FORESIGHTNAV: PERSISTENT BELIEF-AWARE PLANNING FOR MULTI-OBJECT NAVIGATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-Object Navigation (MON) requires an embodied agent to reach an ordered sequence of semantic target categories within a continuous episode under partial observability. When a target category has multiple physical instances, the current-instance choice determines not only the immediate navigation cost, but also the agent's subsequent starting state and the scene evidence available for the remaining goals. However, existing MON methods typically center decision-making on the current goal, without explicitly modeling how alternative current-instance choices influence the completion cost of the remaining ordered task. Thus, we propose ForesightNav, a persistent belief-aware planning framework that turns accumulated cross-target scene evidence into persistent candidate states and explicitly looks ahead over the remaining goal sequence. ForesightNav builds an online open-vocabulary semantic 3DGS scene memory and jointly queries the current and future targets. It then uses Persistent Candidate Cluster Query to associate observations across time and maintain candidate states according to semantic evidence, observation sufficiency, and visitation history. Finally, Belief-Aware Long-Horizon Planning converts these states into candidate beliefs and evaluates each current-instance choice using both its immediate navigation cost and the expected cost of completing the remaining goals. Experiments on HM3D have demonstrated the effectiveness of the proposed method.

1 INTRODUCTION

Multi-Object Navigation (MON) Zeng et al. (2023; 2026); Wani et al. (2020) requires an embodied agent to locate and reach an ordered sequence of semantic target categories within a single continuous episode under partial observability. Unlike independent ObjectNav tasks, MON couples decisions across targets within a single episode. Each target category may have multiple physical instances, and reaching the selected instance determines where the agent begins pursuing the remaining targets. As illustrated in Fig. 1, different instances of the current target lead to different continuation costs: the instance with the shortest immediate path does not necessarily minimize the trajectory over the full target sequence. This creates a long-horizon coupling between current instance selection and future navigation. The key challenge is therefore to make locally effective decisions without sacrificing the efficiency of completing the remaining goals.

Existing MON methods improve sequential navigation through semantic memory, spatial reasoning, target prediction, and hybrid exploration and planning Wani et al. (2020); Marza et al. (2023; 2022); Zeng et al. (2023; 2026); Sadek et al. (2023). More recent approaches further employ persistent open-vocabulary spatial representations to accumulate semantic evidence across observations and support target localization or verification Peng et al. (2023); Zhou et al. (2025). However, their decision processes are still typically organized around the currently active goal: persistent scene knowledge is primarily used to determine where to search for or verify that goal, rather than to compare how different current-instance choices affect the remaining ordered task. Consequently, when multiple instances of the current target are available, there is limited explicit modeling of both their uncertain validity and the expected navigation cost of completing the subsequent target sequence.

Addressing this gap requires a decision state that captures both the possible locations of current and future targets and the strength of the evidence supporting each candidate. Under partial observabil-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

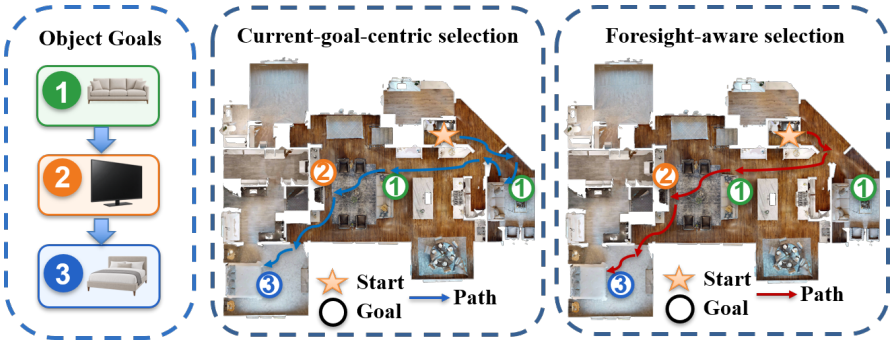


Figure 1: **Current-goal-centric versus Foresight-aware selection in MON.** Current-goal-centric selection minimizes the cost of reaching the current target, but may increase the route cost of subsequent goals. Foresight-aware selection evaluates current instances according to their consequences for the remaining ordered task, potentially producing a shorter overall route.

ity, semantic evidence is distributed across observations collected at different times, viewpoints, and goal stages Marza et al. (2022); Peng et al. (2023); Chen et al. (2022), producing transient target hypotheses that may be incomplete or inconsistent. Simply retaining these hypotheses is insufficient, because occlusion, viewpoint variation, and perceptual noise can produce weakly supported or unstable hypotheses. The agent must therefore associate transient hypotheses over time into persistent candidate states that maintain semantic evidence, observation sufficiency, and inspection history. These states distinguish insufficiently observed candidates from those that have been explicitly inspected and rejected, while providing the basis for estimating candidate validity under uncertainty Kaelbling et al. (1998); Ross et al. (2008). This uncertainty directly affects the consequence of each current-instance choice: reaching a candidate incurs an immediate travel cost; if the candidate is valid, the agent completes the current goal and proceeds to the remaining targets, whereas if it is invalid, the current goal remains unresolved and the agent must continue searching from the reached location. Thus, current-instance selection should minimize the expected cost of completing the remaining ordered suffix, rather than considering only the distance to the active goal.

