LagMemo: Language 3D Gaussian Splatting Memory for Multi-modal Open-vocabulary Multi-goal Visual Navigation

Haotian Zhou[1], Xiaole Wang[1], He Li[1], Jianghuan Xu[1], Zhuo Qi[1], Jinrun Yin[1], Haiyu Kong[2], Huijing Zhao[1]

Abstract—Navigating to a designated goal using visual information is a fundamental capability for intelligent robots. To address the practical demands of multi-modal, open-vocabulary goal queries and multi-goal visual navigation, we propose LagMemo, a navigation system that leverages a language 3D Gaussian Splatting memory. During a one-time exploration, LagMemo constructs a unified 3D language memory with robust spatial-semantic correlations. With incoming task goals, the system efficiently queries the memory, predicts candidate goal locations, and integrates a local perception-based verification mechanism to dynamically match and validate goals. For fair and rigorous evaluation, we curate GOAT-Core, a high-quality core split distilled from GOAT-Bench. Experimental results show that LagMemo's memory module enables effective multi-modal open-vocabulary localization, and significantly outperforms state-of-the-art methods in multi-goal visual navigation. Project page: https://weekgoodday.github.io/lagmemo

I. INTRODUCTION

In real-world applications such as home assistants and service robots, mobile agents are expected to understand user instructions, perceive environments, and navigate to target objects [1][2]. With the advancement of vision-language models, and inspired by the fact that humans primarily rely on vision to navigate, visual navigation has emerged as a prominent research area [1][3]. This task requires robots to rely primarily on visual sensors to reach target destinations in a safe and efficient manner. Recently, the embodied navigation community has increasingly focused on the more complex and practical setting of multi-modal open-vocabulary multi-goal visual navigation, as illustrated in Fig. 1. In such scenarios, robots are required to continuously complete multiple tasks within the same environment, with goals specified in various modalities and potentially involving novel categories. To advance this paradigm, GOAT-Bench [4] has provided baseline algorithms and a dataset based on the Habitat simulator [5].

However, existing methods remain limited. End-to-end approaches like RL GOAT [4] encode the environment implicitly through hidden states, leading to poor generalization. More importantly, while modular approaches like Modular GOAT construct an instance memory to support multi-modal queries, their memory population is fundamentally constrained by an upstream object detector with a predefined

[1]Haotian Zhou, Xiaole Wang, He Li, Jianghuan Xu, Zhuo Qi, Jinrun Yin, and Huijing Zhao are with the Key Laboratory of Machine Perceptions (MOE), School of Intelligence Science and Technology, Peking University, Beijing, 100084, China (e-mail: zhoutiantian19@pku.edu.cn; zhaohj@pku.edu.cn).

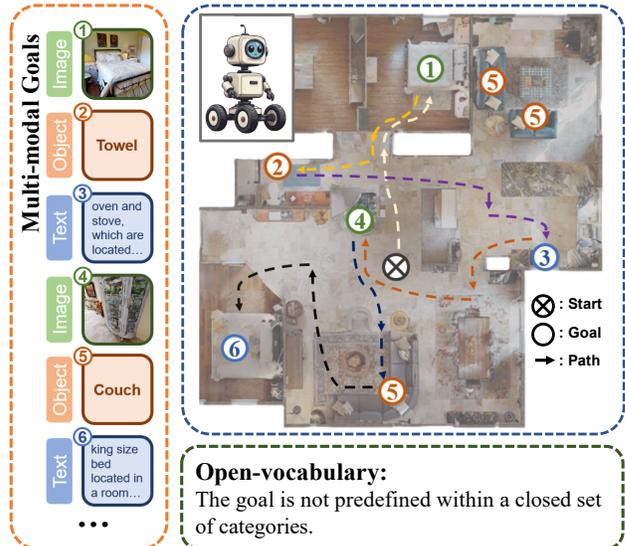[2]Haiyu Kong is with Beijing University of Posts and Telecommunications (BUPT), Beijing, 100876, China.

Fig. 1: Illustration of multi-modal open-vocabulary multi-goal visual navigation task. Multi-modal: the goal can be specified in the forms of an object category, an image or a text description; Open-vocabulary: the agent is not limited to navigating to a predefined closed set of categories; Multi-goal: the agent is required to find multiple goals within the same environment.

category list. From a memory perspective, this prevents open-vocabulary capability, as unforeseen novel targets are ignored during exploration and become irretrievable. For instance, an unforeseen open-set target like a "Mickey Mouse doll" would likely be ignored by the fixed-vocabulary detector during exploration, meaning it is never memorized and thus impossible to retrieve. Furthermore, projecting these features onto 2D semantic maps inevitably loses fine-grained 3D spatial details and lacks 3D spatial-semantic correlation, which prevents global semantic consistency optimization across multiple views. To overcome this upstream detection bottleneck and the limitations of 2D representations, we introduce LagMemo, a novel visual navigation system that integrates 3DGS with a quantized language feature space. During navigation, it supports multi-modal queries to identify candidate waypoints, improving navigation performance.

Our main contributions are summarized as follows:

- We propose LagMemo, a visual navigation system that introduces a unified 3D Gaussian Splatting memory module equipped with codebook-based language feature embeddings. To address the inherently sparse observations collected during rapid pre-exploration, a keyframe retrieval mechanism is incorporated.
- We propose a memory-guided visual navigation framework that incorporates a novel goal verification mechanism to bridge memory and real-time perception. This
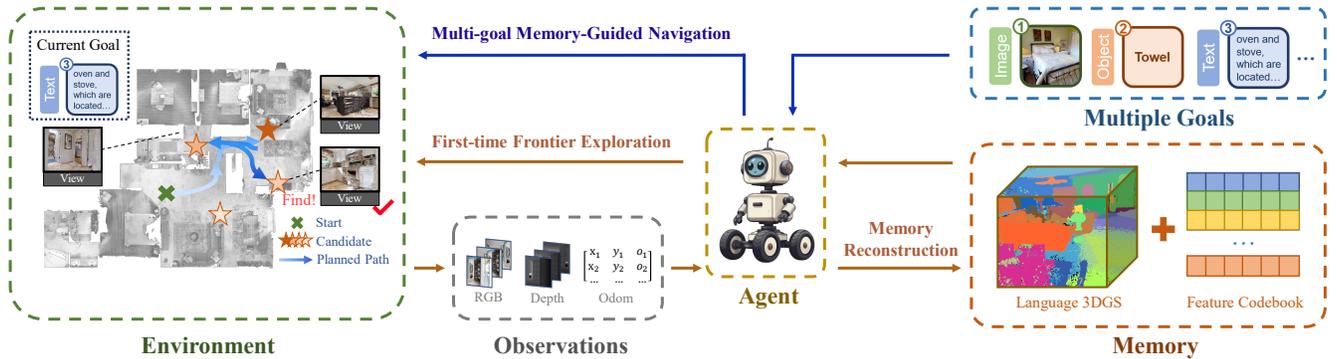
Fig. 2: LagMemo Overview. The agent first performs frontier-based exploration to collect observations from the environment, upon which it reconstructs a language 3DGS memory and a feature codebook. As multi-modal open-vocabulary goals input, the agent queries the memory to generate candidate localization regions and uses real-time perception to verify targets, thereby accomplishing multi-goal visual navigation.

