

DualMap: Online Open-Vocabulary Semantic Mapping for Natural Language Navigation in Dynamic Changing Scenes

Jiajun Jiang¹, Yiming Zhu¹, Zirui Wu¹, and Jie Song^{1,2}

Abstract—We introduce DualMap, an online open-vocabulary mapping system that enables robots to understand and navigate dynamically changing environments through natural language queries. Designed for efficient semantic mapping and adaptability to changing environments, DualMap meets the essential requirements for real-world robot navigation applications. Our proposed hybrid segmentation frontend and object-level status check eliminate the costly 3D object merging required by prior methods, enabling efficient online scene mapping. The dual-map representation combines a global *abstract* map for high-level candidate selection with a local *concrete* map for precise goal-reaching, effectively managing and updating dynamic changes in the environment. Through extensive experiments in both simulation and real-world scenarios, we demonstrate state-of-the-art performance in 3D open-vocabulary segmentation, efficient scene mapping, and online language-guided navigation. Project page: <https://eku127.github.io/DualMap/>

I. INTRODUCTION

Imagine a robot navigating a bustling office when a colleague suddenly requests, “Bring me the MacBook I was using earlier.” This seemingly simple task requires a mapping system that bridges human language and dynamic environments. To succeed, three critical capabilities emerge: 1) **Open-vocabulary**: The system must detect objects in a class-agnostic manner during mapping and interpret various natural language queries (e.g., ‘the MacBook’) without relying on a set of fixed categories; 2) **Efficient online mapping**: The system must incrementally construct and maintain semantic maps in real-time, allowing it to adapt to ongoing changes in the environment; 3) **Navigation with dynamic changes**: In human-centered environments, objects frequently change their location or state—for instance, a laptop might be moved or unfolded. The system should implement a navigation strategy that guides the agent to the relocated object and updates the map for lifelong navigation.

Manuscript received: May, 26, 2025; Revised August, 12, 2025; Accepted October, 4, 2025.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported in part by Guangdong Basic and Applied Basic Research Foundation under Grant 2023A1515111151, in part by Guangzhou Municipal District Science and Technology Bureau under Grant 2025A03J3655, and in part by a research under Grant 2023QN10X121 (Corresponding author: Jie Song.)

¹Thrust of Robotics and Autonomous Systems, The Hong Kong University of Science and Technology (Guangzhou), {jjjiang127, yzhu714, zwu797}@connect.hkust-gz.edu.cn, jsongroas@hkust-gz.edu.cn

²Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE



Fig. 1: DualMap enables online open-vocabulary navigation under dynamic object changes. Given a natural language query, the robot navigates toward the desk guided by the *abstract map* (①), only to discover that the queried object is missing (②). The system updates the queried object’s status and new observations (e.g., cracker box on the table) via the online-built *concrete map* during navigation. This continuously updated map facilitates efficient re-selection of candidates to locate the relocated object (③).

To support such tasks, existing works [1], [2] construct semantic maps online from RGBD input using closed-set detectors. While efficient, these methods are constrained to predefined categories and cannot handle open-ended queries. Recent efforts [3]–[5] address this limitation by integrating vision foundation models [6]–[8]. However, these methods assume static environments. Other open-vocabulary systems [9], [10] tackle dynamic changes but typically require significant time—often several hours—to construct a usable map, making them impractical for lifelong robotic navigation. Consequently, no existing system simultaneously satisfies open-vocabulary understanding, efficient online mapping, and dynamic navigation, which restricts their applicability in real-world robotic applications.

In this work, we present DualMap, an online open-vocabulary map representation that is the first to meet all three aforementioned requirements. To enable efficient open-vocabulary mapping, the system constructs a high-quality 3D semantic *concrete map* through two key designs. A hybrid frontend process facilitates fast, holistic, and open-vocabulary

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE I: Comparison of recent semantic mapping systems.

System	Open-Vocab	Online Mapping	Dynamic Handling
Hydra [1]		✓	
Khronos [2]		✓	✓
ConceptGraphs [4]	✓		
HOV-SG [5]	✓		
CLIO [25]	✓	✓	
OpenIN [9]	✓		✓
DualMap (Ours)	✓	✓	✓

object segmentation. Intra-object self-status checks, including stability checks and object splitting, help remove noisy segments and enhance map fidelity. These innovations also eliminate the need for the costly 3D inter-object merging commonly used in prior works [4], [5].

To support navigation in dynamic environments, we convert the fine-grained *concrete map* into an *abstract map* composed of typically static objects (e.g., beds, desks) and layout information. This abstraction is based on the insight that maintaining precise details of all objects globally is unnecessary for navigation: global structural cues are sufficient for high-level planning, while object details can be retrieved locally through online perception. We then propose a dual-map navigation strategy, where the agent uses the global abstract map for candidate selection and the local concrete map for accurate object localization. The abstract map is continuously updated during navigation, facilitating efficient re-selection when the queried object has moved. Experimental results on Replica [11], ScanNet [12], and HM3D [13] demonstrate that DualMap achieves strong performance in semantic mapping, efficiency, and object navigation across both static and dynamic settings. Real-world experiments in indoor and outdoor scenes confirm its practical applicability. The code will be made available to support future research.

