

Boosting Zero-Shot VLN via Abstract Obstacle Map-Based Waypoint Prediction with TopoGraph-and-VisitInfo-Aware Prompting

Boqi Li^{*}, Siyuan Li^{*}, Weiyi Wang, Anran Li, Zhong Cao, and Henry X. Liu[†], *Fellow, IEEE*

Abstract—With the rapid progress of foundation models and robotics, vision-language navigation (VLN) has emerged as a key task for embodied agents with broad practical applications. We address VLN in continuous environments, a particularly challenging setting where an agent must jointly interpret natural language instructions, perceive its surroundings, and plan low-level actions. We propose a zero-shot framework that integrates a simplified yet effective waypoint predictor with a multimodal large language model (MLLM). The predictor operates on an abstract obstacle map, producing linearly reachable waypoints, which are incorporated into a dynamically updated topological graph with explicit visitation records. The graph and visitation information are encoded into the prompt, enabling reasoning over both spatial structure and exploration history to encourage exploration and equip MLLM with local path planning for error correction. Extensive experiments on R2R-CE and RxR-CE show that our method achieves state-of-the-art zero-shot performance, with success rates of 41% and 36%, respectively, outperforming prior state-of-the-art methods. The source code is available at: <https://github.com/michigan-traffic-lab/OMAP-VLN>.

I. INTRODUCTION

Vision-language navigation (VLN) is the task in which an embodied agent follows a natural language instruction to navigate through an environment and reach a specified destination. It represents an important capability for embodied AI, with applications ranging from search and rescue to autonomous navigation and daily human–robot interaction. However, VLN remains highly challenging: the agent must not only comprehend natural language instructions but also ground them in its surrounding environment [1]–[3]. Existing approaches [4]–[6] typically rely on RGB-D sensory inputs, which demand both a detailed understanding of complex visual scenes and robust reasoning over the task context.

Early research on VLN has largely focused on discrete environment settings, where a navigation graph is predefined and the agent’s movement is restricted to a finite set of checkpoints [1], [2], [7]. More recently, it has been extended to continuous environments, where the agent can move freely [3], [8]. This setting is considerably more challenging and has motivated a line of work that employs waypoint prediction models [9]: candidate waypoints are first generated in

the continuous space, and a navigation policy then selects one to pursue. This decomposition is intended to ease the learning process by separating spatial grounding from long-horizon reasoning.

Despite their promise, existing waypoint prediction approaches exhibit drawbacks. First, the choice of input representations remains underexplored. The original waypoint model relied on RGB and depth images; however, subsequent work found that RGB inputs can sometimes degrade performance, while others introduced increasingly complex architectures for RGB-D processing, further increasing model complexity [5], [10]. Second, predicted waypoints are not guaranteed to be reachable from the agent’s current position and may be obstructed by intervening obstacles.

Our primary goal is to investigate whether navigation-relevant spatial structure can be represented in a more abstract and lightweight form. Therefore, rather than designing waypoint prediction models with increasingly complex sensory inputs and neural architectures, we propose a simplified waypoint prediction model that proves rather effective. Our key insight is that providing raw RGB and depth images introduces large amounts of irrelevant information, which can obscure the underlying spatial relationships and hinder generalization.

To address this, we abstract the input representation: depth images are first processed into an obstacle map centered on the agent, which is then used as the sole input to the waypoint predictor. This abstraction encourages the model to focus on the dynamic patterns of obstacle distribution, enabling it to better identify critical regions where waypoints should be generated. Empirically, we find that our approach produces waypoints that are both more feasible and reliable than those generated by prior methods.

We integrate our proposed waypoint model with a zero-shot navigator based on a multimodal large language model (MLLM) [11]. Zero-shot VLN with large language models (LLMs) has attracted increasing attention in recent years [4], [12]–[15], as purely learning-based navigation models still suffer from data scarcity and struggle to generalize to unseen environments. By leveraging the broad knowledge encoded in LLMs, zero-shot approaches enable agents to better interpret both natural language instructions and perceptual inputs.

