

MapNav: A Novel Memory Representation via Annotated Semantic Maps for VLM-based Vision-and-Language Navigation

Lingfeng Zhang^{1,2,3,*} Xiaoshuai Hao^{2,*,†} Qinwen Xu^{2,6} Qiang Zhang^{1,4}
 Xinyao Zhang¹ Pengwei Wang² Jing Zhang⁵ Zhongyuan Wang²
 Shanghang Zhang^{2,6,✉} Renjing Xu^{1,✉}

¹The Hong Kong University of Science and Technology (Guangzhou)

² Beijing Academy of Artificial Intelligence, ³ Tsinghua University

⁴Beijing Innovation Center of Humanoid Robotics Co., Ltd.

⁵ School of Computer Science, Wuhan University

⁶ State Key Laboratory of Multimedia Information Processing,
 School of Computer Science, Peking University

Abstract

Vision-and-language navigation (VLN) is a key task in Embodied AI, requiring agents to navigate diverse and unseen environments while following natural language instructions. Traditional approaches rely heavily on historical observations as spatio-temporal contexts for decision making, leading to significant storage and computational overhead. In this paper, we introduce **MapNav**, a novel end-to-end VLN model that leverages Annotated Semantic Map (ASM) to replace historical frames. Specifically, our approach constructs a top-down semantic map at the start of each episode and update it at each timestep, allowing for precise object mapping and structured navigation information. Then, we enhance this map with explicit textual labels for key regions, transforming abstract semantics into clear navigation cues and generate our ASM. **MapNav** agent using the constructed ASM as input, and use the powerful end-to-end capabilities of VLM to empower VLN. Extensive experiments demonstrate that **MapNav** achieves state-of-the-art (SOTA) performance in both simulated and real-world environments, validating the effectiveness of our method. Moreover, we will release our ASM generation source code and dataset to ensure reproducibility, contributing valuable resources to the field. We believe that our proposed **MapNav** can be used as a new memory representation method in VLN, paving the way for future research in this field.

1 Introduction

Vision-and-language Navigation (VLN) (Gu et al., 2022; Park and Kim, 2023) is a crucial integra-

* Co-first Authors.

lfzhang715@gmail.com, xshao@baai.ac.cn

† Project Leader.

✉ Corresponding Authors.

renjingxu@hkust-gz.edu.cn, shanghang@pku.edu.cn



Figure 1: **Illustration of our Annotated Semantic Map (ASM).** At each timestep, MapNav agent leverages egocentric observations to capture semantic objects and assign explicit textual labels to key regions, creating the ASM for the current moment. ASM provides information such as physical obstacles, explored regions, the agent’s current position, trajectory and semantic objects.

tion of embodied AI and multimodal understanding (Hao et al., 2025b; Ji et al., 2025; Tang et al., 2025), empowering autonomous agents to interpret natural language instructions and navigate complex, unseen environments. Unlike traditional navigation tasks that rely solely on visual input or predetermined waypoints, VLN requires a complex blend of language understanding and visual perception. For example, agents must interpret subtle instructions (Vasudevan et al., 2021; An et al., 2024; Chen

et al., 2024) such as “walk past the potted plants and turn left in front of the wooden cabinet”, while also processing dynamic visual scenes and making real-time navigation decisions.

Existing VLN methods can be classified into discrete and continuous navigation paradigms, primarily distinguished by their representations of action space. Discrete environment navigation methods leverage MP3D (Chang et al., 2017) and abstract the navigation space into a connected graph structure for waypoint selection. While these methods have demonstrated impressive performance on benchmark datasets (Shah et al., 2023; Zhou et al., 2024; Long et al., 2024b), they fail to reflect the continuity of real-world navigation. To address this gap, the Habitat simulator (Savva et al., 2019) and benchmarks like R2R-CE (Krantz et al., 2020) and RxR-CE (Ku et al., 2020a) have enabled research in continuous navigation. Moreover, to minimize the sim-to-real gap, Habitat implements a low-level action space with forward movement and rotational actions. Existing continuous environment navigation methods (Zhang et al., 2024a; Zheng et al., 2024) heavily rely on historical robot observations as spatio-temporal contexts for decision-making and instruction following. However, these approaches significantly increase storage requirements and lack a structured understanding of past trajectories. Therefore, designing a novel memory representation to effectively replace traditional historical frames is of great significance and thus becomes the motivation for our work.

To fill this gap, we propose a novel end-to-end VLM-based VLN model, *MapNav*, which leverages Annotated Semantic Maps for innovative memory representation, effectively replacing traditional historical frames. Specifically, we first transform RGB-D and pose data into point cloud representations to generate precise top-down visualizations. Then, we align semantic segmentation to construct a base semantic map, which is further enhanced into an Annotated Semantic Map (ASM) by integrating explicit textual annotations for salient regions and abstract semantic concepts. As shown in Fig. 1, the ASM is initialized at the beginning of each episode and updated at each timestep for real-time operation. It encapsulates crucial environmental information for navigation, including physical obstacle distributions, explored areas, agent position, historical trajectories, and semantic object locations. The *MapNav* agent utilizes the constructed ASM as input, harnessing the

powerful end-to-end capabilities of VLM to enhance vision-and-language navigation. Extensive experiments highlight the pivotal role of ASM in VLN tasks, showcasing memory representation capabilities that are comparable to traditional historical frame-based approaches. Additionally, ASM enables VLMs to develop structured spatial understanding and optimize path selection effectively. These advantages make ASM a promising method for enhancing navigation performance in various environments. We will release a dataset containing **1 Million** step-wise samples, featuring RGB frames, corresponding ASMs, VLN instructions, and frame-specific actions to promote this field. We believe that our proposed *MapNav* can serve as an innovative memory representation method in VLN, paving the way for future advancements in this field. Our main contributions are summarized as follows:

- We propose a novel end-to-end VLM-based vision-and-language navigation model, *MapNav*, which leverages Annotated Semantic Maps for innovative memory representation, effectively replacing traditional historical frames.
- We introduce a top-down Annotated Semantic Map (ASM) that we update at each timestep, enabling precise object mapping and structured navigation, while enhancing it with explicit textual labels for key regions to provide clear navigation cues.
- *MapNav* outperforms state-of-the-art (SOTA) methods in both simulated and real-world environments, offering a new memory representation in VLN and paving the way for future research in the field.
- We will release our ASM generation source code and dataset, allowing for the reproducibility of the results presented in this study, which will serve as a valuable contribution to the field.

2 Related Work

Vision-and-Language Navigation Vision-and-Language Navigation (VLN) (Gu et al., 2022; Park and Kim, 2023) in embodied AI (Li et al., 2024b; Zhang et al., 2025; Hao et al., 2025a) focuses on navigating unseen environments by following human instructions, primarily in discretized simulated scenarios (Ku et al., 2020b; Thomason et al., 2020). Agents navigate between predefined nodes on a




Figure 2: **An overview of MapNav framework.** We present a top-down Annotated Semantic Map (ASM), updated at each timestep for precise object mapping and structured navigation. It features explicit textual labels for key regions, providing clear navigation cues. The current RGB observation, ASM, and instruction are used as inputs to an end-to-end VLM framework, which generates navigation actions in natural language.

graph by integrating language and visual input (Qi et al., 2020; Liu et al., 2024), but this reliance complicates real-world deployment. To address this, VLN for Continuous Environments (VLN-CE) (Krantz et al., 2020; Savva et al., 2019) enable unrestricted navigation through low-level control or waypoint-based methods (Hong et al., 2022; Krantz et al., 2021), improving sim-to-real transferability despite added complexity. Recent advancements in vision-language models have significantly influenced VLN development, utilizing large-scale pre-trained models (Zhang et al., 2024a; Zheng et al., 2024; Wu et al., 2025) and VLN-specific pre-training (Hao et al., 2020; Wu et al., 2022). For instance, NavGPT (Zhou et al., 2024) autonomously generates actions using GPT-4o, while DiscussNav (Long et al., 2024b) employs VLN experts to reduce human involvement. InstructNav (Long et al., 2024a) decomposes navigation into subtasks, and Nav-CoT (Lin et al., 2024) uses Chain of Thoughts (CoT) (Wei et al., 2022) for environmental simulation. Some methods (Zheng et al., 2024; Zhang et al., 2024a) fine-tune VLMs for specific navigation tasks, highlighting flexibility. However, existing approaches often depend on hierarchical prompts or historical frames, leading to high memory demands and limited understanding of past data.