II. RELATED WORK

A. Open-Vocabulary Online Semantic Mapping

Fast closed-set detectors like YOLO [14] have been widely used in semantic robotic mapping [1], [15]–[17], but their fixed category sets limit generalization to open-world settings. To enable open-vocabulary perception, two main approaches have emerged. One uses image tagging models [18], [19] to generate class labels from incoming frames, which then guide closed-set detectors [20] for object detection. ConceptGraphs [4] follows this pipeline but remains offline due to the high cost of running GroundedSAM [21] per frame. A more common alternative applies class-agnostic segmentation [7], [8] directly, followed by feature embedding using vision foundation models (VFMs) [22]–[24]. HOV-SG [5] uses this method and merges over-segmented objects at the 3D level, improving semantic accuracy but remaining expensive in computation. CLIO [25] addresses efficiency by merging only task-relevant objects. However, since the map is task-specific, querying non-task objects requires extra map reconstruction, limiting map reusability. Different from prior approaches, DualMap adopts a hybrid segmentation strategy for efficient, holistic open-vocabulary mapping and replaces costly 3D inter-object merging with lightweight intra-object checks, enabling online operation.

B. Mapping and Navigation in Dynamic Environments

Previous 3D semantic mapping methods [4], [5], [26] support closed-set or open-vocabulary object navigation, but assume static environments. While being effective in static settings, they struggle to handle dynamic object movements common in real-world scenarios. To address this challenge, recent works [2], [9], [27], [28] have explored ways to adapt semantic maps to dynamic scene changes. OpenIN [9] proposes an open-vocabulary carrier-relation scene graph that updates memory when objects are relocated, but its offline pipeline limits update speed and hinders long-term operation. Khronos [2] introduces a real-time semantic mapping system for dynamic environments, maintaining a consistent dense mesh. However, it relies on closed-set segmentation and lacks support for arbitrary language queries, and also lacks navigation strategies for locating moved queried objects. The proposed DualMap system addresses the dynamic navigation challenge with a dual-map design. The system updates memory online during operation, enabling efficient navigation candidate re-selection when the queried object moves, and supports lifelong language-guided navigation.

III. ONLINE CONCRETE MAP

In this section, we detail the construction of the online concrete map \mathcal{M}_c (Fig. 2-a). To begin with, we describe how to efficiently generate object-level observations from the RGBD stream using fast, holistic open-vocabulary segmentation. Then, we explain how these observations incrementally update the map to ensure consistent and accurate scene representation.

A. Observation Generation

1) *Hybrid Open-Vocabulary Segmentation*: To efficiently detect and segment all potential objects in the scene, we adopt a hybrid open-vocabulary segmentation pipeline. As in Fig. 3-a, we employ YOLO [20] to rapidly produce object-level detections, and further generate the corresponding segmentation masks using MobileSAM [29]. Since YOLO requires a predefined category list, a task-specific detection list is generated once at the beginning of operation by prompting a large language model [30] based on the robot’s working context (e.g., “Give me a list of common objects when the robot operates in an office.”). We further refine YOLO outputs by merging overlapping bounding boxes with similar color distributions, measured via histogram similarity in RGB space. In parallel (Fig. 3-b), we run FastSAM [8], an open-set segmentation model, to capture objects beyond YOLO’s categories. We prioritize YOLO detections during fusion, retaining all YOLO outputs and supplementing them with non-overlapping FastSAM segments (\oplus in Fig. 3). This hybrid strategy ensures holistic and efficient object-level segmentation suitable for downstream mapping.

2) *Semantic Feature Embedding*: We generate a semantic feature by combining visual and textual information. For each detected object, we crop the image region defined by its 2D bounding box and encode it using CLIP [22] to obtain the image feature $\mathbf{f}_{\text{image}}$. If a class label is available (e.g., ‘chair’),