To realize this decision principle, we propose ForesightNav, a persistent belief-aware planning framework for multi-object navigation. ForesightNav first builds an online open-vocabulary semantic 3DGS scene memory and jointly queries the current and remaining targets to expose spatial evidence for the remaining ordered goal sequence. Persistent Candidate Cluster Query (PCQ) associates transient hypotheses across time into persistent candidates and maintains their semantic evidence, observation sufficiency, and inspection state, allowing candidates to be rejected or reactivated as evidence changes. Belief-Aware Long-Horizon Planning (BLP) then converts the eligible candidate states into validity beliefs. It evaluates each current-instance choice by combining its immediate travel cost with the continuation cost after successful completion and the recovery cost after an invalid candidate. Through receding-horizon replanning, ForesightNav continuously updates candidate states, validity beliefs, and expected suffix costs as new observations become available. Experiments on HM3D Ramakrishnan et al. (2021) demonstrate the effectiveness of ForesightNav.

2 Related Works

Multi-Object Navigation and Persistent Memory. Multi-Object Navigation (MON) requires an agent to complete an ordered sequence of semantic goals, placing particular demands on persistent scene understanding and cross-goal memory Wani et al. (2020). Existing approaches improve MON through spatial or implicit memory Marza et al. (2022; 2023), active perception Chen et al. (2022), target-position modeling and dynamic optimization Zeng et al. (2023; 2026), and hybrid exploration policies Sadek et al. (2023). In parallel, open-vocabulary maps and 3DGS-based representations provide persistent spatial-semantic memories that support long-term navigation Huang et al. (2023); Peng et al. (2023); Lei et al. (2025); Gao et al. (2025); Zheng et al. (2026).

Most closely related, LagMemo Zhou et al. (2025) uses a language-embedded 3DGS memory to retrieve and verify target candidates across multi-goal navigation. However, candidate retrieval and

verification alone do not address the decision coupling introduced by ordered goals: when multiple plausible instances exist for the current target, choosing one determines not only the immediate navigation cost but also the agent's starting state for the remaining goals. ForesightNav therefore focuses on this decision-level problem. It maintains persistent candidate states that distinguish insufficiently observed, validated, and rejected hypotheses, estimates their validity under uncertainty, and evaluates each current-instance choice according to its expected consequence over the remaining ordered task. In this way, persistent memory is used not only to recall where potential targets are, but also to support belief-aware long-horizon instance selection.

Planning under Perceptual Uncertainty. Nonmyopic view planning and active visual search account for future observations and uncertain detections when selecting sensing or search actions Atanasov et al. (2014); Taioli et al. (2024). Recent ObjectNav methods similarly use active perception and semantic uncertainty to improve current-target search and verification Qin et al. (2025); Bajpai et al. (2025); Pei et al. (2026). ForesightNav instead maintains an explicit validity belief for each persistent candidate and evaluates its expected remaining-task cost by combining immediate travel, continuation over future goals after a valid selection, and continued search after an invalid selection. These beliefs and costs are recomputed online as the scene memory and candidate states evolve.

3 METHOD

3.1 METHOD OVERVIEW

As illustrated in Fig. 2, ForesightNav integrates persistent scene understanding, candidate-state maintenance, and long-horizon planning in a closed-loop navigation framework. Given RGB-D observations, agent poses, and the ordered goal sequence, an online semantic 3DGS memory incrementally reconstructs the scene and jointly queries the current and remaining targets. PCQ associates the resulting transient hypotheses across time into persistent candidate states that maintain semantic evidence, observation sufficiency, and inspection history. BLP then converts eligible candidates into validity beliefs and evaluates each current-instance choice according to its immediate travel cost and expected suffix cost under success and failure. The selected candidate is executed for a short horizon, after which new observations update the scene memory, candidate states, and beliefs, enabling rolling replanning throughout the episode.

3.2 ONLINE SEMANTIC 3DGS SCENE MEMORY

Long-horizon MON planning requires spatial evidence about both the current and future goals. However, such evidence is observed incrementally from different viewpoints and may become relevant only after the active goal changes. We therefore construct an online semantic 3DGS memory that preserves geometric and open-vocabulary semantic observations throughout the episode. The memory is built during task execution without requiring a preconstructed scene map Lei et al. (2025); Zhou et al. (2025).

Online semantic 3DGS construction. At navigation step t, a feed-forward reconstructor Wang et al. (2026) converts the current RGB-D observation into a set of local Gaussian primitives. Using the estimated camera pose, these primitives are transformed into the world coordinate system and fused with the global Gaussian set $\mathcal{G}_{t-1}$, according to spatial consistency.. Matched primitives update the existing geometry, whereas unmatched primitives are inserted into the map. We further associate the Gaussians with open-vocabulary visual features following the sparse semantic representation used in EmbodiedSplat Lee et al. (2026). The scene memory is represented as

$$\mathcal{M}_t = (\mathcal{G}_t, \mathcal{F}_t, \mathcal{A}_t)\,, \tag{1}$$

where $\mathcal{G}_t$ is the global Gaussian geometry, $\mathcal{F}_t$ is the open-vocabulary feature bank, and $\mathcal{A}_t$ stores the weighted associations between Gaussian primitives and semantic features. Because $\mathcal{M}_t$ is retained across goal transitions, observations collected for earlier targets remain available during subsequent stages.