This paper introduces a novel memory representation using Annotated Semantic Maps (ASMs) to effectively replace traditional historical frames.

Map Representations for VLN Structured maps in VLN enhance navigation performance by improving environmental understanding (Wang et al., 2023; Hong et al., 2023; Hao et al., 2024a, 2025c, 2024b). Methods like MC-GPT (Zhan et al., 2024) and VoroNav (Wu et al., 2024) utilize topological maps to capture viewpoints and spatial relationships, while InstructNav (Long et al., 2024a) and VLFM (Yokoyama et al., 2024) create value maps for waypoint selection. Semantic maps (Zhang et al., 2024c,b; Hong et al., 2023; Yu et al., 2023) retain object-level information for navigation. However, these maps are often not interpretable by VLMs. To address this, we propose Annotated Semantic Map (ASM), a novel semantic map representation that allows VLMs to explicitly understand rich map information, including obstacle distributions, explored areas, agent positions, historical trajectories, and semantic object locations. ASM aims to establish a new memory representation for VLN.




Figure 3: **ASM Generation Process.** Semantic map generation starts with episode initialization. At each timestep, the RGB image is processed by a semantic segmentation module to create a semantic mask aligned with the depth-converted 3D point cloud. By combining this with the previous pose transformation, we project the 3D point cloud onto a 2D plane to update the semantic map. Finally, we convert the semantic map into the ASM through region clustering and text annotation, yielding a comprehensive memory representation with labeled objects.

3 Methodology

MapNav pursues a novel end-to-end vision-and-language navigation model based on VLM, leveraging Annotated Semantic Maps (ASM) to replace traditional historical frames. We introduce a top-down annotated semantic map at the start of each episode, updating it at each timestep for precise object mapping and structured navigation information, while enhancing it with textual labels for key regions to clarify navigation cues. The MapNav agent uses the constructed ASM as input, harnessing VLM’s powerful end-to-end capabilities to enhance vision-and-language navigation (VLN). The framework overview of MapNav is shown in Fig. 2.

3.1 Task Definition

Vision-and-language Navigation in continuous environments (VLN-CE) refers to a task in which an agent navigates continuous 3D environments using natural language instructions as guidance. The input to a VLN-CE agent comprises two key components: (1) a natural language instruction $\mathcal{I}$ that specifies the desired navigation path (e.g., “Walk down the hallway, turn right at the plant, and stop at the third door on your left”), and (2) a sequence of first-person RGB observations $\mathbf{X}_t$ collected as the agent navigates through the environment. At each timestep t , the agent must predict a continuous action $\mathbf{a}_{t+1} \in \mathcal{A}$ that defines its next movement, where $\mathcal{A}$ denotes the low-level action repertoire. The agent must execute these actions sequentially until it determines that it has reached the target destination specified in the instruction.

3.2 Our MapNav Agent

Our MapNav agent uses the current RGB observation, Annotated Semantic Maps (ASM), and in-

struction to directly generate executable navigation actions. Specifically, we first extract the object mask from the current observation $\mathbf{X}_t$ using a semantic segmentation module, then construct the semantic map by integrating the depth image and pose. To enhance readability for the VLM, we generate Annotated Semantic Maps. We then employ a two-stream encoder to process the observation and ASM features separately, followed by a multi-modal projector to align these modalities in a shared embedding space. Finally, we concatenate the instruction tokens with the aligned features and input them into the VLM, which outputs executable navigation actions in text format. The detailed architectural designs are described as follows.

Annotated Semantic Maps (ASM) Generation.

The generation process of our ASM is shown in Fig. 3. We employ a semantic mapping system to create a rich environmental representation using a multi-channel tensor $\mathbf{M}$ of dimensions $C \times W \times H$, where $C = C_n + 4$ and n represents distinct object categories. The foundational channels (1-4) encode navigational information: physical obstacles, explored regions, the agent’s current position, and historical locations, while the remaining n channels store object-specific semantic information. At the start of each navigation episode, we initialize a new semantic map with the agent at $(\frac{W}{2}, \frac{H}{2})$. The map is constructed by transforming RGB-D data into point clouds, which are projected onto a 2D plane for a top-down view. By aligning semantic segmentation masks with point cloud data, we achieve accurate object-wise mapping in the dedicated channels.

Remarks: A key innovation in our approach is the generation of ASM following semantic map construction, enhancing traditional representations

with explicit natural language annotations. The ASM generation pipeline involves two main stages: (1) semantic region identification via connected component analysis on each object-specific channel, and (2) centroid computation for these regions to determine optimal text placement. For each semantic region exceeding the minimum area threshold τ , we compute its geometric centroid to establish a text anchor point, ensuring optimal annotation placement and readability while preserving visual clarity. As shown in Fig. 3, the resulting ASM transforms abstract semantic representations into linguistically grounded spatial information (e.g., “chair”, “plant”, “bed”), bridging spatial understanding and natural language comprehension. This explicit textual grounding allows the VLM to leverage its pre-trained knowledge of object-language relationships, facilitating intuitive spatial reasoning and improved navigation decision-making. As shown in Fig. 4, we conducted a comparative analysis of the proposed ASM against conventional top-down and semantic maps in simulator. To evaluate comprehension, we input these three map representations into both GPT-4o and MapNav. The results clearly demonstrate that the VLM exhibits superior understanding when processing ASM compared to traditional mapping approaches, validating our linguistic augmentation strategy. The ASM generation process fundamentally transforms the traditional semantic mapping paradigm by introducing a novel layer of linguistic information that aligns with the VLM’s pre-trained capabilities, enabling more effective multi-modal reasoning in navigation tasks. Appendix H demonstrates that our ASM successfully directs the attention of the VLM towards the textual labels.