mechanism operates through a cyclic process of memory query and perception-based validation, effectively leveraging the constructed memory to improve navigation performance.

- Based on GOAT-Bench, we curate a high-quality household environment split named GOAT-Core, which can be used to evaluate both goal localization and multi-modal multi-goal visual navigation tasks.
- Extensive quantitative and qualitative evaluations, alongside real-world deployments, demonstrate that LagMemo achieves superior performance in both goal localization and multi-modal open-vocabulary multi-goal visual navigation.

II. RELATED WORKS

A. Visual Navigation

In terms of methodology, visual navigation approaches can be broadly categorized into modular [6] and end-to-end [7] frameworks, with recent advances in zero-shot navigation methods [8][9]. In terms of goal modality, visual navigation tasks are commonly classified into ObjectGoal [10], ImageGoal [11], and TextGoal [12]. A critical challenge in these settings is whether the target is limited to a closed set of categories. For example, SemExp [6] uses C channels in 2D semantic map to record the location of C predefined goal categories. Recent works [9][13] exploit CLIP or BLIP-2 to enable open-vocabulary identification and localization for novel goals. Moving beyond single-target scenarios, this paper focuses on multi-modal open-vocabulary multi-goal visual navigation, a task formalized as lifelong visual navigation in GOAT [4], which is more aligned with practical applications.

B. Memory in Visual Navigation

Effective memory is critical for long-horizon navigation. While some approaches encode implicit memory using RNNs [7], explicit representations are preferred for complex environments. A common paradigm is to use 2D semantic grid maps [6][9], which discretize the environment into spatial grids, assigning each grid a semantic label or high-dimensional embedding. However, such 2D projections inevitably lose vertical spatial details and are highly susceptible to multi-view feature inconsistencies. Some works [14][15]

construct scene graphs to model spatial relationships between objects and rooms. However, such paradigm often over-abstracts the environment by aggressively compressing dense visual information into sparse node embeddings. Recently, 3DGS has emerged as a promising dense representation (e.g., GaussNav [16], IGL-Nav [17]). Yet, these methods tailored to instance image-goal navigation, rely heavily on high-fidelity RGB rendering for visual matching. Such dense multi-view observations are typically not available during rapid robotic exploration. In contrast, LagMemo utilizes a codebook-quantized language 3DGS memory, preserving 3D spatial-semantic correlations while enabling efficient retrieval directly within the feature space. This design not only enables multi-modal open-vocabulary querying but also ensures robust localization even when geometric reconstruction is imperfect due to sparse exploration views.

C. 3D Gaussian Splatting with Language Embedding

3D Gaussian Splatting (3DGS) [18] is a 3D scene reconstruction method, which has attracted significant attention due to its high-quality and real-time rendering capabilities [19]. Beyond online geometric modeling [20][21], as many embodied tasks require not only geometric representations but also scene understanding, recent works [22][23] attempt to incorporate language features by embedding vision-language features into Gaussians for scene understanding. For instance, LangSplat [24] compresses CLIP features via scene-specific autoencoders. Online Language Splatting [25] enables near real-time incremental language mapping. However, while achieving high-fidelity 2D semantic rendering, these approaches exhibit limited capabilities in explicit 3D spatial indexing as pointed out in [26]. LagMemo employs codebook clustering for robust 3D spatial-semantic association, and specifically tailored to the sparse-view navigation setting with efficient feature retrieval.