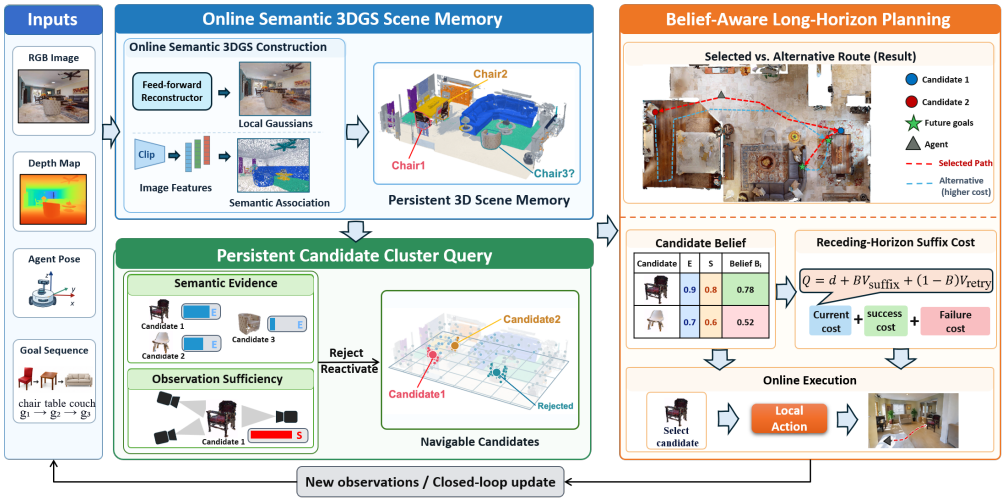


Figure 2: **Overview of ForesightNav.** The online semantic 3DGS memory incrementally reconstructs the scene and queries the current and remaining targets. PCQ associates transient hypotheses into persistent candidates with semantic evidence, observation sufficiency, and rejection state. BLP estimates candidate validity beliefs and selects the current target instance by minimizing its expected long-horizon task cost, which combines immediate travel, success-conditioned continuation, and failure-conditioned recovery. Short-horizon execution and rolling updates enable continuous replanning as new observations become available.

Cross-target semantic query. At target stage k, we query all remaining goals $g_{k:N} = (g_k, \ldots, g_N)$ from the same scene memory. Let $\mathbf{f}_{a_{i\ell}} \in \mathcal{F}_t$ be the ℓ-th semantic feature associated with Gaussian i, and let $w_{i\ell}$, obtained from $\mathcal{A}_t$, denote its association weight, and $\mathbf{e}(g_j)$ the embedding of goal g_j in the shared vision–language space. The semantic relevance of Gaussian i to goal g_j is

$$
\begin{aligned}
s_{t,i}(g_j) &= \sum_{\ell=1}^{L_i} \bar{w}_{i\ell} \cos(\mathbf{f}_{a_{i\ell}}, \mathbf{e}(g_j)) \,, \\
\bar{w}_{i\ell} &= \frac{w_{i\ell}}{\sum_{q=1}^{L_i} w_{iq} + \epsilon}, \qquad j = k, \ldots, N,
\end{aligned}
\tag{2}
$$

where L_i is the number of semantic features associated with Gaussian i. This weighted aggregation allows multiple observations of the same Gaussian to contribute to its target relevance.

Transient candidate generation. For each remaining goal, we retain Gaussians with sufficient semantic and geometric support, group them into object-level hypotheses through 3D spatial clustering, and merge overlapping clusters. This produces

$$
\mathcal{H}_j^t = \{h_{j,m}^t\}_{m=1}^{M_j^t}, \qquad j = k, \ldots, N,
\tag{3}
$$

where $h_{j,m}^t$ is a transient spatial hypothesis for goal g_j. By querying the complete remaining goal sequence, the module exposes spatial evidence for both the active target and future targets. The next module temporally associates these transient hypotheses and converts them into persistent candidate states for belief estimation and remaining-task-aware planning.

3.3 Persistent Candidate Cluster Query

The transient hypotheses $\mathcal{H}_j^t$ produced by the semantic scene memory are reconstructed whenever the map is updated. Consequently, they do not directly preserve whether a region has been insufficiently observed, explicitly inspected, or previously rejected. Treating every query result as a

new candidate may therefore cause repeated visits to the same unsupported region or the premature removal of a candidate observed under unfavorable conditions. We address this problem with Persistent Candidate Cluster Query (PCQ), which associates transient hypotheses across time and maintains an explicit persistent state for each spatial candidate.

Persistent candidate association. For each remaining goal g_j, PCQ associates a transient hypothesis $h_{j,m}^t \in \mathcal{H}_j^t$ with an existing candidate of the same goal category when their spatial centers and 3D supports are consistent. Matched hypotheses update the corresponding candidate, while unmatched hypotheses initialize new candidates. The precise association gates and assignment procedure are provided in Appendix Section 4.1.1.