Inputs Encoding. MapNav utilizes an advanced vision-language architecture based on LLaVA-Onevision (Li et al., 2024a). At each timestep t , we process two primary visual inputs: the observation frame $\mathbf{X}_t$ and the ASM $\mathbf{M}_t$. These inputs are concurrently handled by a shared vision encoder, followed by modality-specific projectors for optimal feature alignment. Employing a SigLIP (Zhai et al., 2023) vision encoder, we derive initial visual embeddings for both inputs. The observation frame produces $\mathbf{X}_t \in \mathbb{R}^{N \times D}$, where N represents the number of image patches and D is the model’s hidden dimension. Similarly, the ASM generates $\mathbf{X}_t^M \in \mathbb{R}^{N \times D}$. Both visual features undergo parallel spatial-aware transformations through a patch merge operation:




Figure 4: Comparison of different VLM’s understanding of different map formats includes top-down, semantic map and our ASM.

$$\begin{aligned} \mathbf{F}_t &= \Phi_{spatial}(\mathbf{X}_t, \mathcal{G}), \\ \mathbf{F}_t^M &= \Phi_{spatial}(\mathbf{X}_t^M, \mathcal{G}), \end{aligned} \quad (1)$$

where $\Phi_{spatial}$ denotes our spatial unpadding patch merge function and $\mathcal{G}$ defines the grid pinpoints for feature extraction. The features are subsequently projected into the language model’s embedding space using modality-specific MLP projector:

$$\mathbf{E}_t = P_{mlp}^{obs}(\mathbf{F}_t), \quad \mathbf{E}_t^M = P_{mlp}^{map}(\mathbf{F}_t^M), \quad (2)$$

where $\mathbf{E}_t, \mathbf{E}_t^M \in \mathbb{R}^{N \times C}$, with C as the language model’s hidden dimension. The projections P_{mlp}^{obs} and P_{mlp}^{map} are two-layer MLPs with GELU activation. The final multimodal representation is formed by concatenating the encoded observations and map features, along with special tokens for task structuring:

$$\mathbf{V}_t = [\text{TASK}; \mathbf{E}_t; \text{OBS}; \mathbf{E}_t^M; \text{MAP}]. \quad (3)$$

This unified representation seamlessly integrates multiple input modalities while preserving essential spatial relationships for navigation. By employing efficient token management and precision optimization, the encoding process ensures computational efficiency while maximizing the model’s capacity to represent complex navigation scenarios.

Action Prediction. Our framework implements a fully end-to-end Action Prediction strategy that directly generates navigation commands in natural language form. Unlike traditional approaches

Table 1: Comparison with state-of-the-art methods on the Val-Unseen split of R2R-CE and RxR-CE. * indicates methods using the waypoint predictor. *Cur. RGB* and *His. RGB* refer to methods utilize current and historical RGB frames, respectively. † indicates that we reproduced the results using the open-source code. *Pano.* represents methods utilize panoramic views.

Method	Observation					R2R Val-Unseen				RxR Val-Unseen			
	Cur. RGB	Pano.	Depth	Odo.	His. RGB	NE ↓	OS ↑	SR ↑	SPL ↑	NE ↓	SR ↑	SPL ↑	nDTW ↑
HPN+DN*		✓	✓	✓	✓	6.31	40.0	36.0	34.0	-	-	-	-
CMA*		✓	✓	✓	✓	6.20	52.0	41.0	36.0	8.76	26.5	22.1	47.0
VLN ² BERT*		✓	✓	✓	✓	5.74	53.0	44.0	39.0	8.98	27.0	22.6	46.7
Sim2Sim*		✓	✓	✓	✓	6.07	52.0	43.0	36.0	-	-	-	-
GridMM*		✓	✓	✓	✓	5.11	61.0	49.0	41.0	-	-	-	-
Ego ² -Map*		✓	✓	✓	✓	5.54	56.0	47.0	41.0	-	-	-	-
DreamWalker*		✓	✓	✓	✓	5.53	59.0	49.0	44.0	-	-	-	-
Reborn*		✓	✓	✓	✓	5.40	57.0	50.0	46.0	5.98	48.6	42.0	63.3
ETPNav*		✓	✓	✓	✓	4.71	65.0	57.0	49.0	5.64	54.7	44.8	61.9
HNR*		✓	✓	✓	✓	4.42	67.0	61.0	51.0	5.50	56.3	46.7	63.5
BEVBert*		✓	✓	✓	✓	4.57	67.0	59.0	50.0	4.00	68.5	-	69.6
HAMT+ScaleVLN*		✓	✓	✓	✓	4.80	-	55.0	51.0	-	-	-	-
AG-CMTP		✓	✓	✓	✓	7.90	39.0	23.0	19.0	-	-	-	-
R2R-CMTP		✓	✓	✓	✓	7.90	38.0	26.0	22.0	-	-	-	-
LAW	✓		✓	✓	✓	6.83	44.0	35.0	31.0	10.90	8.0	8.0	38.0
CM2	✓		✓	✓	✓	7.02	41.0	34.0	27.0	-	-	-	-
WS-MGMap	✓		✓	✓	✓	6.28	47.0	38.0	34.0	-	-	-	-
AO-Planner		✓	✓	✓	✓	5.55	59.0	47.0	33.0	7.06	43.3	30.5	50.1
Seq2Seq	✓		✓	✓	✓	7.77	37.0	25.0	22.0	12.10	13.9	11.9	30.8
CMA	✓		✓	✓	✓	7.37	40.0	32.0	30.0	-	-	-	-
RGB-Seq2Seq	✓				✓	10.10	8.0	0.0	0.0	-	-	-	-
RGB-CMA	✓				✓	9.55	10.0	5.0	4.0	-	-	-	-
NaVid (All RGB Frames)	✓				✓	5.47	49.1	37.4	35.9	8.41	23.8	21.2	-
NaVid (Cur. RGB)†	✓					8.10	24.9	13.0	7.8	11.33	8.7	6.8	-
MapNav(w/o ASM + Cur. RGB)	✓					7.26	41.2	27.1	23.5	9.31	15.6	12.2	30.9
MapNav(w/ ASM + Cur. RGB)	✓					5.22	50.3	36.5	34.3	8.95	22.1	20.2	35.6
MapNav(w/ ASM + Cur. RGB + 2 His. RGB)	✓				✓	4.93	53.0	39.7	37.2	7.62	32.6	27.7	43.5

Table 2: Comparison in diverse real-world environments scenes (*Meeting Room*, *Office*, *Lecture Hall*, *Tea Room*, and *Living Room*). Simple I.F. and Semantic I.F. indicate simple and semantic instruction following tasks, respectively. Our MapNav outperforms all the baselines in both simple instructions and semantic instructions.

Method	Meeting Room				Office				Lecture Hall				Tea Room				Living Room			
	Simple I.F.		Semantic I.F.		Simple I.F.		Semantic I.F.		Simple I.F.		Semantic I.F.		Simple I.F.		Semantic I.F.		Simple I.F.		Semantic I.F.	
	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓	SR↑	NE↓
WS-MGMap	50%	1.62	20%	2.83	60%	1.21	30%	3.11	60%	0.95	20%	3.63	50%	1.11	20%	2.93	70%	0.62	40%	2.86
Navid	70%	0.86	50%	1.93	70%	1.59	60%	1.99	60%	0.75	40%	2.94	80%	0.35	50%	1.61	60%	2.13	30%	3.83
MapNav(Ours)	70%	0.73	60%	1.32	80%	0.96	60%	1.38	80%	0.82	70%	1.15	90%	0.31	70%	0.66	80%	1.03	60%	0.85

that require separate action decoders, we leverage the VLM’s inherent language understanding capabilities to directly parse navigation intentions into discrete actions.