III. METHODOLOGY

A. Task and System Overview

We address the multi-modal, open-vocabulary, multi-goal visual navigation task in a multi-room indoor household environment. In a practical deployment scenario, an agent is expected to execute a continuous sequence of tasks within a given environment where the goal of the k^{th} task is denoted
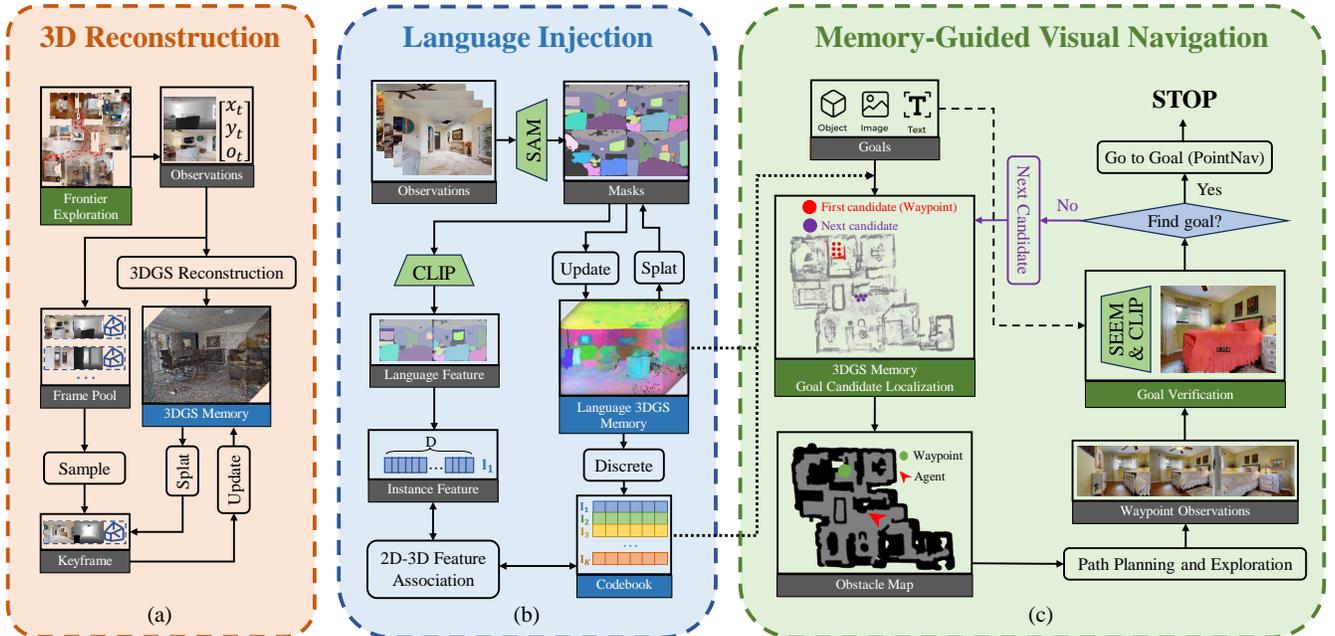
Fig. 3: Language 3DGS Memory Reconstruction and Memory-Guided Visual Navigation Pipeline. (a) 3D Reconstruction. During frontier exploration, the agent collects RGB, depth, and odometry to reconstruct 3DGS memory. A keyframe retrieval mechanism is employed to mitigate the forgetting and surface holes caused by sparse navigation views. (b) Language Injection. For image observations, we leverage SAM and CLIP to extract 2D semantic features. Via 2D-3D association, these features are assigned to Gaussians and discretized into a codebook. (c) Memory-Guided Visual Navigation. During execution, multi-modal open-vocabulary goals query the memory to propose candidate locations (waypoints). Using the obstacle map for path planning, the agent verifies the target to decide success or move to the next candidate.

as g^k. At every timestep t during the task execution, the agent receives observations including an RGB image I_t^k, a depth image D_t^k, and odometry $p_t^k = (x_t^k, y_t^k, o_t^k)$, and outputs an action A_t^k. Upon finishing the current task, a new target goal g^{k+1} is sequentially provided. The navigation goals are specified in one of three modalities: a category, an image, or a text description. Crucially, the task is open-vocabulary: the agent is not limited to navigating to objects from a predefined closed set of categories.

As illustrated in Fig.2, we propose LagMemo. Upon entering a novel environment, the agent conducts a one-time frontier-based exploration to scan the surroundings. During this phase, a language 3D Gaussian Splatting (3DGS) memory is constructed (Sec.III-B). Once the memory is built, it serves as a persistent prior to efficiently support all subsequent multi-goal tasks within that environment. For incoming multi-modal goals, the agent queries the language 3DGS memory, and then navigates to the candidate instances and verifies whether the observed objects match the goal (Sec.III-C).

B. Language 3DGS Memory Reconstruction

Geometry Reconstruction. During the initial one-time exploration phase, the agent performs frontier-based exploration [27] to collect sparse RGB-D and pose sequences. Following [21], We reconstruct 3D Gaussians parameterized by position $\boldsymbol{\mu} \in \mathbb{R}^3$, color $\boldsymbol{c} \in \mathbb{R}^3$, radius r, and opacity o. As illustrated in Fig. 3 (a), given incoming RGB-D frames and camera poses, new Gaussians are inserted in under-covered regions, and all parameters are optimized by the classical geometry loss that supervises both color and depth rendering:

$$\mathcal{L}_{geo} = (1 - \lambda)\mathcal{L}_1(\boldsymbol{C}) + \lambda\mathcal{L}_{SSIM}(\boldsymbol{C}) + \mu\mathcal{L}_1(D) \quad (1)$$

In large-scale, multi-room navigation environments with limited inter-frame overlap, geometric reconstruction often suffers from forgetting and surface holes. To mitigate this, we introduce a keyframe retrieval mechanism. At each timestep, the current frame is optimized for p_1 iterations, followed by p_2 iterations on sampled historical frames (weighted negatively correlated to PSNR), thus frames with lower fidelity are more likely to be revisited. Crucially, we opt for 3DGS over discrete point clouds due to its continuous and differentiable nature. Geometrically, Gaussians support real-time optimization and seamlessly interpolate sparse-view holes. Semantically, although optimization introduces offline overhead, this differentiable structure enables an "analysis-by-synthesis" paradigm. Unlike heuristic multi-view averaging, it actively back-propagates rendering losses to resolve cross-view conflicts, yielding a globally consistent and highly extensible semantic memory.

Language Injection. Following [26], as illustrated in Fig. 3 (b), we incorporate language information by optimizing the feature $\boldsymbol{f} \in \mathbb{R}^d$ of each Gaussian and discretizing it into a codebook $\mathcal{C}$ which is associated with high-dimensional 2D instance features $\boldsymbol{f}_{2d} \in \mathbb{R}^D$ extracted from image observations.

Concretely, instance masks are first obtained by SAM [28], and feature splatting is applied to render per-pixel semantic features. Mask-level features are aggregated and optimized with a feature loss that encourages intra-instance consistency and inter-instance separability. This initial step explicitly encodes instance boundaries, ensuring that Gaussians belonging to the same object share similar embeddings, while different objects remain distinguishable.

To achieve stable retrieval and cross-view instance align-

TABLE I: Data Comparison between GOAT-Core and GOAT-Bench of 4 Representative Rooms.

	Avg. Subtasks per Episode	Avg. Unique Categories per Episode	Avg. Inter-Subtask Distance (m)	Total Episodes	Total Subtasks
GOAT-Core (Ours)	**20**	**13.37**	**6.89**	24	**480**
GOAT-Bench (Original)	7.88	4.82	5.18	**40**	315

ment, Gaussian features are further discretized via a two-level codebook quantization. A coarse partition jointly considers 3D positions and language features, while a fine partition refines categories based solely on language features. Then we conduct a 2D-3D feature association between 2D instance-level features from multiple image views and the discretized 3D Gaussian language features. Specifically, for each instance category, we render all Gaussians assigned to that category and evaluate their spatial and semantic consistency with the 2D instance masks. The instance feature is then assigned to the discrete language category of the Gaussians with the highest similarity score.