The n-th persistent candidate for goal g_j is represented as

$$\mathcal{K}_{j,n}^t = \left(\Omega_{j,n}^t, \mathbf{x}_{j,n}^t, E_{j,n}^t, S_{j,n}^t, R_{j,n}^t\right), \tag{4}$$

where $\Omega_{j,n}^t$ is its supporting Gaussian set, $\mathbf{x}_{j,n}^t$ is its spatial center, $E_{j,n}^t$ denotes target-semantic evidence, $S_{j,n}^t$ measures observation sufficiency, and $R_{j,n}^t \in \{0,1\}$ records whether the candidate has been explicitly rejected. Newly initialized candidates satisfy $R_{j,n}^t = 0$.

Target-semantic evidence. The semantic evidence of a persistent candidate is obtained by aggregating the goal relevance of its supporting Gaussians:

$$E_{j,n}^t = \frac{1}{|\Omega_{j,n}^t|} \sum_{i \in \Omega_{j,n}^t} s_{t,i}(g_j), \tag{5}$$

where $s_{t,i}(g_j)$ is defined in equation 2. Because the underlying Gaussian memory is continuously updated, $E_{j,n}^t$ may increase or decrease as new observations provide additional semantic evidence.

Observation sufficiency. Semantic relevance alone does not indicate whether a candidate has been observed reliably. A high response from a single occluded or distant view should be treated differently from a response that is consistently supported across independent viewpoints. We therefore maintain a viewpoint-based observation-sufficiency score.

Let N_i^t denote the number of valid, sufficiently distinct viewpoints from which Gaussian i has been observed by time t. An observation is counted only when the Gaussian has sufficient visibility and rendering contribution, and its viewing direction differs from previously counted viewpoints by at least a predefined angular threshold. We map this count to a bounded support value:

$$r_i^t = 1 - \exp\left(-\frac{N_i^t}{\tau_{\text{cov}}}\right), \tag{6}$$

where τ_{cov} controls the saturation rate. Candidate-level observation sufficiency is then

$$S_{j,n}^t = \frac{1}{|\Omega_{j,n}^t|} \sum_{i \in \Omega_{j,n}^t} r_i^t. \tag{7}$$

For a fixed supporting set, $S_{j,n}^t \in [0,1]$ increases as the candidate region is observed from additional sufficiently distinct viewpoints, rather than merely when the agent collects many adjacent frames. When the supporting set changes after association, the score is recomputed and may either increase or decrease.

Rejection and reactivation. To avoid repeatedly revisiting unsupported regions, a candidate is rejected only after it has been explicitly inspected while g_j is the active goal. If the candidate is sufficiently observed but its semantic evidence remains weak,

$$S_{j,n}^t \geq \tau_S, \qquad E_{j,n}^t < \tau_E, \tag{8}$$

we set $R^t_{j,n} = 1$ and record its current evidence as $E^{\text{rej}}_{j,n}$. Rejected candidates are excluded from planning but continue to receive evidence from new observations. The explicit inspection event, complete rejection–reactivation state transition, and evidence update for rejected candidates are detailed in Appendix Section 4.1.2. A candidate is reactivated by resetting $R^t_{j,n} = 0$ when

$$E^t_{j,n} - E^{\text{rej}}_{j,n} \geq \tau_{\text{react}}. \tag{9}$$

This mechanism avoids repeated inspection of rejected regions while allowing candidates to recover when substantially stronger evidence becomes available.

Finally, the candidate set eligible for long-horizon planning is

$$\mathcal{K}^t_{k:N,\text{elig}} = \left\{ \mathcal{K}^t_{j,n} \;\middle|\; j \in \{k, \ldots, N\},\ R^t_{j,n} = 0 \right\}. \tag{10}$$

3.4 Belief-Aware Long-Horizon Planning

In MON, the choice of the current target instance not only determines the immediate navigation distance, but also changes the agent's spatial state after completing the current goal, thereby affecting the cost of visiting the remaining ordered targets. We therefore introduce **Belief-Aware Long-Horizon Planning** (BLP). Rather than greedily selecting the current candidate according to local distance, the planner jointly evaluates each candidate according to its validity and its long-term consequences for the remaining target sequence $g_{k+1:N}$.

Candidate Validity belief. Candidates that have not been explicitly rejected may still differ in reliability. We therefore estimate candidate validity from its semantic evidence $E^t_{j,n}$ and observation sufficiency $S^t_{j,n}$. First, semantic evidence is converted into a calibrated target-support score:

$$q^t_{j,n} = \sigma\left(aE^t_{j,n} + b_0\right), \tag{11}$$

where $\sigma(\cdot)$ is the sigmoid function and a, b_0 are learned through logistic calibration. We then define the candidate validity belief as

$$B^t_{j,n} = S^t_{j,n} q^t_{j,n} + \left(1 - S^t_{j,n}\right) \pi_j, \tag{12}$$

where $\pi_j \equiv \pi$ is a global validity prior shared across target categories. With sufficient observation, $B^t_{j,n}$ is dominated by calibrated semantic evidence; otherwise, it remains closer to the prior. The parameters a, b_0, and π are estimated only from the training calibration split and fixed during evaluation; details are provided in Appendix Section 4.2.