Recent works in VLN-CE have explored textual and visual prompt design for LLM-based navigators [10], [13], [15], incorporating different forms of contextual information. However, most of these approaches overlook the structure of the topological graph, causing the LLM agent to quickly lose track of which waypoints have already been explored. To

This work was supported by the DARPA TIAMAT Challenge. (HR0011-24-9-0429)

Boqi Li, Zhong Cao, Henry X. Liu are with the Department of Civil and Environmental Engineering; Siyuan Li and Weiyi Wang are with the Department of Computer Science and Engineering; Anran Li is with the Department of Robotics, University of Michigan, Ann Arbor, MI, USA. (e-mails: [boqili, syuanlee, wweiyi, anranli, zhcao, henryliu]@umich.edu)

^{*}: These authors contributed equally to this work.

[†]: Corresponding author.

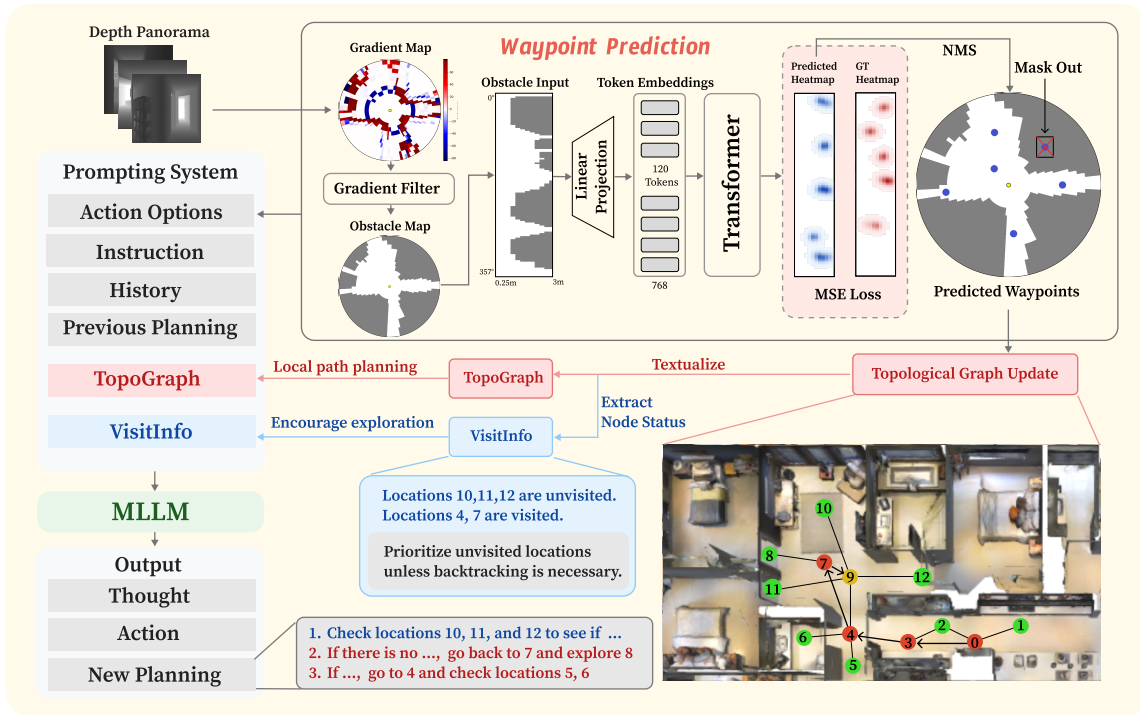


Fig. 1. Overview of the proposed VLN framework. Compared with prior VLN pipelines that rely on high-dimensional RGB-D encoders, our approach introduces an abstraction layer (obstacle map) for waypoint prediction and explicitly encodes spatial connectivity and visitation history for MLLM reasoning. Depth images are processed into gradients and filtered to form an obstacle map, which serves as input to the waypoint predictor. Generated waypoints are dynamically added to a topological graph and textualized as TopoGraph and VisitInfo. VisitInfo helps MLLM prioritize exploration and TopoGraph equips MLLM with local path planning (backtracking one/multiple steps and re-exploration) for error correction.

address this limitation, we propose a dynamically updated topological graph that explicitly encodes node visitation status. This graph is incorporated into the prompt design, providing the LLM with structured knowledge of explored and unexplored regions, alongside perceptual inputs, thereby enabling more consistent and informed reasoning about the agent’s navigation choices.