This process produces a language codebook $\mathcal{C}$ where each entry corresponds to a cluster of 3D Gaussians enriched with CLIP features. By explicitly leveraging spatial correlations to filter out multi-view noise, this discretization ensures instance-level semantic consistency, thereby establishing a robust language-conditioned 3D memory for precise goal localization.

C. Memory-Guided Visual Navigation

Goal Localization via Memory Query. The input goal, provided as either text or an image, is encoded into an embedding vector through the corresponding CLIP encoder. Cosine similarity is then computed against features of all entries stored in the codebook to identify candidate instances. The codebook records the indices of all Gaussians associated with the selected instance. Since navigation essentially requires guiding the agent towards the spatial location of the target instance, we compute the geometric centroid of all Gaussians linked to the instance and project it onto the 2D obstacle map. The resulting projection serves as the candidate position (waypoint) for navigation.

Waypoint Navigation. After calculating the corresponding goal position on the obstacle map, the position is dilated and intersected with the traversable region to define a feasible waypoint. A collision-free path from the agent's current position to the waypoint is then planned using the classical Fast Marching Method (FMM) [29].

Goal Verification and Matching. A straightforward approach to goal verification would be applying perception models to historical frames or images rendered from the 3DGS memory before navigation. However, searching through raw frames is computationally inefficient, and more importantly, images rendered from 3DGS often contain blurring or artifacts due to the sparse views collected during rapid exploration. These artifacts severely degrade the performance of 2D matching models. Therefore, LagMemo adopts an on-site verification strategy. The 3DGS memory serves as a coarse global prior to generate candidate waypoints. Upon reaching a waypoint, the agent first performs a panoramic scan for goal verification. For text and object goals, we

incorporate SEEM [30] model for open-vocabulary instance segmentation, generating pixel-level masks for secondary validation. A waypoint is deemed to contain the target object if either the CLIP similarity or the mask confidence exceeds the corresponding thresholds. For image goals, LightGlue [31] feature matching is employed. If the goal is found, the system transitions into the Goalpoint Navigation stage. Otherwise, the waypoint is marked as invalid, and the memory is re-queried to obtain the next candidate instance. This iterative cycle of waypoint generation, navigation, and goal verification is repeated until the goal is found or the maximum step limit is reached.

Goalpoint Navigation. Once the verification mechanism confirms the target's visibility, the system switches from the coarse memory-guided waypoint to a perception-guided final goalpoint. Using the pixel-level mask provided by SEEM in combination with current depth information, the target object's footprint is projected onto the 2D obstacle map. We define the optimal final goalpoint as the nearest traversable grid cell adjacent to this projected footprint. The agent then executes a local FMM algorithm to navigate to final goal and triggers STOP.

IV. BENCHMARK

The GOAT-Bench [4] dataset, built on Habitat, has been the primary benchmark for evaluating multi-modal, open-vocabulary, multi-goal visual navigation. Despite its diverse scenes, several limitations hinder its ability to reasonably assess navigation capabilities. First, the evaluation of long-term memory and planning is insufficient: each episode contains only about 8 subtasks on average, with high goal repetition and relatively short inter-goal distances. Second, annotation quality issues are present: some text descriptions are inaccurate, object categories are semantically ambiguous, and certain scenes include missing meshes. These factors often cause agents to fail for reasons unrelated to algorithms, thereby obscuring fair comparison of navigation performance.

Instead of reporting on the entire GOAT-Bench validation split, which is broad but contains quality issues, we reorganize the benchmark by sampling a quality-controlled core subset GOAT-Core. As illustrated in Tab. I, we increase the number of subtasks per episode to 20 (compared to 7.88 in the original set) and enhance task diversity. The average distance between subtasks is also extended, imposing greater demands on memory and long-horizon planning. To ensure label reliability, we manually correct inaccurate or ambiguous text descriptions and prioritize objects with clear semantics. We further restrict all subtasks to occur on a single floor and select four reviewer-recommended high-quality multi-room scenes (`5cd`, `4ok`, `Nfv`, `Tee`), each averaging 7.25 rooms, mitigating mesh defects.

In total, GOAT-Core contains 480 multi-modal subtasks, covering 163 images, 158 objects, and 159 text goals. Curated from GOAT-Bench's "seen", "unseen", and "seen-synonyms" validation splits, it ensures open-vocabulary evaluation. The dataset also explicitly emphasizes long-horizon multi-goal navigation, with longer episodes and greater goal diversity.

Beyond multi-goal visual navigation, GOAT-Core also enables the evaluation of goal localization tasks. For each of the four multi-room scenes (average area 218.63 m^2), we conduct full-house exploration via frontier exploration, collecting RGB, depth, and odometry data. The curated 480 multi-modal goal queries with ground-truth positions can thus be used to evaluate the accuracy of goal localization after mapping (Sec. V-A).

For navigation experiments, we use the HelloRobot Stretch model as the agent. The robot has a height of 1.41 m and a base radius of 17 cm. It is equipped with an RGB camera with a resolution of 640×480, a horizontal field of view (HFOV) of $42°$, mounted at a height of 1.31 m. The robot's depth perception range is $0.5 \sim 5.0$ m. Its action space consists of: MOVE_FORWARD (0.25 m), TURN_LEFT/RIGHT (30°), LOOK_UP/DOWN (30°), and STOP. A subtask is considered successful if the agent executes the STOP action within a 200-step limit and its final position is less than 1.0 m from the target object instance.

V. EXPERIMENTS

A. Goal Localization

Settings. **1) Task:** The goal localization task requires mapping an open-vocabulary query to the memory constructed after frontier exploration and estimating the corresponding object's location, evaluated on the GOAT-Core. **2) Baseline:** We use **VLMaps** [13] as baseline, which constructs a visual-language map by first back-projecting depth into a 3D point cloud to aggregate dense VLM features, and then projecting them onto a 2D grid. **3) Metrics:** Localization accuracy is measured by the Euclidean distance between predicted and ground-truth positions, with success defined as any of the top-5 predictions falling within a 1.5 m radius of the target. **4) Hyperparameters:** Geometry optimization runs $p_1 = 30$ iterations per step for new viewpoints and $p_2 = 60$ iterations for keyframe viewpoints. Feature dimension d of Gaussians is 6. The two-level codebook uses coarse $k_1 = 32$ clusters and fine $k_2 = 5$ clusters per coarse cluster (total 160 entries).