Receding-Horizon Expected Suffix Cost. Given the candidate-validity belief $B^t_{j,n}$, we couple candidate reliability with the navigation cost of the remaining ordered goals. At each replanning step, the scene memory, navigation graph, candidate sets, beliefs, and pairwise path costs are treated as a fixed planning snapshot when evaluating the continuation costs. After a short-horizon execution step, these quantities are updated using newly acquired observations and the suffix costs are recomputed.

For a reachable candidate $\mathcal{K}^t_{k,n}$ of the current target g_k, selecting this candidate first incurs the geodesic travel cost from the current agent position $\mathbf{p}_t$ to its candidate location $\mathbf{x}^t_{k,n}$. If the candidate is valid, the current target is completed and planning advances to $g_{k+1:N}$. Otherwise, the current target remains active, while candidate n is excluded only within the corresponding failure branch. We therefore define

$$Q^t_{k,n} = d_t\left(\mathbf{p}_t, \mathbf{x}^t_{k,n}\right) + B^t_{k,n} V^t_{k+1}\left(\mathbf{x}^t_{k,n}\right) + \left(1 - B^t_{k,n}\right) V^{t,-n}_k\left(\mathbf{x}^t_{k,n}\right), \tag{13}$$

where $d_t(\cdot,\cdot)$ is the geodesic path cost on the current navigation graph. The success-conditioned value $V^t_{k+1}(\mathbf{x})$ advances to the remaining goal sequence $g_{k+1:N}$, whereas the failure-conditioned value $V^{t,-n}_k(\mathbf{x})$ keeps g_k active after excluding candidate n.

Both continuation values are computed recursively over the reachable, non-rejected candidates in the current planning snapshot. Success advances the target index, failure removes the selected candidate, and an empty candidate set triggers nearest-frontier exploration followed by replanning. A branch with neither reachable candidates nor reachable frontiers is infeasible; the full recursion and boundary conditions are provided in Appendix Section 4.3.

At each planning update, we retain at most 5 reachable, non-rejected candidates for the current target g_k. We denote the resulting candidate set by $\mathcal{C}_k^t$; its complete construction is provided in Appendix Section 4.3.1. Among these candidates, we select

$$n^\star = \arg\min_{n \in \mathcal{C}_k^t} Q_{k,n}^t, \tag{14}$$

where all terms in Eq. equation 13 are expressed in geodesic path-length units and therefore require no additional weighting coefficient.

Online Execution and Rolling Replanning. After obtaining the selected candidate, the local planner executes only a short-horizon portion of the path toward $\mathbf{x}_{k,n^*}^t$. The resulting observations update the scene memory, navigation graph, persistent candidate states, and validity beliefs. ForesightNav then constructs a new planning snapshot and recomputes the costs. Consequently, candidate selection accounts for the currently estimated cost of the remaining goal sequence, while rolling replanning allows execution to adapt to newly observed geometry and semantic evidence.

4 Experiments

4.1 Experimental Setup

We evaluate ForesightNav on HM3D Ramakrishnan et al. (2021) to assess its effectiveness in multi-object navigation. ForesightNav requires no task-specific navigation-policy training or end-to-end fine-tuning; pretrained perception modules remain frozen during evaluation. All pretrained perception modules remain frozen during evaluation, while scene-memory construction, persistent candidate maintenance, and belief-aware planning are performed online. On HM3D, we consider 2-MON, 3-MON, and 4-MON tasks, in which the agent must reach target categories in a prescribed order within a continuous episode. Our **Baseline** uses only the semantic 3DGS scene map and performs candidate selection for the current target without persistent candidate modeling or long-horizon planning. We compare it against L3MVN Yu et al. (2023), VLFM Yokoyama et al. (2024), ApexNAV Zhang et al. (2025), and GDWO Zeng et al. (2026), using SR and SPL as the primary evaluation metrics. Further details on the evaluation metrics, implementation settings, additional experiments, and visualizations, are provided in the Appendix and Supplementary Material.

Table 1: **Multi-object Navigation Validation Results on HM3D-MON (%).**

	2-MON		3-MON		4-MON	
Method	SR	SPL	SR	SPL	SR	SPL
L3MVN (Yu et al., 2023)	25.83	9.98	16.00	5.43	8.67	2.82
VLFM (Yokoyama et al., 2024)	29.67	15.13	17.00	7.34	8.67	3.24
ApexNAV (Zhang et al., 2025)	**36.17**	**17.14**	20.25	8.15	10.67	4.23
GDWO (Zeng et al., 2026)	28.00	7.60	14.40	3.90	6.40	1.90
Baseline	21.67	8.82	12.00	4.73	8.67	3.67
ForesightNav	32.67	13.60	**23.75**	**9.65**	**12.00**	**5.28**

4.2 Comparison on HM3D-MON

Table 1 compares ForesightNav with representative navigation methods across different ordered target-sequence lengths. Compared with the corresponding 3DGS-map Baseline, ForesightNav consistently improves performance on 2-MON, 3-MON, and 4-MON. The gain is particularly pronounced on 3-MON, where SR/SPL increase from 12.00%/4.73% to 23.75%/9.65%, nearly doubling SR and more than doubling SPL. These results demonstrate the benefit of integrating persistent

candidate maintenance with explicit reasoning over the remaining target suffix for both complete-sequence success and path efficiency.