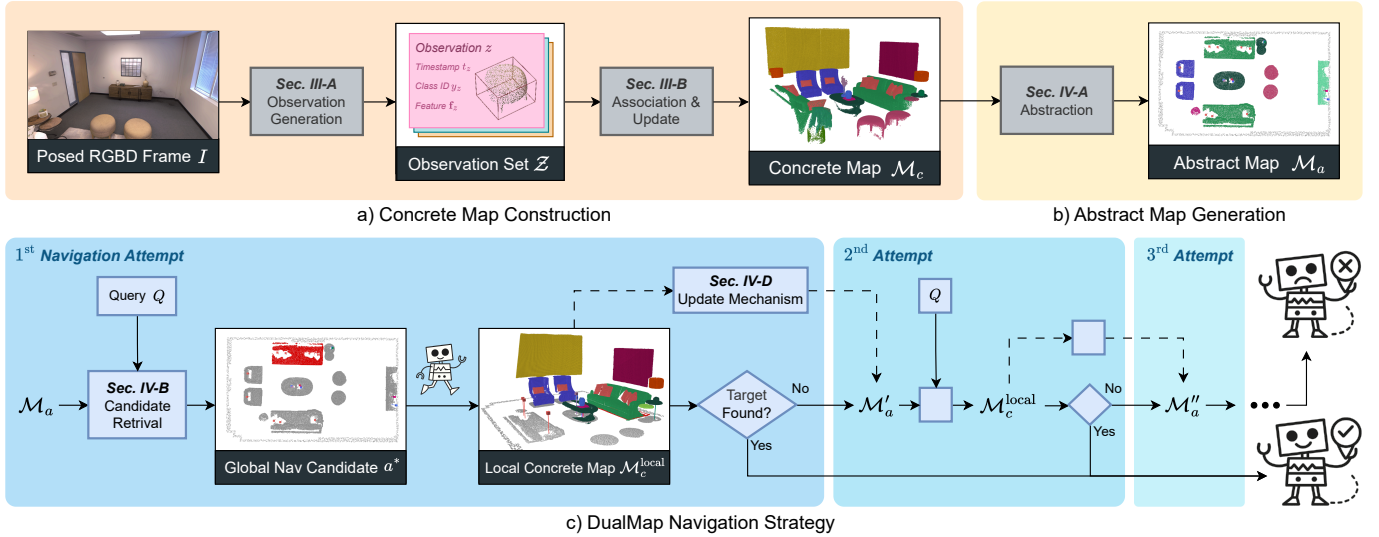


Fig. 2: DualMap system overview: a) A detailed 3D semantic concrete map \mathcal{M}_c is built from online observations of posed RGBD frames; b) An anchor-based abstract map \mathcal{M}_a is derived from \mathcal{M}_c , retaining global layout and static objects; c) Given a natural language query Q , the agent retrieves a global candidate a^* from \mathcal{M}_a and starts navigation. During execution, it incrementally builds a local concrete map $\mathcal{M}_c^{\text{local}}$, checks for target object presence, and updates the abstract map \mathcal{M}_a accordingly. If the target is not found near the a^* , a new navigation attempt is made using the updated map \mathcal{M}'_a . This loop continues until the target is found or the attempt limit is reached.

it is also encoded into a CLIP text feature \mathbf{f}_{text} . The final segment feature \mathbf{f} is computed as a weighted sum:

$$\mathbf{f} = w_{\text{image}}\mathbf{f}_{\text{image}} + w_{\text{text}}\mathbf{f}_{\text{text}}, \quad \mathbf{f}_{\{\text{image}, \text{text}\}} \in \mathbb{R}^d \quad (1)$$

where $w_{\text{image}} = 0.7$ and $w_{\text{text}} = 0.3$. For objects labeled as “null” by FastSAM, we replace \mathbf{f}_{text} with the normalized mean text feature computed over all known class embeddings to eliminate semantic bias.

3) *Observation Structure*: After processing current frame I , we generate an observation set \mathcal{Z} , defined as $\mathcal{Z} = \{z_1, z_2, \dots, z_N\}$, where N denotes the number of segments. Each observation $z = (\mathbf{P}_z, \mathbf{f}_z, y_z, t_z) \in \mathcal{Z}$ is structured with the following components: The 3D point cloud \mathbf{P}_z is generated by projecting the segmented region from the depth map into the world coordinate system using the camera intrinsics and the given pose. The semantic feature $\mathbf{f}_z \in \mathbb{R}^d$ is obtained through the semantic feature embedding process described previously. The class ID y_z specifies the object category detected by YOLO, or is set to “null” for segments detected by FastSAM. The timestamp t_z records when the observation is created. This structured observation set serves as the input for the subsequent object-level map updating process. In the following section, we describe how the system incrementally organizes and maintains objects in the concrete map based on these observations.

4) *Scene Pointcloud Generation*: To support layout estimation for downstream navigation, we maintain a scene point cloud \mathbf{P}_s that captures the overall geometry information of the environment during operation (Fig. 3-c). It is built by selecting RGBD frames with significant pose changes (translation $> 1.0\text{m}$ or rotation $> 20^\circ$), projecting their point clouds, and applying downsampling. This process (blue branch in Fig. 3) runs in a parallel low-frequency thread, providing a compact structural representation of the scene.

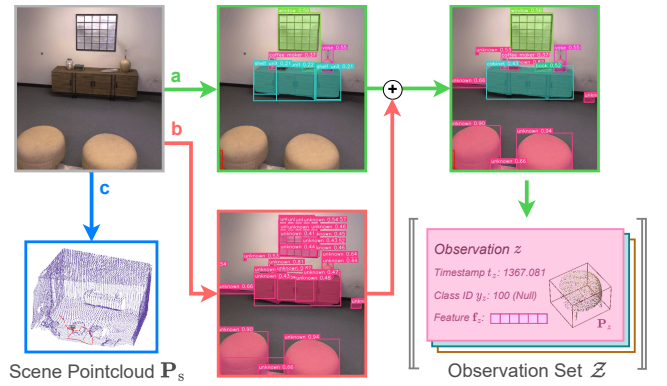


Fig. 3: Observation generation pipeline. From each posed RGBD frame, the system runs three parallel threads: a) *Closed-set detection* (e.g. YOLO) identifies in-list objects in object-level but misses unseen ones (e.g., *stools*); b) *Open-set segmentation* (e.g. FastSAM) captures a broader object range but may over-segment (e.g., *window*). After refinement and merging, the results are embedded into the final observation set. c) *Scene cloud generation* runs at low frequency and generates 3D point clouds for layout calculation.

B. Object Association and Update

1) *Map Initialization*: The concrete map \mathcal{M}_c represents a collection of objects, defined as $\mathcal{M}_c = \{o_1, o_2, \dots\}$. Each object o is structured as $(\mathbf{P}_o, y_o, \mathbf{f}_o, L_o)$, where \mathbf{P}_o , y_o , and \mathbf{f}_o inherit the same definitions as in the observation structure, and L_o records the list of associated observations. The initialization of the concrete map occurs when the first observation set \mathcal{Z}_0 is received. For each observation $z \in \mathcal{Z}_0$, a new object o is created in the map as follows:

$$(\mathbf{P}_o, y_o, \mathbf{f}_o, L_o) = (\mathbf{P}_z, y_z, \mathbf{f}_z, \{z\}). \quad (2)$$