Results. In goal localization, as illustrated in Tab. II, LagMemo achieves an overall 70.8% success rate, significantly outperforming VLMaps (58.8%) and across all modalities. These results highlight the superiority of the globally optimized 3DGS representation over discrete point
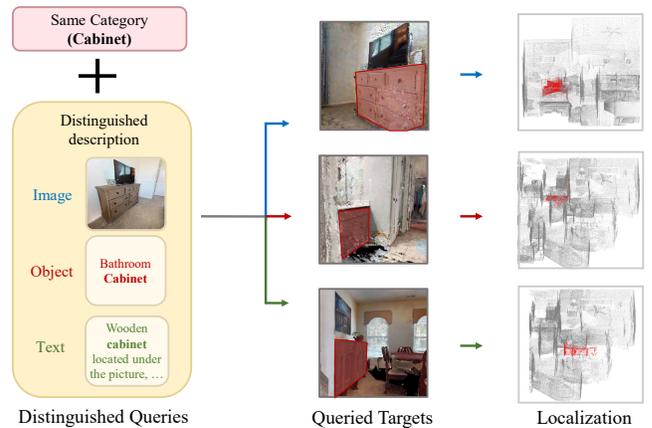

Fig. 4: Distinguished Queries Retrieving Different Instances of the Same Category in Language 3DGS Memory. For the same "cabinet" category, with distinguished queries, the language memory can retrieve the intended target. The middle column shows a geometric rendering containing queried target, and the right column presents the 3D localization of that instance.

cloud aggregation for precise goal localization in complex environments.

We also provide visualizations to demonstrate our method's precise and context-aware goal localization. As shown in Fig. 4, our system is capable of fine-grained discrimination among multiple instances of the same category.

B. Visual Navigation

Settings. **1) Task:** The visual navigation evaluation follows a sequential multi-goal protocol. Specifically, each evaluation episode comprises a sequence of subtasks. At the beginning of each episode, the agent is initialized at a random pose. During an episode, only when the current subtask is finished will the next subtask goal be provided to the agent. At every timestep, given visual observations as input, the agent outputs an action to interact with the environment. **2) Baselines:** We compare LagMemo against five baselines. **RL GOAT** [4]: the official reinforcement learning baseline of GOAT-Bench, which encodes multi-modal goals with CLIP and directly maps observations and goals to actions. **Modular GOAT** [32]: a modular approach that builds a 2D map for navigation by relying on an upstream detector to record instance images of predefined categories. **GOAT – GT Sem***: a variant of Modular GOAT using ground-truth semantic segmentation results from simulator, serving as an upper bound for reference. **GOAT – Full Exp***: a Modular GOAT variant with the same frontier exploration, aligned with LagMemo for fairer comparison. **CoWs***: adapted from the CoWs [8] method that combines frontier exploration with CLIP-based detection, extended here with multi-modal

TABLE II: Goal Localization Results on GOAT-Core (across all scenes).

Method	Average	Object	Image	Text
VLMaps [13]	58.8%	69.7%	43.3%	61.0%
LagMemo (Ours)	**70.8%**	**88.4%**	**56.4%**	**66.8%**

TABLE III: Multi-modal Visual Navigation Results on GOAT-Core.

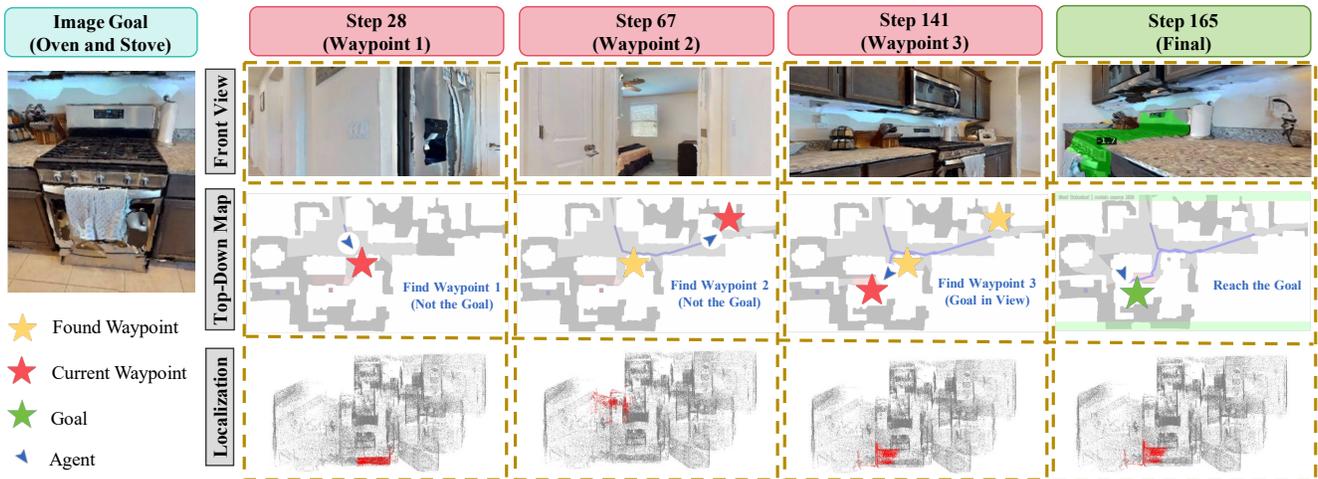| Method	OVERALL		SR BY MODALITY		
	SR (↑)	SPL (↑)	Object	Image	Text
GOAT-GT Sem*	75.0%	60.2%	86.4%	68.8%	76.9%
Modular GOAT [32]	38.3%	29.7%	36.6%	40.8%	37.8%
GOAT Full Exp*	36.3%	28.5%	39.0%	39.5%	30.5%
RL GOAT [4]	11.3%	6.2%	18.3%	5.6%	9.2%
CoWs* [8]	45.8%	28.6%	58.5%	43.3%	35.4%
LagMemo (Ours)	**56.3%**	**35.3%**	**68.3%**	**46.1%**	**53.7%**

Fig. 5: Step-by-step Visualization of Memory-Guided Navigation to an Image Goal. Columns show key steps (28, 67, 141, 165), rows show the front view, the top-down map, and the 3D localization results (red). In this case, the agent reaches waypoint-1/2/3 (yellow star; current waypoint in red). After checking the first two, it arrives at the third where the goal verification module identifies the goal. Then the agent proceeds to the final goal (green star) and the subtask successfully terminates at step 165.