A clearer trend emerges when comparing with ApexNAV Zhang et al. (2025). ForesightNav remains below ApexNAV on 2-MON, suggesting that short sequences are still strongly influenced by single-object navigation capability. However, as the task horizon increases, ForesightNav surpasses ApexNAV on both 3-MON and 4-MON, with gains of 3.50/1.50 and 1.33/1.05 percentage points in SR/SPL, respectively. Notably, the underlying 3DGS-map baseline of ForesightNav is not stronger, suggesting that these gains are more likely attributable to explicit modeling of persistent cross-target candidates and long-term continuation costs rather than superior low-level navigation.

This trend is consistent with the motivation of ForesightNav. As the target sequence becomes longer, current-instance selection affects not only the immediate travel distance, but also the starting state, subsequent observations, and spatial relation to the remaining goals, leading to stronger cross-target coupling. By preserving candidate evidence across targets and jointly considering candidate validity and expected suffix cost, ForesightNav exhibits greater advantages on 3-MON and 4-MON tasks.

4.3 Ablation Studies

Table 2: **Ablation study of ForesightNav on HM3D-3MON task (%).** We progressively introduce each component to evaluate its individual contribution and the complementary gains from their combination.

Variant	3DGS Map	PCQ	BLP	SR	SPL
Baseline	✓			12.00	4.73
+ PCQ	✓	✓		19.25	6.19
+ BLP	✓		✓	14.25	5.62
Full ForesightNav	✓	✓	✓	**23.75**	**9.65**

Table 2 evaluates the contributions of PCQ and BLP on HM3D-3MON. Starting from the 3DGS-map baseline, adding PCQ improves SR/SPL from 12.00%/4.73% to 19.25%/6.19%, showing that persistently maintaining semantic evidence, observation sufficiency, and inspection history helps reduce repeated or unreliable candidate selection and provides more stable candidate states for subsequent decisions.

Adding BLP alone raises SR/SPL to 14.25%/5.62%, indicating that current-instance selection benefits from jointly modeling candidate validity, suffix continuation cost, and failure recovery cost. When combined with PCQ, Full ForesightNav reaches 23.75%/9.65%, with BLP further improving SR/SPL by 4.50%/3.46% over PCQ alone. These results show that PCQ provides reliable persistent candidate states, while BLP exploits them for suffix-aware decision-making, together accounting for the core performance gains of ForesightNav.

4.4 Long-Horizon Analysis

Table 3: **Stage-wise conditional success on HM3D-3MON task(%).** C_k denotes successful completion of the k-th target. Joint success is computed as $P(C_1)P(C_2 \mid C_1)P(C_3 \mid C_1, C_2)$.

Method	$P(C_1)$	$P(C_2 \mid C_1)$	$P(C_3 \mid C_1, C_2)$	Joint
L3MVN (Yu et al., 2023)	49.75	54.77	58.72	16.00
VLFM (Yokoyama et al., 2024)	51.25	58.05	57.14	17.00
ApexNAV (Zhang et al., 2025)	**57.25**	60.70	58.27	20.25
Baseline	36.75	62.59	52.17	12.00
Baseline + PCQ	45.75	65.57	64.17	19.25
Full ForesightNav	47.75	**64.92**	**76.61**	**23.75**

Stage-wise Conditional Success. To determine whether ForesightNav improves complete-sequence success through stronger immediate target navigation or better decisions for the remaining

targets, Table 3 decomposes the 3-MON success rate into stage-wise conditional probabilities. Compared with the Baseline, introducing PCQ increases $P(C_1)$ from 36.75% to 45.75%, $P(C_2 \mid C_1)$ from 62.59% to 65.57%, and $P(C_3 \mid C_1, C_2)$ from 52.17% to 64.17%. These consistent improvements show that persistently associating candidate observations and maintaining their inspection states benefits not only the current target but also subsequent navigation stages.

Adding BLP produces a distinct improvement pattern. From Baseline + PCQ to Full ForesightNav, $P(C_1)$ increases only slightly from 45.75% to 47.75%, while $P(C_2 \mid C_1)$ remains nearly unchanged. In contrast, $P(C_3 \mid C_1, C_2)$ rises from 64.17% to 76.61%, increasing joint success from 19.25% to 23.75%. Although ForesightNav has a lower first-stage success rate than ApexNAV Zhang et al. (2025) (47.75% vs. 57.25%), it achieves higher joint success (23.75% vs. 20.25%) and a 18.34% higher final-stage conditional success rate. These results suggest that the gains of ForesightNav emerge primarily in later stages, where preserved cross-target information and the long-term consequences of current-instance selection become more important, supporting the motivation of suffix-aware planning.

Table 4: **GOAT-Bench results under the standard sequential-goal protocol (%).** We report SR and SPL on the Seen, Synonyms, and Unseen splits. GOAT-GTSem† uses ground-truth semantic information and is included only as a privileged reference.

