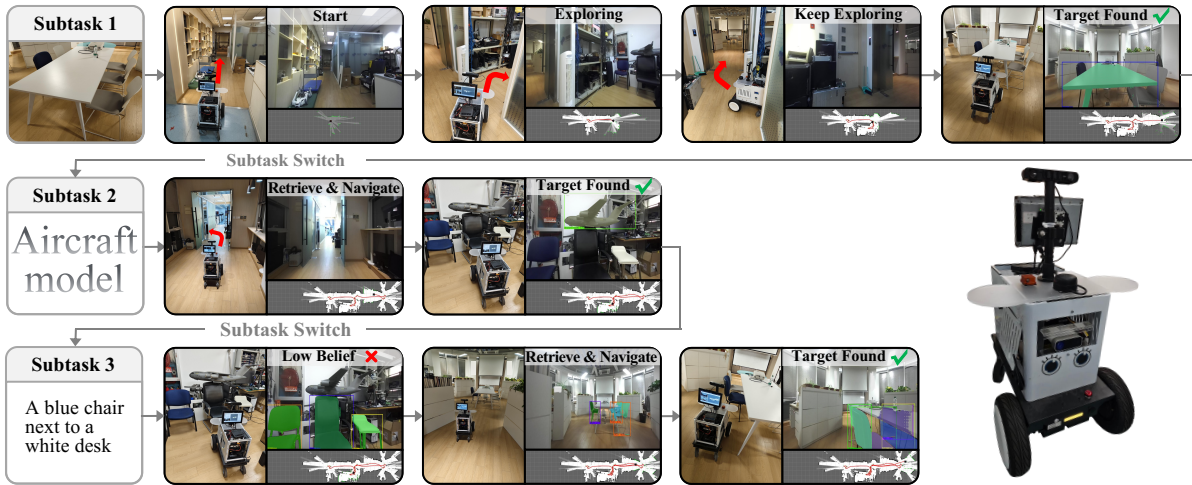


Fig. 4: Real-world demonstration of SSMG-Nav in an office environment with three subtasks.

TABLE III: Ablation Study.

Method		s-SR	e-SR	SPL
Map	Planner			
Semantic Map	Greedy	0.389	0.056	0.278
Value Map	Greedy	0.442	0.083	0.258
SSMG	Greedy	0.453	0.056	0.328
SSMG	LHP w/o revisiting	0.479	0.083	0.361
SSMG	LHP	0.502	0.083	0.376

baselines, unlike on GOAT-Bench. We posit that the non-lifelong nature of this benchmark reduces the marginal benefit of persistent memory. We further attribute our SPL advantage primarily to the proposed long-horizon planner. Notably, the best-performing zero-shot method, WMNav, ingests a panoramic observation composed of six cameras at each step, whereas our approach relies solely on a single camera.

C. Ablation Study

In the ablation study, we evaluate how map representations and planner strategies affect performance. Results on a downsampled subset of *GOAT val_unseen* are reported in Table III. ‘Semantic Map’ records only object semantics. ‘Value Map’ relies solely on the value map derived from the current target (i.e., no reusable long-term memory). ‘Greedy’ selects the area with highest value or nearest distance. Despite lower s-SR and e-SR than “Value Map,” the simplest “Semantic Map” method attains a higher SPL by leveraging memory, highlighting the importance of persistent memory for lifelong tasks. For our method, combining ‘SSMG’ with ‘Greedy’ yields substantial gains over other map formats, demonstrating the advantage of SSMG as a memory representation. Moreover, performance improves with ‘LHP w/o revisiting’ and improves further with full ‘LHP’, indicating a progressive benefit and validating our module design. The full version of our method achieves the best overall results.

D. Real-world deployment

Our method is deployed on a wheeled robotic platform. A ZED stereo camera provides RGB and depth data, and a 2D LiDAR supplies localization via Cartographer. We adopt [39] as the local planner to navigate to waypoints. The experiment is conducted in an office environment with three subtasks: (i) an image-modality target specified as a “white desk,” (ii) an object-modality target specified as an aircraft model, and (iii) a text-modality target described as “a blue chair next to a white desk.” An overview of the process is shown in Fig. 4. This indicates that our method is able to perform lifelong multimodal object navigation on physical platforms.

V. CONCLUSION

In this paper, we introduced SSMG-Nav, a novel lifelong multimodal object navigation framework based on Semantic Skeleton Memory Graph (SSMG) that consolidates historical observations into a spatially grounded, reusable memory. By anchoring semantics to topological keypoints and clustering nearby entities into subgraphs, SSMG yields compact, decision-ready candidate destinations. Building on this representation, we synthesize multimodal prompts to elicit target beliefs from a VLM and plan with a long-horizon planner (LHP) that balances belief and traversability to minimize expected path length. Experiments demonstrate that our method surpasses strong baselines on lifelong multimodal object navigation and standard object navigation, achieving substantial improvements in the challenging lifelong setting. We further confirm real-world feasibility.