We evaluate our proposed waypoint prediction + MLLM navigation framework in the R2R-CE [8] and RxR-CE [7] datasets, achieving 41% and 36% success rates and outperforming recent zero-shot methods. Our contributions are threefold:

- Abstract obstacle-map-based waypoint prediction: A lightweight predictor using obstacle maps as input, improving waypoint feasibility and reachability.
- TopoGraph-and-VisitInfo aware prompting: A dynamically updated topological graph with visited-node tracking, incorporated into prompts to help the MLLM reason over spatial structure and exploration history.
- Integrated zero-shot VLN framework: A unified system combining our waypoint predictor with an MLLM navigator, achieving state-of-the-art results on both R2R-CE and RxR-CE datasets.

II. RELATED WORK

A. Vision-Language Navigation

Early research in VLN focused on discrete settings. Research on discrete VLN spans imitation and reinforcement

learning [16]–[18], pretraining strategies [19], [20], and more recently, zero-shot methods with LLMs [4], [12], [21]–[23], collectively laying the foundation for benchmark evaluation and methodological advances in VLN.

Despite progress in discrete VLN, models in this setting remain limited, as agents are restricted to predefined navigation graphs that oversimplify the task. To bridge the gap toward real-world navigation, Savva et al. introduced Habitat [3], a flexible 3D simulator enabling continuous movement and Krantz et al. proposed the more challenging VLN-CE [8], where agents must also handle low-level control. Numerous approaches have been explored, including waypoint prediction models [5], [24], [25], hierarchical planning with subgoals [26], and semantic/spatial map representations [27], [28]. Recent work has also explored safety-aware reasoning mechanisms for VLN [29].

B. Waypoint Prediction

Waypoint prediction is a widely adopted strategy for VLN-CE. Krantz et al. [9] first introduced this framework, and several subsequent works have extended it [5], [10], [24], [30]. Typically, waypoint models are pretrained on navigation graphs from the training environments and then kept fixed while the downstream planner is learned. While waypoint prediction has helped close the performance gap in VLN-CE, it still has key limitations. Pretrained models are often trained on limited navigation graphs [24], restricting generalization; they typically generate only local waypoints and may even

predict infeasible ones that collide with obstacles [5]. Moreover, the choice of input representation remains unsettled: An et al. [5] showed that RGB inputs can harm generalization by introducing unnecessary semantics, whereas Shi et al. [10] reported improved performance by using stronger RGB-D encoders with an obstacle-awareness loss.

C. Foundation models for VLN

Foundation models [11], [31]–[33], with their strong reasoning and generalization abilities, are particularly well-suited for embodied tasks such as VLN. Existing works have integrated them into navigation pipelines through either zero-shot approaches [4], [10], [12]–[15], which directly leverage LLMs for decision-making, or fine-tuning strategies [34]–[36], which adapt open-source vision language models (VLMs) to navigation-specific objectives.

As for the recent VLN-CE, representative approaches include A2Nav [13], which decomposes instructions into action-specific subtasks, and waypoint-based methods such as OpenNav [14], which leverages open-source LLMs/VLMs, and SmartWay [10], which integrates a waypoint predictor with an MLLM capable of backtracking. AO-Planner [15] takes a fully zero-shot route, generating low-level actions from affordance estimation via Grounded SAM [37] and refining candidate paths with an LLM planner. Beyond direct action generation, LLMs have also been explored as value estimators, e.g., InstructNav [38] constructs value maps from textualized observations, while CA-Nav [6] enforces sub-instruction constraints to guide navigation.

III. METHOD

A. Problem Formulation

We consider the VLN-CE task within the Habitat simulator, where an embodied agent is equipped with an odometry sensor and an RGB-D camera operating in a 3D mesh environment. At each time step t , the agent has access to its current position (x_t, y_t, z_t) and a panoramic observation of its surroundings, denoted as $O_t = \{(O_n^{\text{rgb}}, O_n^{\text{depth}})\}_{n=1}^{12}$, which consists of twelve RGB and depth image pairs captured at uniformly spaced headings $\{0^\circ, 30^\circ, \dots, 330^\circ\}$. The agent’s action space includes low-level controls: `Move-forward(0.25m)`, `Turn-left/right(15°)`, and `Stop`. At the beginning of an episode, the agent is provided with a natural language instruction $I = \{w_1, w_2, \dots, w_N\}$, which implicitly specifies a target location $(x_{\text{target}}, y_{\text{target}}, z_{\text{target}})$. The goal of the agent is to interpret the instruction and execute a sequence of low-level actions that navigate it successfully to the destination.