To robustly interpret the model’s natural language outputs, we employ a comprehensive pattern matching system capable of recognizing various linguistic expressions of the same navigational intent:

$$\mathcal{A}(t) = \Psi(\mathcal{T}(t), \mathcal{P}), \quad (4)$$

where $\mathcal{T}(t)$ denotes the model’s text output at time t , $\mathcal{P}$ represents our pattern matching ruleset, and Ψ is the action parsing function that translates natural language into discrete actions. For each action type, we maintain a comprehensive collection of synonymous expressions:

$$\begin{aligned} \mathcal{P}_{FORWARD} &= \{“move forward”, “proceed”, \dots\}, \\ \mathcal{P}_{TURN-LEFT} &= \{“turn left”, “rotate left”, \dots\}, \\ \mathcal{P}_{TURN-RIGHT} &= \{“turn right”, “rotate right”, \dots\}, \\ \mathcal{P}_{STOP} &= \{“stop”, “halt”, “wait”, \dots\}. \end{aligned}$$

This natural language interpretation approach presents several advantages. First, it eliminates the need for additional decoder networks, thereby simplifying the architecture. Second, it preserves the end-to-end nature of the system, enhancing efficiency. Finally, it offers robustness against variations in language expression. Pattern matching is achieved through case-insensitive regular expressions that accommodate various word arrangements, which enhances the system’s resilience to the inherent variations in VLM’s output formatting.

4 Experiments

4.1 Experimental Details

Dataset. We constructed a comprehensive dataset of approximately 1 Million training pairs using a hybrid collection strategy, which includes ground truth trajectories from the R2R and RxR datasets ($\approx 300k$ pairs from both), DAgger-collected data ($\approx 200k$ pairs from both), and specialized collision recovery samples ($\approx 25k$ pairs from both). This approach ensures diverse coverage of navigation scenarios and incorporates samples from general vision-language tasks for co-training.

For fair comparison, we trained on 500k pairs from R2R and evaluated Zero-Shot capability on RxR, in addition to conducting separate training and evaluation on the RxR dataset. Detailed strategies for dataset construction are provided in Appendix A.

Simulated Environments. We evaluate our *MapNav* agent using the VLN-CE benchmark in Habitat, which offers a continuous environment for navigation in reconstructed, photo-realistic indoor scenes. We focus on the val-unseen split of the R2R and RxR datasets in VLN-CE, which are two of the most recognized benchmarks in VLN, comprising 1,839 and 1,517 episodes, respectively.

Real-world environments. To evaluate the sim-to-real performance of our model, we designed diverse real-world experiments featuring 50 instructions across five environments: office, conference room, lecture hall, living room, and tea room. Simple instructions require basic actions like moving forward and turning, while semantic instructions involve navigating to specific objects (e.g., “move forward to the refrigerator and then turn right”).

Metrics. We utilize several widely used evaluation metrics for VLN tasks: Navigation Error (NE), Oracle Success rate (OS), Success Rate (SR), Success weighted Path Length (SPL), and normalized Dynamic Temporal Wrapping (nDTW). *SPL* is the primary metric, as it effectively reflects both navigation accuracy and efficiency.

Implementation Details. Our model is trained on a cluster server with 8 NVIDIA A100 GPUs for about 30 hours, totaling 240 GPU hours. We employ the LLaVA-Onevision (Li et al., 2024a) architecture, using Google’s SigLIP-so400M (Zhai et al., 2023) as the vision encoder, Qwen2-7B-Instruct (Wang et al., 2024) as the language model backbone, and Mask2Former (Cheng et al., 2022) for semantic segmentation during ASM generation.

Table 3: Memory consumption and average processing time comparison between Navid and MapNav (our method) across different numbers of navigation steps.

Method	Memory Consumption↓				Avg. Time↓
	1 Step	10 Steps	100 Steps	300 Steps	
Navid	0.92MB	9.2 MB	92 MB	276 MB	1.22s
MapNav (Ours)	0.17MB	0.17MB	0.17MB	0.17MB	0.25s

For detailed experimental details, see Appendix B.

4.2 Comparisons with SOTA Methods

Results in Simulated Environment. To evaluate cross-dataset performance, we firstly train solely on R2R samples and then assess its zero-shot performance on the RxR Val-Unseen split. To ensure fairness and consistency in our comparison, we first conducted experiments on single-frame RGB, where *MapNav (w/o ASM + Cur. RGB)* denotes MapNav without ASM, and NaVid (Cur. RGB) refers to NaVid without historical frames. As shown in Tab. 1, we achieved improvements of 14.1% in SR and 15.7% in SPL for R2R, and 6.9% in SR and 5.4% in SPL for RxR. After incorporating ASM, *MapNav (w/ ASM + Cur. RGB)*, which uses single-frame RGB images and ASM as input, shows improvements of 23.5% in SR and 26.5% in SPL for R2R compared to Navid (Cur. RGB).

Furthermore, our performance is competitive with Navid (All Frames), demonstrating that ASM effectively serves as a new historical representation method for large models, replacing traditional historical frames.

After incorporating just two historical RGB frames and training on both R2R and RxR, *MapNav (w/ ASM + Cur. RGB + 2 His.RGB)* outperforms state-of-the-art methods that use all historical frames. Specifically, compared to the previous SOTA method, Navid (All RGB Frames), we achieved improvements of 2.5% in SR and 1.3% in SPL for R2R, as well as 8.8% in SR and 6.5% in SPL for RxR.

Results in Real-World Environment. We utilize two metrics (*SR* and *NE*) to compare *MapNav(w/ ASM + Cur. RGB)* with WS-MGMAP (Chen et al., 2022) and Navid (Zhang et al., 2024a), which both use all the historical frames. As shown in Tab. 2, MapNav significantly outperforms both WS-MGMAP and Navid in SR and NE across simple and semantic instructions. Specifically, our method surpasses Navid under the semantic instruction settings in the lecture hall and living room, where we improve the *SR* by 30% in each setting. These results highlight the exceptional performance of

Table 4: Comparison of different map representation methods.

Method	VLN-CE R2R Val-Unseen			
	NE↓	OS↑	SR↑	SPL↑
MapNav (w/o Map)	7.26	41.2	27.3	23.2
MapNav (Original Map)	8.93	35.1	26.4	21.9
MapNav (Semantic Map)	6.56	43.2	29.1	24.5
MapNav (ASM)	5.22	50.3	36.5	34.3

Table 5: Comparison of different training dataset composition.

Method	VLN-CE R2R Val-Unseen			
	NE↓	OS↑	SR↑	SPL↑
MapNav (300k)	6.38	38.2	23.9	19.5
MapNav (300k+DAgger)	6.02	46.1	33.5	30.7
MapNav (300k+DAgger+RxR)	5.89	48.2	34.4	31.7
MapNav (300k+DAgger+Collision)	5.22	50.3	36.5	34.3

MapNav in real-world scenarios and validate the effectiveness of our proposed ASM.

4.3 Ablation Studies

Efficiency Analysis. For the efficiency analysis, we compare memory consumption and inference time between MapNav and Navid. As shown in Tab. 3, MapNav demonstrates significant improvements in both areas. Specifically, our semantic map-based method maintains a constant memory footprint of 0.17MB, regardless of trajectory length, while Navid’s frame-based approach scales linearly, reaching 276MB at 300 steps. This is because MapNav only stores and updates a compact ASM, whereas Navid accumulates all historical RGB observations. In terms of inference speed, **MapNav** reduces average processing time by 79.5% (from 1.22 seconds per step to 0.25 seconds). This improvement arises because MapNav only calculates features from the current RGB frame and ASM, avoiding the need to process all historical frames.

Effect of ASM. To validate the efficacy of our proposed ASM, we conducted a comparative analysis using multiple map variants. We evaluated three conditions: (1) the original top-down map from the simulator, (2) a semantic map with categorical information but no textual annotations, and (3) complete removal of map-based features. Quantitative results in Tab. 4 show a clear performance hierarchy: the RGB-only baseline performed the worst, with moderate improvements from the original and semantic map variants. The ASM outperformed all evaluated metrics, highlighting the benefits of semantic enrichment and the importance of textual annotations for spatial reasoning tasks.