However, several limitations remain. First, the approach relies on large models, which introduces latency and constrains runtime speed in real-world deployments; as a result, the robot operates at relatively low velocities. Second, the current method does not handle multi-floor scenarios. A promising direction is to extend SSMG with a voxel-based representation to build a hierarchical SSMG.

REFERENCES

- [1] C. Lu, M. Liu, Z. Luan, Y. He, and B. Chen, "Multi-view spatial context and state constraints for object-goal navigation," *IEEE Robot. Automat. Lett.*, vol. 10, no. 3, pp. 2207–2214, 2025.
- [2] R. Ramrakhya, E. Undersander, D. Batra, and A. Das, "Habitat-web: Learning embodied object-search strategies from human demonstrations at scale," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 5163–5173.
- [3] S. Y. Gadre, M. Wortsman, G. Ilharco, L. Schmidt, and S. Song, "Cows on pasture: Baselines and benchmarks for language-driven zero-shot object navigation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2023, pp. 23 171–23 181.
- [4] N. Yokoyama, S. Ha, D. Batra, J. Wang, and B. Bucher, "Vlfn: Vision-language frontier maps for zero-shot semantic navigation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2024.
- [5] Y. Long, W. Cai, H. Wang, G. Zhan, and H. Dong, "Instructnav: Zero-shot system for generic instruction navigation in unexplored environment," in *Proc. Conf. Robot Learn.*, 2024.
- [6] H. Yin, X. Xu, L. Zhao, Z. Wang, J. Zhou, and J. Lu, "Unigoal: Towards universal zero-shot goal-oriented navigation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2025.
- [7] M. Chang, T. Gervet, M. Khanna, S. Yenamandra, D. Shah, S. Y. Min, K. Shah, C. Paxton, S. Gupta, D. Batra, *et al.*, "Goat: Go to any thing," *arXiv preprint arXiv:2311.06430*, 2023.
- [8] M. Khanna*, R. Ramrakhya*, G. Chhablani, S. Yenamandra, T. Gervet, M. Chang, Z. Kira, D. S. Chaplot, D. Batra, and R. Motlaghi, "Goat-bench: a benchmark for multi-modal lifelong navigation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024.
- [9] S. K. Ramakrishnan, D. S. Chaplot, Z. Al-Halah, J. Malik, and K. Grauman, "Poni: Potential functions for objectgoal navigation with interaction-free learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2022, pp. 18 868–18 878.
- [10] K. Zhou, K. Zheng, C. Pryor, Y. Shen, H. Jin, L. Getoor, and X. E. Wang, "Esc: exploration with soft commonsense constraints for zero-shot object navigation," in *Proc. Int. Conf. Mach. Learn.*, vol. 202, Honolulu, Hawaii, USA, July 2023, pp. 42 829–42 842.
- [11] Y. Kuang, H. Lin, and M. Jiang, "Openfmnav: Towards open-set zero-shot object navigation via vision-language foundation models," in *Findings Assoc. Comput. Linguist.: NAACL 2024*, Mexico City, Mexico, June 2024, pp. 338–351.
- [12] H. Yin, X. Xu, Z. Wu, J. Zhou, and J. Lu, "Sg-nav: Online 3d scene graph prompting for llm-based zero-shot object navigation," in *Proc. Adv. Neural Inform. Process. Syst.*, vol. 37, 2024, pp. 5285–5307.
- [13] Z. Zhan, L. Yu, S. Yu, and G. Tan, "Mc-gpt: Empowering vision-and-language navigation with memory map and reasoning chains," *arXiv preprint arXiv:2405.10620*, 2024.
- [14] H. Niu, X. Ji, L. Zhang, F. Wen, R. Ying, and P. Liu, "A skeleton-based topological planner for exploration in complex unknown environments," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2025, pp. 11 766–11 772.
- [15] S. Zhang, X. Yu, X. Song, X. Wang, and S. Jiang, "Imagine before go: Self-supervised generative map for object goal navigation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Seattle, WA, USA, June 2024, pp. 16 414–16 425.
- [16] A. Majumdar, G. Aggarwal, B. Devnani, J. Hoffman, and D. Batra, "Zson: Zero-shot object-goal navigation using multimodal goal embeddings," in *Proc. Adv. Neural Inform. Process. Syst.*, vol. 35, 2022, pp. 32 340–32 352.
- [17] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
- [18] S. Nasiriany, F. Xia, W. Yu, T. Xiao, J. Liang, I. Dasgupta, A. Xie, D. Driess, A. Wahid, Z. Xu, Q. Vuong, T. Zhang, T.-W. E. Lee, K.-H. Lee, P. Xu, S. Kirmani, Y. Zhu, A. Zeng, K. Hausman, N. Heess, C. Finn, S. Levine, and B. Ichter, "Pivot: iterative visual prompting elicits actionable knowledge for vlms," in *Proc. Int. Conf. Mach. Learn.*, 2024.