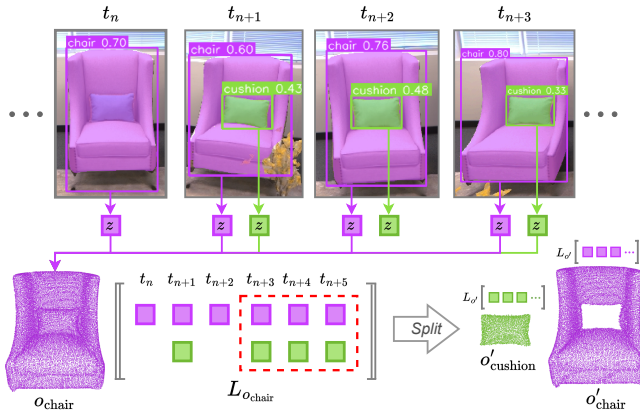


Fig. 4: Split detection. At t_n , under-segmentation merges the cushion into the chair’s mask, producing an erroneous observation. Updating o_{chair} with this observation corrupts its pointcloud, causing frequent misassignment of later cushion observations. The system detects such errors by checking $L_{o_{\text{chair}}}$ for co-occurring observations with different class IDs at the same timestamps. If such conflicts persist across frames (red box), a split is triggered to create new objects.

The initialized concrete map \mathcal{M}_c sets up the basis for further incremental updates.

2) *Observation Matching*: When a new observation set $\mathcal{Z} = \{z_1, z_2, \dots, z_N\}$ arrives, we match the observations to the existing objects $\mathcal{M}_c = \{o_1, o_2, \dots, o_M\}$ in the concrete map. The matching relies on both the geometric representation \mathbf{P} and the semantic feature \mathbf{f} of each observation and object. We define a similarity matrix $S \in \mathbb{R}^{N \times M}$, where each element $S(z_i, o_j)$ measures the similarity between observation z_i and object o_j as follows:

$$S(z_i, o_j) = \cos(\mathbf{f}_{z_i}, \mathbf{f}_{o_j}) + \text{Overlap}(\mathbf{P}_{z_i}, \mathbf{P}_{o_j}), \quad (3)$$

where $\cos(\cdot, \cdot)$ computes the cosine similarity between semantic features, and $\text{Overlap}(\cdot, \cdot)$ measures the 3D point cloud overlap ratio, following the calculation in [4].

For each observation z_i , we identify the best matching object based on the maximum similarity score:

$$o^* = \arg \max_{o_j} S(z_i, o_j). \quad (4)$$

If the score $S(z_i, o^*)$ exceeds a threshold τ , we associate z_i with o^* and update the object. We update o^* by averaging \mathbf{f}_{o^*} with \mathbf{f}_{z_i} , weighted by the number of previous observations (i.e., the size of the observation list L_{o^*}), augmenting \mathbf{P}_{o^*} with \mathbf{P}_{z_i} , and appending z_i to L_{o^*} . Otherwise, a new object is created following the initialization procedure.

3) *Object Status Check*: Our hybrid segmentation frontend produces holistic object observations. This avoids the need for expensive 3D inter-object merging, which typically combines fragmented segments into complete objects at high cost. Instead, after a map update with a new observation set \mathcal{Z} , we perform lightweight intra-object status checks—**stability check** and **split detection**—to maintain map fidelity. Both checks operate by analyzing the distribution of class IDs within an object’s observation list L_o .

Stability check filters out insufficiently observed or noisy objects. An object triggers a stability check if it has not been updated for an extended period. It passes the check if it has accumulated a number of observations over the threshold and if its most frequent class ID accounts for at least two-thirds of the total observations in its observation list. Otherwise, it is considered unstable and removed from the map.

Although we have refined object-level segmentation results, under-segmentation can still occur, leading subsequent observations from adjacent objects to be incorrectly associated with a single object, as shown in Fig. 4. When observations with different class IDs appear at the same timestamp across consecutive frames, we trigger a split operation, partitioning the object’s observation list by class IDs and creating new objects. This process preserves small objects and improves scene fidelity.

IV. ABSTRACT MAP AND NAVIGATION

This section details how the abstract map \mathcal{M}_a is constructed from the concrete map \mathcal{M}_c (Fig. 2-b), and how this dual-map design supports language-guided navigation under dynamic object changes (Fig. 2-c). We first describe the abstraction process, then explain how to retrieve navigation candidates from the abstract map given a language query. Next, we present a navigation strategy for static and dynamic environments. Finally, we show how the abstract map is updated online with a concrete map built during navigation.

A. Map Abstraction

The abstraction of the concrete map involves two steps: first, identifying which objects are selected as anchor objects; and second, associating the remaining volatile objects with anchor objects based on spatial relations. Anchor objects refer to static items that rarely move, such as beds or desks, while volatile objects (e.g., books or cups) frequently change positions. Only anchor objects are retained in the abstract map. We first classify objects using two lists of representative anchor and volatile categories. For each object $o \in \mathcal{M}_c$, we compute its similarity to both lists using CLIP features. If the maximum similarity to one list exceeds the other by a margin of $\Delta\tau = 0.05$, the object is assigned to the corresponding category. For cases without a clear margin, we further compute the similarity between the object’s feature \mathbf{f}_o and the anchor template feature \mathbf{f}_t , encoded from descriptive phrases (e.g., “furniture that is not often moved”). If the similarity exceeds a threshold τ_a , the object is classified as an anchor object a ; otherwise, it is a volatile v . The threshold τ_a is set to 0.5 based on classification accuracy experiment.