Impact of different training dataset composi-




Figure 5: Comparison of MapNav using different numbers of historical RGB frames. Cur. RGB and His. RGB refer to methods using the current and historical RGB frames, respectively.

tion. To evaluate the impact of different training data compositions, we conducted a dataset ablation study with five configurations: (1) MapNav trained on 100K R2R samples (baseline), (2) 300K R2R samples, (3) DAgger-generated samples on R2R, (4) integration of DAgger and RxR datasets, and (5) all previous components plus collision-aware training. As shown in Tab. 5, increasing from 100K to 300K R2R samples resulted in substantial performance improvements. The integration of DAgger samples yielded the most significant gains, underscoring the importance of interactive learning. The RxR dataset provided modest improvements, particularly in diverse linguistic instructions. Our final configuration, with collision-aware training, achieved marginally better performance across metrics, setting a new performance benchmark. These results highlight that data diversity and interactive learning enhance model performance, with the greatest benefits stemming from ample base training data and DAgger augmentation.

Effect of Varying Numbers of Historical RGB Frames. We conducted a systematic ablation study to evaluate the effect of varying historical RGB frames. As shown in Fig. 5, the most significant improvement occurs with ASM in the single current RGB frame setting, enhancing SR from 27% to 36%, SPL from 23% to 34%, OS from 41% to 50%, and NE from 7.2m to 5.2m. This highlights the effectiveness of our ASM-based approach in capturing spatial information. While adding historical RGB frames (1, 2, and 4 His.RGB) results in gradual improvements, these gains are modest compared to the initial ASM integration. Notably, with just two historical RGB frames, our model achieves state-of-the-art performance (SR: 39%, SPL: 37%) compared to the baseline NaVid (Zhang et al., 2024a). The final configuration with four his-

torical RGB frames shows only slight gains (SR: 41%, SPL: 38%), indicating that the core advantage lies in the ASM-based representation rather than temporal accumulation. These findings validate our choice of ASM as a more efficient alternative to extensive historical RGB frame sequences, delivering superior performance while maintaining computational efficiency.

5 Conclusion

In this paper, we introduce MapNav, a novel method using Annotated Semantic Maps (ASMs) for Vision-and-Language Navigation. It achieves state-of-the-art performance while significantly reducing memory usage compared to traditional approaches. By converting RGB-D and pose data into enriched top-down maps with annotations, MapNav integrates efficiently with Vision-Language Models. Experiments show it delivers superior navigation while maintaining a constant memory footprint of 0.17MB, regardless of trajectory length.

Limitations

While MapNav demonstrates promising results, the semantic segmentation module may occasionally produce imprecise object labels under challenging conditions like occlusions or varying lighting. To address this aspect, we have identified several promising directions for future research. First, we aim to explore more advanced semantic understanding approaches that can effectively handle dynamic scenes and partial observations through multi-view scene understanding. Second, we plan to investigate methods to enhance real-world generalization by leveraging the extensive real-world image data in VLM pre-training. Additionally, we intend to extend MapNav to tackle more complex embodied AI tasks, such as interactive navigation and manipulation, which will require integrating object affordances and physical interaction capabilities into the ASM representation.

Acknowledgement

This work was supported by the National Science and Technology Major Project (No. 2022ZD0117800).

References

Dong An, Hanqing Wang, Wenguan Wang, Zun Wang, Yan Huang, Keji He, and Liang Wang. 2024. Etpnav:

Evolving topological planning for vision-language navigation in continuous environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niebner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. 2017. Matterport3d: Learning from rgb-d data in indoor environments. In *International Conference on 3D Vision*, pages 667–676.

Jiaqi Chen, Bingqian Lin, Xinmin Liu, Lin Ma, Xiaodan Liang, and Kwan-Yee K Wong. 2024. Affordances-oriented planning using foundation models for continuous vision-language navigation. *arXiv preprint arXiv:2407.05890*.

Peihao Chen, Dongyu Ji, Kunyang Lin, Runhao Zeng, Thomas Li, Mingkui Tan, and Chuang Gan. 2022. Weakly-supervised multi-granularity map learning for vision-and-language navigation. *Advances in Neural Information Processing Systems*, pages 38149–38161.

Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Jing Gu, Eliana Stefani, Qi Wu, Jesse Thomason, and Xin Wang. 2022. Vision-and-language navigation: A survey of tasks, methods, and future directions. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 7606–7623.

Peng Hao, Chaofan Zhang, Dingzhe Li, Xiaoge Cao, Xiaoshuai Hao, Shaowei Cui, and Shuo Wang. 2025a. Tla: Tactile-language-action model for contact-rich manipulation. *arXiv preprint arXiv:2503.08548*.

Weituo Hao, Chunyuan Li, Xiujun Li, Lawrence Carin, and Jianfeng Gao. 2020. Towards learning a generic agent for vision-and-language navigation via pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13137–13146.

Xiaoshuai Hao, Yunfeng Diao, Mengchuan Wei, Yifan Yang, Peng Hao, Rong Yin, Hui Zhang, Weiming Li, Shu Zhao, and Yu Liu. 2025b. Mapfusion: A novel bev feature fusion network for multi-modal map construction. *Information Fusion*, 119:103018.

Xiaoshuai Hao, Ruikai Li, Hui Zhang, Dingzhe Li, Rong Yin, Sangil Jung, Seung-In Park, ByungIn Yoo, Haimei Zhao, and Jing Zhang. 2024a. Mapdistill: Boosting efficient camera-based hd map construction via camera-lidar fusion model distillation. In *European Conference on Computer Vision*, pages 166–183. Springer.

Xiaoshuai Hao, Guanqun Liu, Yuting Zhao, Yuheng Ji, Mengchuan Wei, Haimei Zhao, Lingdong