	Seen		Synonyms		Unseen	
Method	SR	SPL	SR	SPL	SR	SPL
GOAT-GTSem† (Khanna* et al., 2024)	56.7	40.3	58.4	43.5	54.3	41.0
Modular GOAT (Chang et al., 2024)	26.3	17.5	33.8	24.4	24.9	17.2
Modular CoWs (Gadre et al., 2023)	14.8	8.7	18.5	11.5	16.1	10.4
SenseAct-NN (SC) (Khanna* et al., 2024)	29.2	12.8	38.2	15.2	29.5	11.3
SenseAct-NN (Mono) (Khanna* et al., 2024)	16.8	9.4	18.5	10.1	12.3	6.8
TANGO (Ziliotto et al., 2025)	–	–	–	–	32.1	16.5
LagMemo (Zhou et al., 2025)	36.8	20.7	44.8	**32.1**	37.9	**26.3**
ForesightNav(w/o BLP)	**45.4**	**23.7**	**46.9**	24.6	**45.1**	18.4

4.5 Broader Evaluation on GOAT-Bench

We further evaluate ForesightNav on GOAT-Bench Khanna* et al. (2024) under its standard sequential-goal protocol, where only the current goal is revealed at each stage. We follow the benchmark's original instance-level success criterion. Since future goals are unavailable, BLP is disabled to ensure that no unrevealed target information is used. Accordingly, ForesightNav (w/o BLP) selects, from the current-goal candidates maintained by PCQ, the candidate with the highest semantic query score as the navigation target, without querying future goals or computing suffix costs.

As shown in Table 4, ForesightNav(w/o BLP) achieves the highest SR among methods without privileged semantics on Seen, Synonyms, and Unseen targets, reaching 45.4%, 46.9%, and 45.1%, respectively. These results exceed the strongest non-privileged baseline on each split by 8.6, 2.1, and 7.2 percentage points, supporting the effectiveness of persistent candidate maintenance under diverse semantic conditions rather than BLP. ForesightNav(w/o BLP) also improves Seen SPL over LagMemo Zhou et al. (2025) from 20.7% to 23.7%, but remains lower on Synonyms (24.6% vs. 32.1%) and Unseen (18.4% vs. 26.3%). Thus, under stronger semantic uncertainty, it primarily improves goal-reaching reliability rather than uniformly reducing path length.

5 Conclusion

We presented ForesightNav, a persistent belief-aware framework for Multi-Object Navigation. By maintaining persistent candidate states and jointly considering candidate validity and downstream task cost, ForesightNav enables more reliable long-horizon decision making. Experiments on HM3D-MON and GOAT-Bench demonstrate its effectiveness, while ablation studies confirm the complementary benefits of persistent candidate reasoning and foresight-aware planning.

AI USE STATEMENT

In this work, LLMs were used solely to assist with language editing, including improving grammar, clarity, and readability of the manuscript. They were not used to develop the research methodology, formulate hypotheses or mathematical claims, design experiments, implement the proposed method, or interpret experimental results. All AI-assisted edits were reviewed and verified by the authors, who take full responsibility for the final content of the paper.

REPRODUCIBILITY STATEMENT

We provide the information necessary to reproduce our method and experiments throughout the main paper and supplementary material. The method section specifies the construction of the online semantic 3DGS scene memory, persistent candidate cluster query, and belief-aware long-horizon planning. The experimental section describes the datasets, task settings, evaluation metrics, baselines, and implementation configurations used in our evaluation. Additional implementation details, hyperparameters, and experimental settings are provided in the appendix and supplementary material to facilitate reproduction of the reported results.

REFERENCES

Nikolay Atanasov, Bharath Sankaran, Jerome Le Ny, George J. Pappas, and Kostas Daniilidis. Nonmyopic view planning for active object classification and pose estimation. *IEEE Transactions on Robotics*, 30(5):1078–1090, 2014. doi: 10.1109/TRO.2014.2320795.

Utkarsh Bajpai, Julius Rückin, Cyrill Stachniss, and Marija Popović. Uncertainty-informed active perception for open vocabulary object goal navigation. In *2025 European Conference on Mobile Robots (ECMR)*, pp. 1–7, 2025. doi: 10.1109/ECMR65884.2025.11162987.

Matthew Chang, Theophile Gervet, Mukul Khanna, Sriram Yenamandra, Dhruv Shah, So Yeon Min, Kavit Shah, Chris Paxton, Saurabh Gupta, Dhruv Batra, Roozbeh Mottaghi, Jitendra Malik, and Devendra Singh Chaplot. Goat: Go to any thing. In *Robotics: Science and Systems (RSS)*, 2024. doi: 10.15607/RSS.2024.XX.073.

Peihao Chen, Dongyu Ji, Kunyang Lin, Weiwen Hu, Wenbing Huang, Thomas H Li, Mingkui Tan, and Chuang Gan. Learning active camera for multi-object navigation. *arXiv preprint arXiv:2210.07505*, 2022.

Samir Yitzhak Gadre, Mitchell Wortsman, Gabriel Ilharco, Ludwig Schmidt, and Shuran Song. 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)*, pp. 23171–23181, 2023.

Jianzhe Gao, Rui Liu, and Wenguan Wang. 3d gaussian map with open-set semantic grouping for vision-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9252–9262, October 2025.

Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 10608–10615, 2023. doi: 10.1109/ICRA48891.2023.10160969.

Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X.