TABLE IV: Visual Navigation Results on the Full GOAT-Bench.

	SEEN		SYNONYMS		UNSEEN	
Method	SR (↑)	SPL (↑)	SR (↑)	SPL (↑)	SR (↑)	SPL (↑)
GOAT-GTSem [4]	56.7%	40.3%	58.4%	43.5%	54.3%	41.0%
Modular GOAT [32]	26.3%	17.5%	33.8%	24.4%	24.9%	17.2%
Modular CoWs [8]	14.8%	8.7%	18.5%	11.5%	16.1%	10.4%
SenseAct-NN (SC) [4]	29.2%	12.8%	38.2%	15.2%	29.5%	11.3%
SenseAct-NN (Mono) [4]	16.8%	9.4%	18.5%	10.1%	12.3%	6.8%
LagMemo (Ours)	**36.8%**	**20.7%**	**44.8%**	**32.1%**	**37.9%**	**26.3%**

inputs. For closed set methods, we provide the list of target categories in advance to ensure feasibility, as these methods cannot operate without predefined categories. We do not compare against recent 3DGS-based image-goal methods GaussNav[16] and IGL-Nav[17] due to their incomplete open-source and reliance on high-fidelity RGB rendering, which are incompatible with our sparse-viewpoint setting. **3) Metrics:** Performance is assessed using two common metrics: SR (Success Rate), which measures the proportion of successful subtasks and SPL (Success weighted by Path Length), which evaluates efficiency. **4) Hyperparameters:** The goal verification module is triggered when the agent is within 1.2 m of the candidate waypoint. For object/text goals, goal verification passes if either threshold is met: SEEM score $\tau_{SEEM} \geq 1.1$ or Mobile-CLIP cosine similarity $\tau_{CLIP} \geq 0.23$. For image goals, LightGlue declares target presence if the inlier match ratio $r_{match} \geq 5\%$.

Evaluation on GOAT-Core. We first evaluate our method on the GOAT-Core split (Sec. IV) to ensure a high-quality, statistically meaningful comparison. As shown in Tab. III, LagMemo significantly outperforms all baselines in overall navigation, surpassing the second-best method (CoWs*) by 10.5% in SR and achieving the highest SPL (35.3%). Compared to GOAT-Full Exp under the same pre-exploration setting, LagMemo improves SR by 20.0% and SPL by 6.8%, confirming the crucial benefit of our 3DGS memory over 2D projections. Furthermore, the modality breakdown demonstrates LagMemo's robustness across diverse query

modalities. Its notable advantage on text queries specifically highlights the superiority of our language-quantized code-book for open-vocabulary retrieval.

Evaluation on Full GOAT-Bench. While GOAT-Core evaluates long-horizon memory, we further test on the full GOAT-Bench validation set to assess broad generalization. As shown in Tab. IV, LagMemo consistently achieves the highest SR across all three splits. These comprehensive results demonstrate that LagMemo's architectural advantages persist even under the noisy and complex original benchmark.

Qualitative Case Study. Fig. 5 visualizes a step-by-step navigation task for an image goal ("oven and stove"), which demonstrates an effective interplay between long-term memory guidance and goal verification mechanism.

Failure Case Analysis. We compare LagMemo against GOAT-Full Exp across three failure types: Memory Indexing Error (failure to retrieve the correct target location from memory), Verification Error (failure to confirm the target locally), and Navigation Error (e.g., collisions or loops). Overall, compared to Modular GOAT, LagMemo reduces the total failure rate from 63.7% to 43.7%. Specifically, Memory Indexing Error drops significantly from 37.1% to 22.1%, and Verification Error decreases from 22.1% to 18.3%. These reductions highlight two core advantages: first, our codebook-quantized 3DGS memory effectively preserves 3D spatial context and filters multi-view noise to ensure reliable retrieval; second, our multimodal on-site verification module robustly handles local target confirmation.

C. Ablation Study

Memory Reconstruction. Tab. V(a) validates our memory design. Removing the keyframe mechanism degrades geo-metric quality and lowers localization accuracy, confirming that accurate geometry is foundational for goal localization. Furthermore, replacing the discrete codebook with a 2D-trained autoencoder drastically reduces accuracy. This high-lights that explicit 3D spatial clustering is indispensable for consistent multi-room memory.

a) Environment   **b) Memory**   **c) Goals**   **d) Localization**   **e) Navigation**
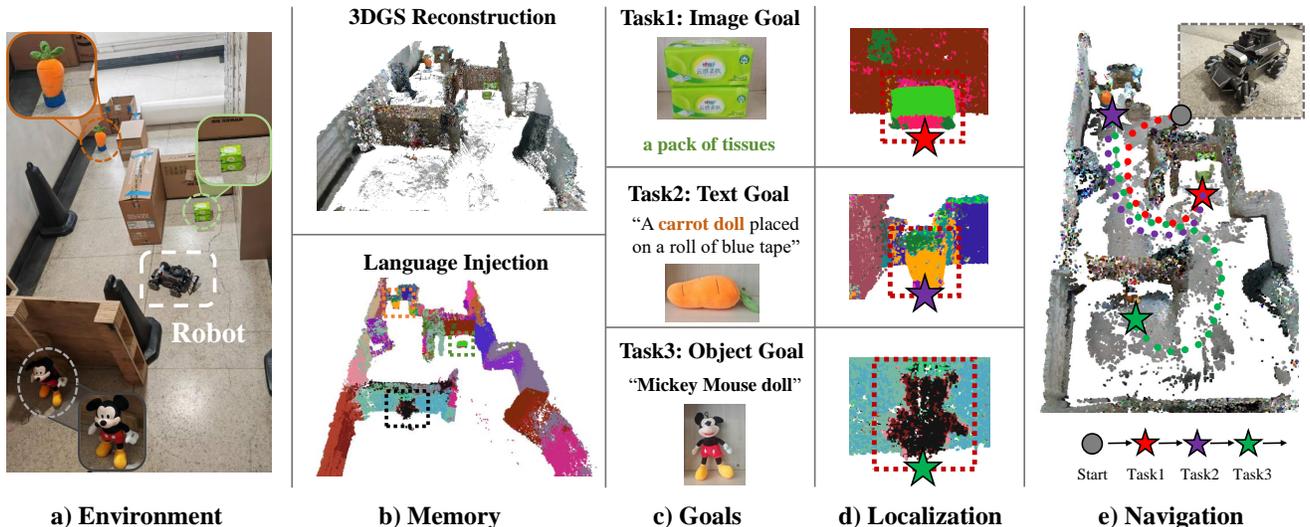
Fig. 6: Real-world Deployment of LagMemo. In a physical indoor environment, we reconstruct a 3DGS memory, and successfully locate and navigate to sequential multi-modal open-vocabulary goals, such as a "Mickey Mouse" doll.