- Kong, Rong Yin, and Yu Liu. 2025c. Msc-bench: Benchmarking and analyzing multi-sensor corruption for driving perception. *arXiv preprint arXiv:2501.01037*.
- Xiaoshuai Hao, Hui Zhang, Yifan Yang, Yi Zhou, Sangil Jung, Seung-In Park, and ByungIn Yoo. 2024b. Mb-fusion: A new multi-modal bev feature fusion method for hd map construction. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 15922–15928. IEEE.
- Yicong Hong, Zun Wang, Qi Wu, and Stephen Gould. 2022. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15439–15449.
- Yicong Hong, Yang Zhou, Ruiyi Zhang, Franck Dérioncourt, Trung Bui, Stephen Gould, and Hao Tan. 2023. Learning navigational visual representations with semantic map supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3055–3067.
- Yuheng Ji, Huajie Tan, Jiayu Shi, Xiaoshuai Hao, Yuan Zhang, Hengyuan Zhang, Pengwei Wang, Mengdi Zhao, Yao Mu, Pengju An, et al. 2025. Robobrain: A unified brain model for robotic manipulation from abstract to concrete. *arXiv preprint arXiv:2502.21257*.
- Glenn Jocher, Jing Qiu, and Ayush Chaurasia. 2023. [Ultralytics YOLO](#).
- Jacob Krantz, Aaron Gokaslan, Dhruv Batra, Stefan Lee, and Oleksandr Maksymets. 2021. Waypoint models for instruction-guided navigation in continuous environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15162–15171.
- Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. 2020. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *European Conference on Computer Vision*, pages 104–120.
- Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. 2020a. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 4392–4412.
- Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. 2020b. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 4392–4412.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. 2024a. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Dingzhe Li, Yixiang Jin, Yuhao Sun, Hongze Yu, Jun Shi, Xiaoshuai Hao, Peng Hao, Huaping Liu, Fuchun Sun, Jianwei Zhang, et al. 2024b. What foundation models can bring for robot learning in manipulation: A survey. *arXiv preprint arXiv:2404.18201*.
- Bingqian Lin, Yunshuang Nie, Ziming Wei, Jiaqi Chen, Shikui Ma, Jianhua Han, Hang Xu, Xiaojun Chang, and Xiaodan Liang. 2024. Navcot: Boosting llm-based vision-and-language navigation via learning disentangled reasoning. *arXiv preprint arXiv:2403.07376*.
- Rui Liu, Wenguan Wang, and Yi Yang. 2024. Volumetric environment representation for vision-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16317–16328.
- Yuxing Long, Wenzhe Cai, Hongcheng Wang, Guanqi Zhan, and Hao Dong. 2024a. Instructnav: Zero-shot system for generic instruction navigation in unexplored environment. *arXiv preprint arXiv:2406.04882*.
- Yuxing Long, Xiaoqi Li, Wenzhe Cai, and Hao Dong. 2024b. Discuss before moving: Visual language navigation via multi-expert discussions. In *IEEE International Conference on Robotics and Automation*, pages 17380–17387.
- Sang-Min Park and Young-Gab Kim. 2023. Visual language navigation: A survey and open challenges. *Artificial Intelligence Review*, pages 365–427.
- Yuankai Qi, Zizheng Pan, Shengping Zhang, Anton van den Hengel, and Qi Wu. 2020. Object-and-action aware model for visual language navigation. In *European Conference on Computer Vision*, pages 303–317.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the international conference on artificial intelligence and statistics*, pages 627–635.
- Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. 2019. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9339–9347.
- Dhruv Shah, Błażej Osioński, Sergey Levine, et al. 2023. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conference on Robot Learning*, pages 492–504.
- Yingbo Tang, Shuaike Zhang, Xiaoshuai Hao, Pengwei Wang, Jianlong Wu, Zhongyuan Wang, and Shanghang Zhang. 2025. Affordgrasp: In-context affordance reasoning for open-vocabulary task-oriented grasping in clutter. *arXiv preprint arXiv:2503.00778*.

- Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. 2020. Vision-and-dialog navigation. In *Conference on Robot Learning*, pages 394–406.
- Arun Balajee Vasudevan, Dengxin Dai, and Luc Van Gool. 2021. Talk2nav: Long-range vision-and-language navigation with dual attention and spatial memory. *International Journal of Computer Vision*, pages 246–266.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Zihan Wang, Xiangyang Li, Jiahao Yang, Yeqi Liu, and Shuqiang Jiang. 2023. Gridmm: Grid memory map for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15625–15636.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, pages 24824–24837.
- Pengying Wu, Yao Mu, Bingxian Wu, Yi Hou, Ji Ma, Shanghang Zhang, and Chang Liu. 2024. Voronav: Voronoi-based zero-shot object navigation with large language model. *arXiv preprint arXiv:2401.02695*.
- Siyang Wu, Xueyang Fu, Feng Wu, and Zheng-Jun Zha. 2022. Cross-modal semantic alignment pre-training for vision-and-language navigation. In *Proceedings of the ACM International Conference on Multimedia*, pages 4233–4241.
- Yujie Wu, Huaihai Lyu, Yingbo Tang, Lingfeng Zhang, Zhihui Zhang, Wei Zhou, and Siqi Hao. 2025. Evaluating gpt-4o’s embodied intelligence: A comprehensive empirical study. *TechRxiv preprint techrxiv.174495686.69962588/v1*.
- Naoki Yokoyama, Sehoon Ha, Dhruv Batra, Jiuguang Wang, and Bernadette Bucher. 2024. Vlfm: Vision-language frontier maps for zero-shot semantic navigation. In *IEEE International Conference on Robotics and Automation*, pages 42–48.
- Bangguo Yu, Hamidreza Kasaei, and Ming Cao. 2023. L3mvm: Leveraging large language models for visual target navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3554–3560.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986.
- Zhaohuan Zhan, Lisha Yu, Sijie Yu, and Guang Tan. 2024. Mc-gpt: Empowering vision-and-language navigation with memory map and reasoning chains. *arXiv preprint arXiv:2405.10620*.
- Chaofan Zhang, Peng Hao, Xiaoge Cao, Xiaoshuai Hao, Shaowei Cui, and Shuo Wang. 2025. Vtla: Vision-tactile-language-action model with preference learning for insertion manipulation. *arXiv preprint arXiv:2505.09577*.
- Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. 2023. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*.
- Jiazhao Zhang, Kunyu Wang, Rongtao Xu, Gengze Zhou, Yicong Hong, Xiaomeng Fang, Qi Wu, Zhizheng Zhang, and He Wang. 2024a. Navid: Video-based vlm plans the next step for vision-and-language navigation. In *Proceedings of Robotics: Science and Systems*.
- Lingfeng Zhang, Hao Wang, Erjia Xiao, Xinyao Zhang, Qiang Zhang, Zixuan Jiang, and Renjing Xu. 2024b. Multi-floor zero-shot object navigation policy. *arXiv preprint arXiv:2409.10906*.
- Lingfeng Zhang, Qiang Zhang, Hao Wang, Erjia Xiao, Zixuan Jiang, Honglei Chen, and Renjing Xu. 2024c. Trihelper: Zero-shot object navigation with dynamic assistance. *arXiv preprint arXiv:2403.15223*.
- Duo Zheng, Shijia Huang, Lin Zhao, Yiwu Zhong, and Liwei Wang. 2024. Towards learning a generalist model for embodied navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13624–13634.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, pages 46595–46623.
- Gengze Zhou, Yicong Hong, and Qi Wu. 2024. Navgpt: Explicit reasoning in vision-and-language navigation with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7641–7649.

Appendix

This supplementary material provides additional details on the proposed method and experimental results that could not be included in the main manuscript due to page limitations.

Specifically, this appendix is organized as follows:

- Sec. A provides a comprehensive overview of the dataset construction process.
- Sec. B outlines the detailed training procedures used for our models.
- Sec. C describes our real-world MapNav robot setup along with implementation specifics.
- Sec. D presents the qualitative results in both the simulator and real-world environments.
- Sec. E showcases additional experimental results that further validate our findings.
- Sec. F offers further qualitative results specifically from the simulator.
- Sec. G provides additional qualitative results obtained from real-world scenarios.
- Sec. H analyzes VLM attention visualization across different map representations.

A Dataset Construction

Motivation and Overview. Training data for VLN tasks faces two primary challenges: limited diversity and insufficient scale. To address these challenges, we propose a comprehensive data collection strategy that combines three complementary approaches: expert demonstrations based on ground truth trajectories, interactive learning through DAgger, and specialized collision recovery data. This hybrid strategy maximizes the utility of the available training data and significantly enhances the model’s generalization capabilities across diverse scenarios.

Phase I: Expert Demonstration Collection. Initially, we collected around 300k high-quality training pairs from both R2R and RxR datasets. Each step-wise pair consists of three components: the current RGB observation frame, the corresponding Annotated Semanti Map (ASM), and the associated action. These pairs serve as the foundation for our initial supervised fine-tuning (SFT) process for the

VLM, providing the model with expert demonstration data from diverse indoor environments.