B. Waypoint Prediction Model

1) *The Abstraction of Input Space*: Our idea is that both RGB and depth images encode overly complex, high-dimensional information that is not well aligned with the waypoint space, which is fundamentally a 2D radial surface. In particular, RGB inputs often carry semantic-level details that are irrelevant to spatial traversability, while depth images

also introduce unnecessary complexity. Importantly, this abstraction is applied only to the waypoint prediction stage. The MLLM planner still receives RGB observations for semantic grounding of language instructions, while the obstacle map provides a compact representation for geometric waypoint generation.

In the original waypoint prediction formulation, the space around the agent is discretized into angular intervals of 3° and radial steps of 0.25m, extending up to 3.0m in each direction. This produces a grid of candidate waypoints represented as a heatmap of size 120×12 , where each entry corresponds to the likelihood of a waypoint at that location.

To better align the input and output spaces, we propose to construct an obstacle map with the same resolution as the waypoint heatmap. Each cell in this binary map is marked as 1 if occupied by an obstacle and 0 if free space. By using this obstacle map as the sole input to the waypoint predictor, we encourage the model to focus exclusively on spatial traversability and to learn the geometric patterns of obstacles and free space surrounding the agent in 2D, ultimately producing more feasible waypoint predictions.

2) *Obstacle Map Construction*: The obstacle map is not directly available and must be constructed from the agent’s depth observations. As is shown in Fig. 1, we first convert panoramic depth images into a 3D point cloud and project them onto the discretized 120×12 radial 2D grid used for waypoint prediction. For each grid cell, we estimate local elevation by computing the maximum height of its points, then calculate elevation gradients along each radial direction, treating it as if the agent were “climbing” along that path. A cell is marked as an obstacle if its gradient exceeds a threshold. Unlike prior work [30] that applies a fixed height threshold from the agent’s footpoint, which often misclassifies stairs or slopes as obstacles, our gradient-based method captures such variations more accurately, yielding obstacle maps that better reflect true traversability.

3) *The Simplified Model Structure*: Since our input and output spaces are aligned, we are able to eliminate the complex model architectures used in previous waypoint predictors. Instead, we adopt a lightweight self-attention model with two layers of transformer encoders and a classification head that outputs waypoint logits over the discretized grid. The model is trained using ground-truth waypoint heatmaps as supervision. This lightweight design not only reduces architectural complexity but also improves waypoint feasibility, ensuring that predicted waypoints are more likely to be located in open and traversable space.

4) *Masking for linear reachability*: To ensure that generated waypoints are truly feasible, they must be compatible with the agent’s low-level motion pattern, which follows a “turn-then-move” strategy. This requires that waypoints be linearly reachable: even if a waypoint lies in open space, it is undesirable if it is located behind an intervening obstacle. To address this, we apply a logit-masking procedure, where for each direction all locations beyond the first detected obstacle are masked out in the waypoint heatmap. From the resulting masked heatmap, we then apply non-maximum suppression

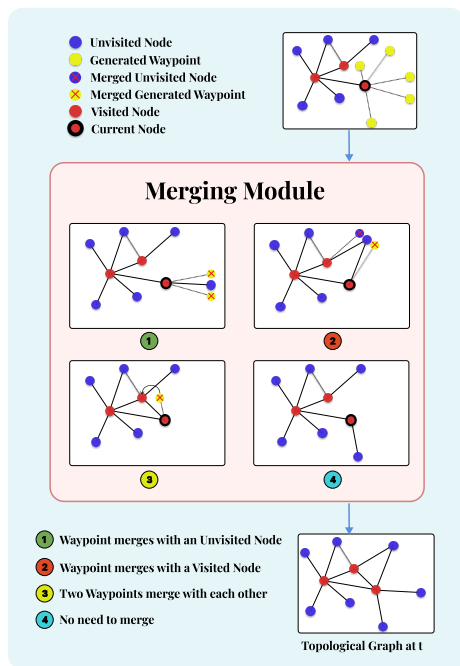


Fig. 2. To maintain the topological graph in a simple, effective, and consistent form, the merging module merges a generated waypoint with an existing node, or merges two generated waypoints, if their Euclidean distance falls below a specified threshold. No merging occurs when the generated waypoint lies in an unexplored region.

(NMS) to select the top K candidate waypoints.