During abstraction, all non-semantic attributes of volatile objects (e.g., \mathbf{P}_v , L_v) are discarded. We retain only the features of volatile objects that are spatially related to anchor objects. While our framework supports multiple relation types, we focus on the “on” relation in this work. To determine whether a volatile object is on an anchor, we first extract the anchor’s supporting plane by analyzing the Z-axis histogram of its point cloud. We then compute the 2D overlap ratio in the XY-plane to ensure the object

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

lies within the anchor’s projected footprint, and verify that the vertical distance from the object’s bottom to the supporting plane is within a threshold δ (we set $\delta = 0.1\text{m}$ in practice). Once a spatial relation is established, the volatile object’s semantic feature \mathbf{f}_v is stored in the associated feature list of the corresponding anchor object, denoted as $L_a = \{\mathbf{f}_{v_1}, \mathbf{f}_{v_2}, \dots, \mathbf{f}_{v_K}\}$. Each anchor object $a \in \mathcal{M}_a$ is thus represented as $a = (\mathbf{P}_a, \mathbf{f}_a, y_a, L_a)$ ¹. After abstraction, the abstract map \mathcal{M}_a consists solely of anchor objects, and is represented as $\mathcal{M}_a = \{a_1, a_2, \dots\}$.

Scene layout, important for navigation, is also extracted in map abstraction. We project the scene point cloud \mathbf{P}_s onto a bird’s-eye-view plane and discretize it into spatial bins. A binary occupancy map is then generated by computing a histogram of point counts across bins and using the 90th percentile as a threshold, highlighting scene structures like walls. The layout is visualized as gray point clouds in Fig. 1, and is updated only when a new navigation query is issued.

B. Candidate Retrieval

Given a natural language query Q , our goal is to retrieve the most relevant candidate from the abstract map \mathcal{M}_a to serve as a navigation target. We first encode the query using CLIP to obtain a semantic embedding, denoted as \mathbf{f}_q . We then compare the query feature \mathbf{f}_q with both the anchor’s own feature \mathbf{f}_a and the features of its associated volatile objects $L_a = \{\mathbf{f}_{v_1}, \dots, \mathbf{f}_{v_K}\}$. We compute the anchor score as:

$$s(a) = \max\left(\cos(\mathbf{f}_q, \mathbf{f}_a), \max_i \cos(\mathbf{f}_q, \mathbf{f}_{v_i})\right). \quad (5)$$

We compute scores for all anchors in \mathcal{M}_a and select the one with the highest score as the global navigation candidate:

$$a^* = \arg \max_{a \in \mathcal{M}_a} s(a). \quad (6)$$

The selected anchor a^* suggests that the queried object is most likely situated nearby. Both the anchor a^* and its similarity score $s(a^*)$ are used for further navigation.

C. Navigation Strategy

A global path toward a^* is planned using a Voronoi-based planner over the abstract map [31]. As the agent moves along this path, it incrementally builds a local concrete map $\mathcal{M}_c^{\text{local}}$ and evaluates objects in $\mathcal{M}_c^{\text{local}}$ for potential matches. For each object $o \in \mathcal{M}_c^{\text{local}}$, we compute the cosine similarity $s(o)$ between its feature \mathbf{f}_o and the query feature \mathbf{f}_q . If $s(o)$ is within a margin ϵ of $s(a^*)$ and o lies within the projected footprint of a^* , it is considered a confident match. A local path is then planned via RRT* [32] to reach the target.

If no confident match is found near a^* , suggesting the queried object may have changed location, the system re-executes the candidate retrieval over the updated abstract map $\mathcal{M}'_a \setminus \{a^*\}$ and selects a new anchor a^* . Here, \mathcal{M}'_a is obtained by updating the original abstract map \mathcal{M}_a with the local concrete map $\mathcal{M}_c^{\text{local}}$, as will be detailed in the next

section. The agent then resumes navigation toward a^* , using the original similarity score $s(a^*)$ to remain consistent with the initial query. This strategy enables the agent to leverage contextual cues encountered in the earlier navigation process, increasing the likelihood of success if the target object was partially observed along the way.

D. Map Update Mechanism

To update the abstract map \mathcal{M}_a , we extract stable objects from $\mathcal{M}_c^{\text{local}}$ and apply the same abstraction process to obtain new anchors a^{new} . Each a^{new} is compared against all existing anchors $a \in \mathcal{M}_a$ by computing the geometric overlap score $s_{\text{overlap}} = \text{Overlap}(\mathbf{P}_{a^{\text{new}}}, \mathbf{P}_a)$. If $s_{\text{overlap}} > \tau_1$, we treat a^{new} as an additional observation of a and update its representation. The point cloud \mathbf{P}_a is augmented with $\mathbf{P}_{a^{\text{new}}}$ and downsampled, and the semantic feature is updated using a size-weighted average:

$$\mathbf{f}_a \leftarrow \frac{|\mathbf{P}_a| \cdot \mathbf{f}_a + |\mathbf{P}_{a^{\text{new}}}| \cdot \mathbf{f}_{a^{\text{new}}}}{|\mathbf{P}_a| + |\mathbf{P}_{a^{\text{new}}}|}. \quad (7)$$

The symbol $|\mathbf{P}|$ refers to the size of the pointcloud \mathbf{P} . If the overlap score further exceeds a stricter threshold τ_2 , indicating that a^{new} provides a more complete observation of the existing anchor, we replace the volatile feature list L_a with $L_{a^{\text{new}}}$. If no anchor satisfies $s_{\text{overlap}} > \tau_1$, we insert a^{new} into the abstract map as a new anchor. The updated abstract map \mathcal{M}'_a maintains the up-to-date scene representation throughout navigation.