Phase II: Interactive Learning via DAgger. Following the Dataset Aggregation (DAgger) methodology (Ross et al., 2011), we deployed our pre-trained model to collect additional trajectory data, amassing approximately 200k new training pairs from both R2R and RxR environments. This phase was crucial for two reasons: First, combining expert trajectories with DAgger-collected data creates a more robust training dataset that reflects both ideal navigation behavior and realistic agent interactions, including potential navigation errors and recovery strategies. Second, this hybrid approach bridges the gap between training and deployment conditions, enhancing the model’s ability to tackle novel scenarios and unexpected situations during navigation.

Phase III: Specialized Collision Recovery Data. To overcome the limitation of expert trajectories rarely including collision scenarios, we supplemented our dataset with specialized collision recovery data. During the DAgger collection strategy, we specifically gathered instances where the agent encountered obstacles and needed recovery actions. This additional dataset, consisting of approximately 25k step-wise training pairs from both R2R and RxR, captures the agent’s collision recovery behavior through rotational movements and obstacle avoidance strategies. This collision recovery subset was crucial for enhancing the model’s robustness in deployment scenarios. While ground truth trajectories provide optimal navigation patterns and DAgger data captures general interaction scenarios, the collision recovery dataset specifically addresses edge cases where the agent must navigate obstacles. During the final fine-tuning phase, we integrated this collision recovery data with the primary dataset, enabling the model to learn effective obstacle avoidance and recovery strategies.

Overall Dataset Structure and Scale. The complete dataset composition can be formalized as:

$$\begin{aligned}\mathcal{D}_{total_R2R} &= \mathcal{D}_{total_RxR} \\ &= \mathcal{D}_{GT} \cup \mathcal{D}_{DAgger} \cup \mathcal{D}_{Collision},\end{aligned}\tag{5}$$

where $|\mathcal{D}_{GT}| \approx |\mathcal{D}_{DAgger}| \approx 200k$, and $|\mathcal{D}_{collision_R2R}| \approx |\mathcal{D}_{collision_RxR}| \approx 25k$ for both R2R and RxR environments respectively.

This comprehensive dataset collection strategy significantly improves the agent’s ability to handle unforeseen obstacles and navigate challenging

situations autonomously. The enhanced training dataset, now totaling approximately **1 Million** training pairs, offers a more complete representation of navigation scenarios, encompassing both optimal pathfinding and practical recovery strategies. This hybrid data collection strategy substantially contributes to the model’s resilience and adaptability across diverse indoor environments and navigation challenges.

Incorporating General Vision-Language Datasets. To enhance our model’s visual understanding and general reasoning capabilities, we implement a co-training strategy that leverages both navigation-specific data and general-purpose vision-language datasets. Following the approach similar to LLaVA-OneVision (Li et al., 2024a), we incorporate a diverse set of high-quality visual understanding datasets alongside our navigation training data. Specifically, for R2R task, the co-training process utilizes approximately 600k samples from general vision-language tasks, which complement our 500k navigation training samples. The auxiliary dataset encompasses diverse visual understanding tasks including Visual Question Answering (VQA) samples for enhanced visual reasoning, multi-image reasoning tasks for improved cross-image understanding and relationship inference, and video-based tasks for strengthening temporal reasoning and dynamic scene comprehension. For more details, please refer to (Li et al., 2024a).

B Details of Training

Model Setting. Our model builds upon the LLaVA-Onevision (Li et al., 2024a) framework and consists of three main components: the visual encoder, projector, and large language model (LLM). For the visual encoder, we employed Google’s SigLIP-so400m-patch14-384 (Zhai et al., 2023), which processes input images at a resolution of 384x384 using patches of size 14x14. The SigLIP model serves as our vision backbone, converting input images into visual embeddings with a hidden dimension of 1152. The projector is implemented as a two-layer MLP with GELU activation, which maps the visual features to the language model’s embedding space. For the language model, we utilized Qwen2-7B-Instruct (Wang et al., 2024) as our backbone, which features 28 transformer layers with 28 attention heads and a hidden size of 3584. The model supports a context length of up




Figure 6: The real-world MapNav robot setup.

to 32,768 tokens and incorporates sliding window attention with a window size of 131,072 tokens.

Training Setting. We conducted our training on 8 NVIDIA A100 GPUs for approximately 30 hours, totaling 240 GPU hours ($\approx 500k$ step-wise samples). During the fine-tuning process, we froze the vision encoder and only fine-tuned the multi-modal projector and language model components for one epoch with a learning rate of $1e-6$. We utilized bfloat16 precision for training efficiency. The model processes images using bilinear spatial pooling and selects features from the penultimate layer of the vision tower. For multi-modal integration, we employed patch-level features without using image patch tokens or explicit image start/end tokens. The projector adopts a spatial unpadding patch merge strategy to handle varying image sizes effectively.

C Real-World Robot Setup

As shown in Fig. 6, our real-world experiments were conducted using a Unitree Go2 Edu quadrupedal robot equipped with an Intel RealSense D435i depth camera mounted 40cm above the robot. The system streams RGB-D images from a D435i camera to a server powered by NVIDIA A100 GPUs, where our MapNav agent processes the observations to generate our ASMs. Based on the generated ASM, the system infers executable actions which are then transmitted back to the Go2 for execution. Notably, our system demonstrated robust sim-to-real transfer capabilities, maintaining accurate ASM generation and high task performance despite real-world challenges such as depth measurement uncertainties, pose estimation errors, and environmental variations.

D Qualitative Results.

Fig. 7 and Fig. 8 illustrate the performance of our MapNav agent in simulated and real-world environments. In the simulator, the agent with ASM




Figure 7: Visualization results of *MapNav* in the simulator.




Figure 8: Visualization results of *MapNav* in the real-world.

successfully identifies the shortest path while following complex instructions involving multiple waypoints. In contrast, without ASM, the agent struggles to find the correct path, underscoring ASM’s importance in semantic understanding and path planning. In real-world tests, the agent effectively executes simple navigation instructions and excels at complex tasks involving semantic landmarks. The ASM visualization reveals its ability to adapt paths in real-time, confirming robust sim-to-real transfer. These results demonstrate that our approach maintains strong performance across both environments, particularly in tasks requiring semantic understanding and adaptive navigation.

E More Experiments

Ablation on Different Input Modals. To systematically investigate the contribution of each input modality, we conducted an extensive ablation study on the input representations. The experimental configurations encompassed three variants: (1) RGB-only input, serving as the base-

line condition, (2) RGB augmented with depth information (RGB+Depth), and (3) our complete model utilizing the ASM alongside RGB inputs. Quantitative results, as presented in Tab . 6, reveal an interesting pattern across these configurations. While the RGB-only baseline achieved moderate performance, the addition of depth information (RGB+Depth) led to a notable performance degradation. This counterintuitive result can be attributed to the Vision-Language Model’s inherent limitations in processing depth modalities, as these models are primarily trained on RGB images and natural language. In contrast, the integration of our proposed ASM demonstrated substantial performance gains across all evaluation metrics, significantly outperforming both baseline configurations. These results not only validate the effectiveness of our ASM approach but also highlight the importance of selecting input modalities that align with the pre-trained model’s capabilities.

Ablation Experiment on Different Semantic Segmentation Modules. We conducted another ab-

Table 6: Comparison of different input modals.

Method	VLN-CE R2R Val-Unseen			
	NE↓	OS↑	SR↑	SPL↑
MapNav (Only RGB)	7.26	41.2	27.1	23.5
MapNav (RGB+Depth)	8.82	35.6	23.1	19.9
MapNav (Annotated Semantic Map)	5.22	50.3	36.5	34.3

Table 7: Comparison of different semantic segmentation modules’ performance on R2R Val-Unseen split.