TABLE V: Ablation Studies on Memory and Verification.

(a) Memory Reconstruction (Goal Localization Accuracy)

Keyframe	Codebook	PSNR	Avg.	Obj.	Img.	Text
✕	✓	21.15	66.3%	77.5%	57.5%	63.4%
✓	✕	**27.20**	34.6%	41.6%	21.0%	37.1%
✓	✓	**27.20**	**70.8%**	**88.4%**	**56.4%**	**66.8%**

(b) Goal Verification (Navigation SR & SPL)

Image Match	Text Match	Average		Obj.	Img.	Text
		SR (↑)	SPL (↑)	SR (↑)	SR (↑)	SR (↑)
✕ (No Verif.)	✕ (No Verif.)	41.3%	30.4%	45.1%	32.9%	45.1%
CLIP	CLIP	46.7%	30.3%	52.4%	43.4%	43.9%
LightGlue	SEEM + CLIP	**56.3%**	**35.3%**	**68.3%**	**46.1%**	**53.7%**

Goal Verification. Tab. V(b) analyzes the verification module. Stopping without verification yields only a 41.3% SR. A naive CLIP-based similarity marginally improves SR to 46.7%. Our modality-specific strategy (LightGlue for images, SEEM+CLIP for text/objects) achieves the highest SR (56.3%) and SPL (35.3%), proving its necessity in robustly confirming targets and mitigating memory noise.

D. Efficiency Analysis

We report the memory construction time, storage, and query latency on an NVIDIA RTX A6000 GPU, operating in an average 200 m^2 scene. As shown in Tab. VI(a), LagMemo requires higher offline build time and storage due to the dense 3DGS optimization. While this time is primarily dominated by comprehensive multi-view feature learning, it represents a strictly one-time offline cost per environment. By investing computational time in global spatial-semantic optimization during pre-exploration, LagMemo explicitly resolves multi-view feature inconsistencies and establishes robust 3D spatial correlations. Consequently, our discrete codebook enables fast index lookups against the established memory, achieving a near-instantaneous 0.5s query latency.

Furthermore, Tab. VI(b) details the per-step inference latency during navigation. The goal verification module operates conditionally (executing specific matching models only when necessary). Combined with the planner, the total inference time is 626ms per step. This confirms that despite the mapping cost, LagMemo ensures real-time navigation,

TABLE VI: System Efficiency.

(a) Memory Building and Query Efficiency

Method	Build Time (s) ↓	Query Latency (s) ↓	Storage (MB) ↓
GOAT [32]	**1260**	>10*	∼400
VLMaps [13]	∼2000	1.1	**∼200**
LagMemo (Ours)	∼4200	**0.5**	∼500

* Matching takes 0.23s per image; GOAT stores hundreds of images, leading to high query latency.

(b) Per-Step Inference Time during Navigation

Module Category	Component	Latency (ms)
Goal Verification	Mobile CLIP	225
	SEEM (Text Goal)	190
	LightGlue (Image Goal)	152
	Subtotal	*396*
Action Planning	FMM Planner	230
Total Inference Time	*per step*	**626**

making it well-suited for deployment on real-world robotic platforms.

E. Real-world Deployment

As shown in Fig. 6, we deploy our system on a physical differential-drive robot. The hardware setup comprises an onboard NVIDIA Jetson Orin NX and a Realsense D435i RGB-D camera. The language 3DGS memory construction is offloaded to a remote server equipped with an NVIDIA RTX A6000 GPU, while the real-time perception, goal verification, and path planning run entirely onboard the Jetson Orin NX. Although the inaccuracy of depth camera and inherent odometry drift lead to a sub-optimal geometric reconstruction, LagMemo's codebook-quantized language memory demonstrated robustness, and successfully localizes multi-modal open-vocabulary queries (e.g., "Mickey Mouse doll") and navigates to the intended instances.

VI. CONCLUSIONS

We present LagMemo, a navigation system that integrates language 3D Gaussian Splatting to address the practical demands of visual navigation with multi-modal and open-vocabulary goal inputs as well as multi-goal tasks. By encoding language features into a codebook-quantized memory

through a one-time exploration and employing an on-site perception-based goal verification mechanism, LagMemo effectively bridges global scene memory with local perception, enabling efficient goal localization and reliable navigation. Extensive experiments and real-world deployment demonstrate the superiority and practical efficiency of LagMemo.

Despite the encouraging results in integrating semantic 3D memory with visual navigation, several directions remain open. (1) Memory-aware Exploration. LagMemo's goal localization capability hinges on geometric fidelity. Insufficient view coverage leaves blind spots in memory. Future work includes developing memory-aware active exploration that leverages geometric and semantic uncertainty to select informative views. (2) Incremental Semantics and Online Memory Construction. Currently, our global optimization and codebook clustering are performed post-exploration to ensure maximum consistency. Real-world environments, however, are dynamic, requiring adaptation. Drawing inspiration from recent advances in online language splatting, our future work will focus on transitioning to a fully incremental language 3DGS architecture, enabling both geometric optimization and online codebook growth to run concurrently during navigation. (3) Task-specific Memory Compression. Full-scene 3D representations are memory-intensive and compute-intensive. A promising direction to further improve both time efficiency and storage overhead is to pursue hierarchical or multi-resolution Gaussians, uncertainty-aware pruning, and feature compression explicitly tailored to embodied navigation needs.