V. EXPERIMENTS

A. Benchmarks and Implementation Details

1) *Datasets*: We evaluate DualMap on semantic segmentation, efficiency, and object navigation. For semantic segmentation and efficiency, we use the ScanNet [12] and Replica [11] datasets, following scene selections from prior work [4], [5], with all experiments conducted at the original image resolutions. For static object navigation, we evaluate on the HM3D dataset [13], selecting three representative scenes (00829, 00848, 00880) and augmenting object diversity by integrating YCB objects [33]. For dynamic navigation, as current benchmarks lack object-level dynamics, we construct new environments for dynamic object navigation using a custom tool built on Habitat Simulator [34], which enables dynamic placement of YCB objects. We define two types of object changes: *In-anchor relocation*, where an object moves within the same anchor (e.g., moving a cup on a table); *Cross-anchor relocation*, where an object is moved to a different anchor (e.g., from a table to a shelf). For each static HM3D scene, we create dynamic variants by manually repositioning added objects to reflect both types of changes. Agents are deployed in the Habitat Simulator for both static and dynamic navigation experiments. For real-world experiments, we design a rigid perception module with a LiDAR and an RGBD camera, where LiDAR provides pose estimation via FastLIO2 [35]. This module is deployed on both wheeled and quadruped platforms, as shown in Fig. 6-a.

¹To support the navigation task, we simplify anchor object geometry \mathbf{P}_a by preserving only its 2D projected point cloud.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 5: Qualitative comparison of semantic segmentation results on ScanNet scene0011_00. Red arrows highlight semantically inaccurate predictions, while blue arrows indicate meaningless segmentations, suggesting noisy predictions.

TABLE II: Open-vocabulary 3D Semantic Segmentation and Efficiency

Dataset	Method	mIoU \uparrow	FmIoU \uparrow	mAcc \uparrow	ODR ≈ 1	Avg. Mem (MB) \downarrow	Peak Mem (MB) \downarrow	TPF (s) \downarrow
Replica	ConceptGraphs	0.1501	0.3858	0.3559	2.02	7148.9	23551.9	4.188
	HOV-SG	0.2050	0.4846	0.3835	3.81	73368.0	158126.6	42.005
	Ours	0.2538	0.5207	0.4024	0.97	3095.2	4564.0	0.276
ScanNet	ConceptGraphs	0.0882	0.3077	0.3538	6.97	9780.3	26155.2	6.301
	HOV-SG	0.1333	0.3381	0.3714	20.34	9223.0	25735.0	8.039
	Ours	0.1604	0.3288	0.3794	2.56	2120.9	2820.2	0.163

2) *Evaluation Metrics*: For semantic segmentation, we adopt standard metrics from [5], including mean IoU (mIoU), frequency-weighted IoU (F-mIoU), and mean accuracy (mAcc). In both qualitative and quantitative evaluations, we exclude background classes like ["wall", "floor", "ceiling"], which occupy large portions of the scene but provide little semantic information. We also report the Object Density Ratio (ODR), calculated as $ODR = N_{pred}/N_{GT}$, where a value closer to 1 indicates that the predicted object count (N_{pred}) is more aligned with the ground-truth count (N_{GT}). For efficiency, we measure average memory usage, peak memory consumption, and time per frame (TPF) during operation. For navigation, we report Success Rate (SR), defined as the percentage of queries in which the agent stops within 1 meter of the queried object. In dynamic scenes, success further requires finding the target within three attempts. Unless otherwise specified, all reported metric values in tables are averaged across scenes for each dataset.

3) *Compared Methods*: We compare against two open-vocabulary semantic mapping systems: ConceptGraphs [4], which supports online object-level mapping, and HOV-SG [5], known for strong semantic performance and built-in navigation ability. We extend both baselines in Habitat Simulator by enabling agents to navigate using their query results and the same planning algorithm as ours.

4) *Implementation Details*: All experiments are conducted on an NVIDIA RTX 4090 GPU and an Intel i7-12700KF CPU. We adopt the YOLOv8l-world model [20] for closed-set detection, using a class list generated by GPT-4 [30], and supplement it with FastSAM-s [8] to enable open-vocabulary segmentation. For feature embedding, we use MobileCLIP-S2 [24], shared across all methods for fair comparison.

B. Concrete Mapping Evaluation

We present quantitative and qualitative results in Table II and Fig. 5. DualMap achieves the best semantic segmentation results across datasets, improving mIoU on Replica by 2.8% over HOV-SG and 10.3% over ConceptGraphs. This gain stems from object-level splitting that preserves the diversity of objects and text-enhanced feature embedding. Fig. 5 further illustrates these improvements, with red arrows showing misclassified objects that are better handled by DualMap. For ODR, DualMap achieves the closest alignment with ground truth (0.97 on Replica), thanks to holistic segmentation that reduces over-segmentation and our stability check that removes noisy objects. While HOV-SG and ConceptGraphs apply inter-object merging to control segment counts, they still suffer from redundant fragments. This difference is clearly illustrated in Fig. 5, where noisy RGBD inputs from real sensors result in meaningless predictions (blue arrows) in the baselines. DualMap filters out such errors via stability check, demonstrating robustness in real-world scenarios.

In efficiency, DualMap also outperforms other methods: it reduces peak memory usage by over **96%** and TPF by **99.3%** compared to HOV-SG on Replica. This is enabled by a hybrid open-vocabulary detection strategy and lightweight 2D refinement, which avoids costly 3D post-processing.

C. Navigation in Simulated Environments

For static navigation, we evaluate all methods across three HM3D scenes. As shown in Table III-Static, DualMap achieves the highest success rate. ConceptGraphs performs competitively due to its object-centric mapping. In contrast, HOV-SG's segment-merging strategy results in fragmented maps, leading to more query and navigation failures.