VLMs	VLN-CE R2R Val-Unseen			
	NE↓	OS↑	SR↑	SPL↑
YOLOv8 (Jocher et al., 2023)	6.43	45.2	31.5	29.6
MobileSAM (Zhang et al., 2023)	6.02	48.9	34.8	31.6
Mask2Former (Our Use) (Cheng et al., 2022)	5.22	50.3	36.5	34.3

lation study to evaluate the impact of different semantic segmentation modules on our ASM generation process, with results presented in Tab. 7. We compared three state-of-the-art semantic segmentation approaches: YOLOv8 (Jocher et al., 2023), MobileSAM (Zhang et al., 2023), and our chosen Mask2Former (Cheng et al., 2022). The experiments were conducted on the R2R Val-Unseen split to assess generalization capability. Results show that Mask2Former achieves the best performance across all metrics, with a Navigation Error (NE) of 5.22m, Success Rate (SR) of 36.5%, and SPL of 34%, outperforming both YOLOv8 and MobileSAM. The superior performance of Mask2Former can be attributed to its more precise segmentation boundaries and better handling of complex indoor scenes. While MobileSAM shows moderate performance with an SR of 34.8% and SPL of 31.6%, YOLOv8 exhibits relatively lower performance with an SR of 31.5% and SPL of 29.6%. These results suggest that the quality of semantic segmentation significantly impacts the overall navigation performance, highlighting the importance of accurate semantic objects segmentation in ASM generation.

F More Simulated Demos

We conducted extensive visualization experiments to demonstrate the effectiveness of our approach, selecting representative successful cases from both R2R and RxR datasets. Specifically, as shown in Fig. 11–Fig. 16, we visualized 24 distinct navigation trajectories across six pages, with each case highlighting the agent’s navigation process and corresponding Abstract Semantic Maps (ASM). These visualizations showcase the robust generalization capability of our ASM-based approach across diverse environments and navigation scenarios. The

demonstrated cases showcase a variety of room layouts, navigation objectives, and complex multi-step instructions, all of which our agent successfully interpreted and executed. This performance across diverse scenarios validates our approach’s spatial understanding capabilities, as reflected in the agent’s ability to generate accurate ASMs and perform appropriate navigation actions. Additionally, these examples highlight our ASMs’ strong memory representation abilities, illustrating how the agent maintains and updates its understanding of the environment throughout the navigation process. By integrating new observations with previously stored spatial information, the agent can make informed navigation decisions.

G More Real-World Demos

To extensively validate our approach in real-world settings, we conducted comprehensive experiments and selected six representative episodes for visualization as shown in Fig. 17 and Fig. 18. These episodes span five distinct real-world environments: office, meeting room, lecture hall, tea room, and living room, encompassing both simple and semantic instructions. The successful generation of ASMs in these diverse real-world scenarios and the agent’s ability to complete assigned tasks demonstrate our system’s effective sim-to-real transfer capabilities and the practical applicability of ASM in real-world navigation. Additionally, we conducted an experiment in a living room setting to test our VLM’s semantic understanding capabilities. When presented with sofas of different colors, the agent was instructed to locate and navigate to the gray sofa. The successful execution of this task validates our model’s advanced semantic understanding and its ability to differentiate between visually similar objects based on specific attributes.

H VLM Attention Visualization Analysis

To quantitatively evaluate our Annotated Semantic Map (ASM) representation, we conducted attention visualization analysis using VLM Visualizer¹ across different map representations. Using Vicuna-13B (Zheng et al., 2023) as our base model, we generated attention heatmaps to examine how the model attends to different regions of the input maps. As shown in Fig. 9 and Fig. 10, the attention visualization reveals that when processing ASMs, the model exhibits significantly stronger

¹<https://github.com/zjysteven/VLM-Visualizer>




Figure 9: **Visualization of VLM Attention Across Different Map Representations.** A comparison of different map representations showing that while Semantic Map exhibits sparse attention patterns without convergence on semantic objects, our ASM successfully leverages textual labels to guide attention towards semantic objects, as evidenced by concentrated attention distributions and the VLM’s responses.

attention alignment with semantically meaningful regions, showing sharp attention peaks (> 0.8) precisely aligned with labeled objects and navigation-relevant areas. Moreover, our attention maps demonstrate that the red trajectory lines in ASM receive substantial attention focus, providing structured navigational cues that complement the semantic understanding. In contrast, attention patterns remain notably diffuse when processing the original top-down map (peak attention < 0.3) and semantic map (peak attention < 0.4), suggesting limited semantic understanding. This quantitative difference in attention patterns is further reflected in the model’s descriptive outputs - while basic map descriptions indicate only geometric recognition (“a black and white map with a blue line”) for the original format, ASM enables sophisticated spatial-semantic understanding with precise object localization (e.g., “bed in the upper left part”). These visualization results demonstrate that ASM’s explicit textual annotations and structured trajectory representation successfully guide the model’s at-

tention to both semantic objects and navigation paths, enabling effective grounding of navigational features through the model’s pre-trained language understanding capabilities.




Figure 10: Additional Visualizations of VLM Attention Across Different Map Representations.




Figure 11: (1/6) Simulator demo results visualization.




Figure 12: (2/6) Simulator demo results visualization.




Move through the kitchen with the square table to your right. Continue moving forward to the staircase just after the refrigerator. Walk up the stairs and stop 18 inches before the end of the stair railing to your right.

Observations
ASMs



Walk past the dining room table, through the sitting room and out the sliding glass doors. Wait on the patio near the table and chairs.

Observations
ASMs



Take a left and take a right at the doll hutch. Walk into the bedroom and take a left. Walk into the bathroom and wait in front of the sink.

Observations
ASMs



Go up the steps. Turn right and go straight until you get to the phone and vase with flower. Then turn right. Wait in that room.

Observations
ASMs






Figure 13: (3/6) Simulator demo results visualization.


With the window behind you, walk straight past the table and kitchen area and through the doorway. Keep going until you get to a separate sitting room on the left with blue walls and white couches. Walk into the room and stop once you step on the carpet.



Continue up the stairs, walk towards the sitting area, go straight passed the table on the left. Stop by the double doors.



Walk towards the ovens and take a left. Walk towards the fireplace and enter the dining room to the right of the fireplace. Stop in front of the white chair.



Walk down the hallway and walk into the kitchen and take a left at the island counter into the hallway and take a right into the dining room and walk through the dining room into the living room and stop nest to the edge of the couch behind the end table.




Figure 14: (4/6) Simulator demo results visualization.



Figure 15: (5/6) Simulator demo results visualization.



Walk past the pool table and turn right. Walk into the next room and turn left. At the door to the bathroom, turn left, step forward to wait in the doorway to the wine cellar.

Observations
ASMs




Steps



Turn to the left, walk between the massage tables. Go out of the door, and straight just a little way and turn right. You'll see a big bathtub-like thing. You'll stop right there.

Observations
ASMs




Steps



Turn around and walk into the dining room. Take a right. Take a left after the table and wait in the doorway.

Observations
ASMs




Steps



Walk around the bed and chairs and enter the bathroom on your left. Stop near the shower.

Observations
ASMs



Steps

Figure 16: (6/6) Simulator demo results visualization.




Figure 17: (1/2) Real-world demo results visualization.




Figure 18: (2/2) Real-world demo results visualization.