Mukul Khanna*, Ram Ramrakhya*, Gunjan Chhablani, Sriram Yenamandra, Theophile Gervet, Matthew Chang, Zsolt Kira, Devendra Singh Chaplot, Dhruv Batra, and Roozbeh Mottaghi. Goat-bench: A benchmark for multi-modal lifelong navigation. In *CVPR*, 2024.

Seungjun Lee, Zihan Wang, Yunsong Wang, and Gim Hee Lee. Embodiedsplat: Online feed-forward semantic 3dgs for open-vocabulary 3d scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23774–23784, 2026.

Xiaohan Lei, Min Wang, Wengang Zhou, and Houqiang Li. Gaussnav: Gaussian splatting for visual navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Pierre Marza, Laetitia Matignon, Olivier Simonin, and Christian Wolf. Teaching agents how to map: Spatial reasoning for multi-object navigation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1725–1732. IEEE, 2022.

Pierre Marza, Laetitia Matignon, Olivier Simonin, and Christian Wolf. Multi-object navigation with dynamically learned neural implicit representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11004–11015, 2023.

Xuetong Pei, Jian Liu, Vidura Munasinghe, Bo Miao, U-Xuan Tan, Wenrui Ding, and Na Zhao. SAP-Nav: Spatial semantic representation meets active perception for hierarchical open-vocabulary object navigation, 2026.

Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, and Thomas Funkhouser. OpenScene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 815–824, 2023.

Liang Qin, Min Wang, Peiwei Li, Wengang Zhou, and Houqiang Li. Active perception meets rule-guided rl: A two-phase approach for precise object navigation in complex environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 7603–7612, 2025. doi: 10.1109/ICCV51701.2025.00713.

Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. URL https://arxiv.org/abs/2109.08238.

Stephane Ross, Joelle Pineau, Sebastien Paquet, and Brahim Chaib-draa. Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Research*, 32:663–704, 2008. doi: 10.1613/jair.2567.

Assem Sadek, Guillaume Bono, Boris Chidlovskii, Atilla Baskurt, and Christian Wolf. Multi-object navigation in real environments using hybrid policies. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4085–4091. IEEE, 2023.

Francesco Taioli, Francesco Giuliari, Yiming Wang, Riccardo Berra, Alberto Castellini, Alessio Del Bue, Alessandro Farinelli, Marco Cristani, and Francesco Setti. Unsupervised active visual search with monte carlo planning under uncertain detections. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):11047–11058, 2024. doi: 10.1109/TPAMI.2024.3451994.

Yunsong Wang, Tianxin Huang, Hanlin Chen, and Gim Hee Lee. FreeSplat++: Generalizable 3D gaussian splatting for efficient indoor scene reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(7):7749–7765, 2026. doi: 10.1109/TPAMI.2026.3665771. URL https://doi.org/10.1109/TPAMI.2026.3665771.

Saim Wani, Shivansh Patel, Unnat Jain, Angel Chang, and Manolis Savva. Multion: Benchmarking semantic map memory using multi-object navigation. *Advances in Neural Information Processing Systems*, 33:9700–9712, 2020.

Naoki Yokoyama, Sehoon Ha, Dhruv Batra, Jiuguang Wang, and Bernadette Bucher. Vlfm: Vision-language frontier maps for zero-shot semantic navigation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 42–48. IEEE, 2024. doi: 10.1109/ICRA57147.2024.10610712.

Bangguo Yu, Hamidreza Kasaei, and Ming Cao. L3mvn: Leveraging large language models for visual target navigation. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3554–3560. IEEE, 2023. doi: 10.1109/IROS55552.2023.10342512.

Haitao Zeng, Xinhang Song, and Shuqiang Jiang. Multi-object navigation using potential target position policy function. *IEEE Transactions on Image Processing*, 2023.

Haitao Zeng, Xinhang Song, and Shuqiang Jiang. Goal-oriented dynamic weight optimization for multi-object navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(6): 6452–6469, 2026. doi: 10.1109/TPAMI.2026.3657778.

Mingjie Zhang, Yuheng Du, Chengkai Wu, Jinni Zhou, Zhenchao Qi, Jun Ma, and Boyu Zhou. Apexnav: An adaptive exploration strategy for zero-shot object navigation with target-centric semantic fusion. *IEEE Robotics and Automation Letters*, 10(11):11530–11537, 2025. doi: 10.1109/LRA.2025.3606388.

Wancai Zheng, Hao Chen, Xianlong Lu, Linlin Ou, and Xinyi Yu. 3dgsnav: Enhancing vision-language model reasoning for object navigation via active 3d gaussian splatting, 2026. URL https://arxiv.org/abs/2602.12159.

Haotian Zhou, Xiaole Wang, He Li, Zhuo Qi, Jinrun Yin, Haiyu Kong, Jianghuan Xu, and Huijing Zhao. Lagmemo: Language 3d gaussian splatting memory for multi-modal open-vocabulary multi-goal visual navigation. *arXiv preprint arXiv:2510.24118*, 2025.

Filippo Ziliotto, Tommaso Campari, Luciano Serafini, and Lamberto Ballan. TANGO: Training-free embodied AI agents for open-world tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24603–24613, 2025.