For dynamic scenes (Table III-Dynamic), DualMap naturally handles *in-anchor* changes, as the queried object

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 6: Real-world navigation in dynamic environments. a) Robotic platforms equipped with a perception module that integrates a LiDAR and an RGB-D camera, both mounted on a rigid 3D-printed mount. b–c) Two examples of language-guided navigation in dynamic real-world scenes, where the agent tries to locate relocated objects across multiple attempts.

TABLE III: Object Navigation Success Rate (SR) on HM3D.

Scene Type	Method	00829	00848	00880	Trials	Avg. SR
Static	ConceptGraphs	69.2%	53.8%	61.5%		61.5%
	HOV-SG	53.8%	46.2%	57.7%	78	52.6%
	Ours	73.1%	69.2%	69.2%		70.5%
Dynamic	In-anchor (Ours)	66.7%	66.7%	61.1%	54	64.8%
	Cross-anchor (Ours)	55.6%	61.1%	64.7%	53	60.3%

TABLE IV: Semantic Mapping Ablation on Replica

Options	FmIoU	mAcc	mIoU
w/o FastSAM	0.4753	0.3685	0.2344
w/o YOLO Refinement	0.5043	0.3886	0.2399
w/o Weighted Feats	0.4209	0.3348	0.1814
w/o Object Split Detection	0.5065	0.3998	0.2496
Full System	0.5207	0.4024	0.2538

remains within its original anchor and can be accurately localized through the local concrete map. For more challenging *cross-anchor* cases, it updates the abstract map online and reselects reasonable candidates accordingly. This strategy maintains high success rates even under large object movements. Among failure cases in *cross-anchor* setting, 28.3% are due to false matches, 9.4% to exceeding navigation attempt limits, and 1.9% to planning errors.

D. Ablation Study

We perform ablations to validate the impact of key components in DualMap, focusing on semantic mapping, efficiency, and navigation under dynamic scenes.

1) *Mapping Performance*: In Table IV and Fig. 5, we ablate key components of the concrete mapping pipeline on Replica and ScanNet. We first remove FastSAM and rely only on YOLO for detection. Since YOLO detects only predefined categories, many unseen objects are missed, leading to a drop in performance. Next, we disable the refinement step that merges overlapping YOLO detections based on color histogram similarity. Without merging overlaps, objects are over-segmented. The performance drops slightly, as YOLO produces relatively holistic detections. Then, we remove the weighted feature merging and use only image features. This causes a notable drop, indicating the importance of incorporating global semantic cues from text. Finally, we disable the object-level splitting mechanism. This causes

TABLE V: Mapping Efficiency Ablation on ScanNet. Gray row indicates configuration used in our proposed system.

Obj. Merging	Stability Check	Closed-set Det.	ODR	TPF(s)
✓	✗	✗	34.2	27.04
✓	✗	✓	26.5	2.841
✓	✓	✓	2.41	0.3056
✗	✓	✓	2.56	0.163

TABLE VI: Success Rates under Different Candidate Selection Strategies for Relocated Objects on HM3D

Strategy	Random Pick	Based on \mathcal{M}_a	Based on \mathcal{M}'_a
SR	13.2%	47.2%	60.3%

under-segmentation, where nearby objects are grouped into one, leading to inaccurate matching and reduced mapping precision. Qualitative results in Fig. 5 show that without object stability check, sensor noise introduces noisy segments, degrading overall semantic quality.

2) *Efficiency*: In Table V, we analyze three key modules that are related to our mapping pipeline to evaluate their impact on efficiency: Inter-object merging in 3D, intra-object level stability check, and closed-set detection assistance in open-vocabulary segmentation. HOV-SG follows the configuration in Row 1, resulting in the highest time cost and ODR. This is because SAM-style segmentation produces a large number of segments, making 3D-level merging computationally expensive. Configurations in Row 2 improve speed by using YOLO to assist segmentation, but still suffer from noisy objects caused by sensor noise. In Row 3, enabling our proposed stability check significantly improves ODR by removing noisy objects, which also reduces the object merging time cost. Our final configuration (Row 4) disables inter-object merging, slightly sacrificing map accuracy but achieving notable efficiency gains—an important trade-off for real-world robotic applications.

3) *Dynamic Scene Navigation*: To evaluate navigation under dynamic changes, we focus on the more challenging *cross-anchor* setting and test different candidate selection strategies when the queried object has changed position. As shown in Table VI, random selection often fails due to excessive retries. Using the original abstract map \mathcal{M}_a provides anchor-based guidance and improves success. Our final strategy updates \mathcal{M}_a during navigation to form \mathcal{M}'_a ,

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE VII: Real-World Object Navigation Results

Platform	Scene Type	Static		Dynamic	
		Trials	SR	Trials	SR
Wheeled	Meeting Room	14	85.7%	27	70.3%
	Apartment	46	69.6%	33	51.5%
Quadruped	Indoor Hallway	19	78.9%	27	55.6%
	Outdoor	12	75.0%	18	50.0%

which makes full use of past navigation information. This helps the agent find the relocated object more effectively.

E. Navigation in Real-World Environments

We evaluate our system on two robotic platforms across four diverse real-world scenes, as shown in Table VII and Fig. 6. To construct test environments, we place diverse objects (e.g., bottles, backpacks, keyboards) within each scene. During mapping, the robot is manually driven to explore and build an abstract map. We then perform static object navigation using the same pipeline as in simulation. To test dynamic scenarios, we manually rearrange the objects and initiate navigation queries. The results show that DualMap achieves performance levels similar to those in the simulations, confirming its robustness in real-world settings across scenes and platforms.