REFERENCES

[1] M. Deitke, D. Batra, Y. Bisk, T. Campari, A. X. Chang, D. S. Chaplot, C. Chen, C. P. D'Arpino, K. Ehsani, A. Farhadi *et al.*, "Retrospectives on the embodied ai workshop," *arXiv preprint arXiv:2210.06849*, 2022.

[2] J. Sun, J. Wu, Z. Ji, and Y.-K. Lai, "A survey of object goal navigation," *IEEE Transactions on Automation Science and Engineering*, 2024.

[3] L. H. K. Wong, X. Kang, K. Bai, and J. Zhang, "A survey of robotic navigation and manipulation with physics simulators in the era of embodied ai," *arXiv preprint arXiv:2505.01458*, 2025.

[4] M. Khanna, R. Ramrakhya, G. Chhablani, S. Yenamandra, T. Gervet, M. Chang, Z. Kira, D. S. Chaplot, D. Batra, and R. Mottaghi, "Goat-bench: A benchmark for multi-modal lifelong navigation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 16 373–16 383.

[5] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik *et al.*, "Habitat: A platform for embodied ai research," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9339–9347.

[6] D. S. Chaplot, D. P. Gandhi, A. Gupta, and R. R. Salakhutdinov, "Object goal navigation using goal-oriented semantic exploration," *Advances in Neural Information Processing Systems*, vol. 33, pp. 4247–4258, 2020.

[7] R. Ramrakhya, E. Undersander, D. Batra, and A. Das, "Habitat-web: Learning embodied object-search strategies from human demonstrations at scale," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 5173–5183.

[8] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 23 171–23 181.

[9] N. Yokoyama, S. Ha, D. Batra, J. Wang, and B. Bucher, "Vlfm: Vision-language frontier maps for zero-shot semantic navigation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 42–48.

[10] D. Batra, A. Gokaslan, A. Kembhavi, O. Maksymets, R. Mottaghi, M. Savva, A. Toshev, and E. Wijmans, "Objectnav revisited: On evaluation of embodied agents navigating to objects," *arXiv preprint arXiv:2006.13171*, 2020.

[11] J. Krantz, S. Lee, J. Malik, D. Batra, and D. S. Chaplot, "Instance-specific image goal navigation: Training embodied agents to find object instances," *arXiv preprint arXiv:2211.15876*, 2022.

[12] X. Sun, L. Liu, H. Zhi, R. Qiu, and J. Liang, "Prioritized semantic learning for zero-shot instance navigation," in *European Conference on Computer Vision*. Springer, 2024, pp. 161–178.

[13] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 10 608–10 615.

[14] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard, "Hierarchical open-vocabulary 3d scene graphs for language-grounded robot navigation," *arXiv preprint arXiv:2403.17846*, 2024.

[15] H. Yin, X. Xu, L. Zhao, Z. Wang, J. Zhou, and J. Lu, "Unigoal: Towards universal zero-shot goal-oriented navigation," *arXiv preprint arXiv:2503.10630*, 2025.

[16] X. Lei, M. Wang, W. Zhou, and H. Li, "Gaussnav: Gaussian splatting for visual navigation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

[17] W. Guo, X. Xu, H. Yin, Z. Wang, J. Feng, J. Zhou, and J. Lu, "Igl-nav: Incremental 3d gaussian localization for image-goal navigation," *arXiv preprint arXiv:2508.00823*, 2025.

[18] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering." *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.

[19] G. Chen and W. Wang, "A survey on 3d gaussian splatting," 2025.

[20] H. Matsuki, R. Murai, P. H. Kelly, and A. J. Davison, "Gaussian splatting slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 18 039–18 048.

[21] N. Keetha, J. Karhade, K. M. Jatavallabhula, G. Yang, S. Scherer, D. Ramanan, and J. Luiten, "Splatam: Splat track & map 3d gaussians for dense rgb-d slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21 357–21 366.

[22] S. Zhou, H. Chang, S. Jiang, Z. Fan, Z. Zhu, D. Xu, P. Chari, S. You, Z. Wang, and A. Kadambi, "Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21 676–21 685.

[23] J.-C. Shi, M. Wang, H.-B. Duan, and S.-H. Guan, "Language embedded 3d gaussians for open-vocabulary scene understanding," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 5333–5343.

[24] M. Qin, W. Li, J. Zhou, H. Wang, and H. Pfister, "Langsplat: 3d language gaussian splatting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 20 051–20 060.

[25] S. Katragadda, C.-Y. Wu, Y. Guo, X. Huang, G. Huang, and L. Ren, "Online language splatting," *arXiv preprint arXiv:2503.09447*, 2025.

[26] Y. Wu, J. Meng, H. Li, C. Wu, Y. Shi, X. Cheng, Z. Zhao, H. Feng, E. Ding, J. Wang *et al.*, "Opengaussian: Towards point-level 3d gaussian-based open vocabulary understanding," *arXiv preprint arXiv:2406.02058*, 2024.

[27] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Proceedings of the second international conference on Autonomous agents*, 1998, pp. 47–53.

[28] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, "Segment anything," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4015–4026.

[29] J. A. Sethian, "A fast marching level set method for monotonically advancing fronts." *proceedings of the National Academy of Sciences*, vol. 93, no. 4, pp. 1591–1595, 1996.

[30] X. Zou, J. Yang, H. Zhang, F. Li, L. Li, J. Wang, L. Wang, J. Gao, and Y. J. Lee, "Segment everything everywhere all at once," *Advances in neural information processing systems*, vol. 36, pp. 19 769–19 782, 2023.

[31] P. Lindenberger, P.-E. Sarlin, and M. Pollefeys, "Lightglue: Local feature matching at light speed," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 17 627–17 638.

[32] M. Chang, T. Gervet, M. Khanna, S. Yenamandra, D. Shah, S. Y. Min, K. Shah, C. Paxton, S. Gupta, D. Batra, R. Mottaghi, J. Malik, and D. S. Chaplot, "Goat: Go to any thing," 2023.