- [19] D. Goetting, H. G. Singh, and A. Loquercio, "End-to-end navigation with vision-language models: Transforming spatial reasoning into question-answering," in *Proc. Int. Conf. Neuro-Symbolic Syst.*, ser. Proceedings of Machine Learning Research, G. Pappas, P. Ravikumar, and S. A. Seshia, Eds., vol. 288. PMLR, 28–30 May 2025, pp. 22–35.
- [20] Z. Ji, H. Lin, and Y. Gao, "Dynavlm: Zero-shot vision-language navigation system with dynamic viewpoints and self-refining graph memory," *arXiv preprint arXiv:2506.15096*, 2025.
- [21] D. Nie, X. Guo, Y. Duan, R. Zhang, and L. Chen, "Wmnav: Integrating vision-language models into world models for object goal navigation," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2025, pp. 2392–2399.
- [22] B. Yu, H. Kasaei, and M. Cao, "L3mvn: Leveraging large language models for visual target navigation," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2023, pp. 3554–3560.
- [23] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," in *Proc. Adv. Neural Inform. Process. Syst.*, vol. 35, 2022, pp. 24 824–24 837.
- [24] Q. Gu, A. Kuwajerwala, S. Morin, K. M. Jatavallabhula, B. Sen, A. Agarwal, C. Rivera, W. Paul, K. Ellis, R. Chellappa, C. Gan, C. M. De Melo, J. B. Tenenbaum, A. Torralba, F. Shkurti, and L. Paull, "Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2024, pp. 5021–5028.
- [25] A. Rajvanshi, K. Sikka, X. Lin, B. Lee, H.-p. Chiu, and A. Velasquez, "Saynav: Grounding large language models for dynamic planning to navigation in new environments," in *Proc. Int. Conf. Autom. Plan. Sched.*, Mar. 2024.
- [26] Y. Tang, M. Wang, Y. Deng, Z. Zheng, J. Deng, S. Zuo, and Y. Yue, "Openin: Open-vocabulary instance-oriented navigation in dynamic domestic environments," *IEEE Robot. Automat. Lett.*, vol. 10, no. 9, pp. 9256–9263, 2025.
- [27] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 19730–19742.
- [28] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su, *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2024, pp. 38–55.
- [29] C. Zhang, D. Han, Y. Qiao, J. U. Kim, S.-H. Bae, S. Lee, and C. S. Hong, "Faster segment anything: Towards lightweight sam for mobile applications," *arXiv preprint arXiv:2306.14289*, 2023.
- [30] G. Zhao, G. Li, W. Chen, and Y. Yu, "Over-nav: Elevating iterative vision-and-language navigation with open-vocabulary detection and structured representation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Seattle, WA, USA, June 2024, pp. 16 296–16 306.
- [31] B. Chen, Z. Xu, S. Kirmani, B. Ichter, D. Sadigh, L. Guibas, and F. Xia, "Spatialvlm: Endowing vision-language models with spatial reasoning capabilities," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Seattle, WA, USA, June 2024, pp. 14 455–14 465.
- [32] M. Aghzal, E. Plaku, G. J. Stein, and Z. Yao, "A survey on large language models for automated planning," *arXiv preprint arXiv:2502.12435*, 2025.
- [33] J. Zhang, L. Tang, Y. Song, Q. Meng, H. Qian, J. Shao, W. Song, S. Zhu, and J. Gu, "Fltrnn: Faithful long-horizon task planning for robotics with large language models," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2024, pp. 6680–6686.
- [34] J. Bai, S. Bai, S. Yang, S. Wang, S. Tan, P. Wang, J. Lin, C. Zhou, and J. Zhou, "Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond," *arXiv preprint arXiv:2308.12966*, 2023.
- [35] G. A. Croes, "A method for solving traveling-salesman problems," *Operations research*, vol. 6, no. 6, pp. 791–812, 1958.
- [36] S. K. Ramakrishnan, A. Gokaslan, E. Wijmans, O. Maksymets, A. Clegg, J. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang, *et al.*, "Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai," *arXiv preprint arXiv:2109.08238*, 2021.
- [37] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. S. Chaplot, O. Maksymets, *et al.*, "Habitat 2.0: Training home assistants to rearrange their habitat," in *Proc. Adv. Neural Inform. Process. Syst.*, vol. 34, 2021, pp. 251–266.
- [38] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niebner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," in *Proc. Int. Conf. 3D Vis.*, 2017, pp. 667–676.
- [39] J. Zhang, C. Hu, R. G. Chadha, and S. Singh, "Falco: Fast likelihood-based collision avoidance with extension to human-guided navigation," *Journal of Field Robotics*, vol. 37, no. 8, pp. 1300–1313, Dec. 2020.