VI. CONCLUSION

We present DualMap, an online open-vocabulary semantic mapping system for language-guided navigation in dynamic environments. The system enables efficient and accurate mapping, and supports robust navigation through online updates and re-selection. While effective, our system has several limitations. First, DualMap does not perform camera pose estimation and instead relies on external localization systems. Second, the system does not yet incorporate human-object interactions, limiting its ability to reason about human activities in the scene. Third, even in simple outdoor scenes, system performance degrades compared to indoor settings due to higher sensor noise. We consider these limitations promising directions for future work.

REFERENCES

- [1] N. Hughes, Y. Chang, and L. Carlone, "Hydra: A real-time spatial perception system for 3D scene graph construction and optimization," in *Proc. Robotics: Science and Systems (RSS)*, 2022.
- [2] L. Schmid, M. Abate, Y. Chang, and L. Carlone, "Khronos: A Unified Approach for Spatio-Temporal Metric-Semantic SLAM in Dynamic Environments," in *Proc. Robotics: Science and Systems (RSS)*, 2024.
- [3] K. M. Jatavallabhula, et al., "Conceptfusion: Open-set multimodal 3d mapping," in *Proc. Robotics: Science and Systems (RSS)*, 2023.
- [4] Q. Gu, et al., "Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 5021–5028.
- [5] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard, "Hierarchical open-vocabulary 3d scene graphs for language-grounded robot navigation," in *Proc. Robotics: Science and Systems (RSS)*, 2024.
- [6] S. Liu, et al., "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *European Conference on Computer Vision (ECCV)*, 2024, pp. 38–55.
- [7] A. Kirillov, et al., "Segment anything," in *Proc. IEEE/CVF international conference on computer vision*. PMLR, 2023, pp. 4015–4026.
- [8] X. Zhao, et al., "Fast segment anything," 2023, arXiv:2306.12156.
- [9] Y. Tang, et al., "Openin: Open-vocabulary instance-oriented navigation in dynamic domestic environments," *IEEE Robotics and Automation Letters*, no. 99, pp. 1–8, 2025.
- [10] Z. Yan, et al., "Dynamic open-vocabulary 3d scene graphs for long-term language-guided mobile manipulation," *IEEE Robotics and Automation Letters*, vol. 10, no. 5, pp. 4252–4259, 2025.
- [11] J. Straub, et al., "The replica dataset: A digital replica of indoor spaces," 2019, arXiv:1906.05797.
- [12] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5828–5839.
- [13] K. Yadav, et al., "Habitat-matterport 3d semantics dataset," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 4927–4936.
- [14] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.
- [15] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "Semanticfusion: Dense 3d semantic mapping with convolutional neural networks," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4628–4635.
- [16] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger, "Fusion++: Volumetric object-level slam," in *International Conference on 3D Vision (3DV)*. IEEE, 2018, pp. 32–41.
- [17] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1689–1696.
- [18] Y. Zhang, et al., "Recognize anything: A strong image tagging model," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024, pp. 1724–1732.
- [19] X. Huang, et al., "Tag2text: Guiding vision-language model via image tagging," 2023, arXiv:2303.05657.
- [20] T. Cheng, L. Song, Y. Ge, W. Liu, X. Wang, and Y. Shan, "Yolo-world: Real-time open-vocabulary object detection," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 16901–16911.
- [21] T. Ren, et al., "Grounded sam: Assembling open-world models for diverse visual tasks," 2024, arXiv:2401.14159.
- [22] A. Radford, et al., "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [23] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer, "Sigmoid loss for language image pre-training," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023, pp. 11941–11952.
- [24] P. K. A. Vasu, H. Pouransari, F. Faghri, R. Venulapalli, and O. Tuzel, "Mobileclip: Fast image-text models through multi-modal reinforced training," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 15963–15974.
- [25] D. Maggio, et al., "Clio: Real-time task-driven open-set 3d scene graphs," *IEEE Robotics and Automation Letters*, vol. 9, no. 10, pp. 8921–8928, 2024.
- [26] J. Zhang, et al., "3d-aware object goal navigation via simultaneous exploration and identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 6672–6682.
- [27] L. Schmid, et al., "Panoptic multi-tsdFs: a flexible representation for online multi-resolution volumetric mapping and long-term dynamic scene consistency," in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8018–8024.
- [28] J. Fu, et al., "Planesdf-based change detection for long-term dense mapping," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9667–9674, 2022.
- [29] C. Zhang, et al., "Faster segment anything: Towards lightweight sam for mobile applications," 2023, arXiv:2306.14289.
- [30] J. Achiam, et al., "Gpt-4 technical report," 2023, arXiv:2303.08774.
- [31] S. Thrun and A. Bücken, "Integrating grid-based and topological maps for mobile robot navigation," in *Proceedings of the National Conference on Artificial Intelligence*, 1996, pp. 944–951.
- [32] K. Sertac and F. Emilio, "Incremental sampling-based algorithms for optimal motion planning," in *Proc. Robotics: Science and Systems (RSS)*, 2010.
- [33] B. Calli, et al., "Yale-cmu-berkeley dataset for robotic manipulation research," *The International Journal of Robotics Research*, vol. 36, no. 3, pp. 261–268, 2017.
- [34] X. Puig, et al., "Habitat 3.0: A co-habitat for humans, avatars and robots," 2023, arXiv:2310.13724.
- [35] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "Fast-lid2: Fast direct lidar-inertial odometry," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.