C. Topological Graph

In prior work, the memory mechanisms of LLM-based navigators in continuous environments [10], [14], [15] have been limited to step-wise trajectory history, represented either through visual observations or textual prompts. Such histories record only visit order, without capturing the topological structure of explored areas. As a result, prompts provide merely sequential trajectories, limiting the LLM’s ability to reason about spatial relationships and connectivity. This often leads to two issues: (1) the agent getting stuck or looping in the same area, and (2) the inability to backtrack and re-explore unvisited regions after a navigation error.

Our approach, drawing inspiration from ETPNav [5], dynamically updates the topological graph by incorporating newly generated waypoints at each time step. The topological graph at each step is denoted as $G_t = (N_t, E_t)$, where N_t denotes the set of all nodes and there exists an edge $e_{ij} \in E_t$ between node n_i and n_j if they are directly reachable via a straight line. The nodes $n_i \in N_t$ are categorized as visited if the agent has already explored them, and as unvisited if they have only been observed but not yet explored.

The topological graph update mechanism should maintain **simplicity**, **effectiveness**, and **consistency**, ensuring the MLLM receives a clear spatial representation of explored versus unexplored regions:

- 1) **Avoid regenerating waypoints for visited nodes:**
When returning to a previously visited node, no new

waypoints are generated. Instead, the original waypoints are reused. This ensures that the same action options remain available and preserves a consistent spatial structure within that region.

- 2) **Reuse nearby nodes:** Assign new waypoints close in Euclidean distance to existing nodes the same ID and location, preventing clutter and reducing reasoning overhead.
- 3) **Remove duplicate waypoints:** Keep only one waypoint among those close in distance, avoiding redundancy and confusion during selection.

Algorithm 1: GRAPHUPDATE

Input: previous topological graph G_{t-1} ; current node n_t ; current observations o_t

Output: Updated graph G_t

- 1 $G_t \leftarrow G_{t-1}$
 - 2 **if** n_t *is visited* **then**
 - 3 **return** G_t
 - 4 $m_t \leftarrow \text{get_obstacle_map}(o_t)$
 - 5 $\mathcal{W}_t \leftarrow \text{waypoint_predictor}(m_t)$
 - 6 $G_t \leftarrow \text{MERGINGMODULE}(G_t, \mathcal{W}_t, n_t)$;
 - 7 **return** G_t
-

Based on these principles, the overall pipeline of the graph update algorithm first determines whether the current location has been visited to decide whether new waypoints should be generated. If new waypoints are required, they are incorporated into the graph through a merging module. The detailed procedures are presented in Algorithm 1 and Fig. 2 illustrates the process of the merging module.

D. MLLM-Based Planner

In discrete VLN, methods such as NavGPT [12] and DiscussNav [22] used multi-model pipelines, where VLMs (e.g., BLIP [39]) generated scene descriptions, LLMs summarized navigation history, and reasoning was performed over textualized inputs. This design suffers from two limitations: compressing visual observations into text loses critical information, and splitting the task across models weakens memory integration. With the advent of MLLMs such as GPT-4o [31], recent works like MapGPT [4] instead leverage raw visual histories and a single MLLM for joint reasoning, while AO-Planner [15] demonstrated the effectiveness of such prompting in continuous environments. However, history-aware prompting alone cannot fully exploit exploration: it lacks (1) a topological representation of spatial connectivity and (2) explicit tracking of visited/unexplored waypoints. To address this, we propose a TopoGraph-and-VisitInfo-aware prompting system for VLN-CE, which equips the MLLM with structured spatial awareness and enables local path planning and re-exploration of visited regions.

1) TopoGraph-and-VisitInfo-Aware Prompting System:

Our prompting system receives: (1) Instruction I , a global step-by-step guidance; (2) History H_t , the sequence of visited locations visualized in their corresponding images;

TopoGraph-VisitInfo-Aware Prompting System

History lists the locations you have visited along with their images. It may include correct landmarks mentioned in the 'Instruction' and also some past erroneous explorations.

Trajectory records the sequence of visited Location IDs. You always start from Location 0, and your current position is the last ID in 'Trajectory'.

Topological Graph refers to the connectivity between the locations you have visited and other locations you have observed but not yet visited.

Supplementary Info records unvisited locations that are not in current action options in topological graph and their images.

Previous Planning records your earlier long-term multi-step planning, which you can now refer to.

Action Options are navigable location IDs with observed images from front, back, left, and right views. Each location is provided with a direction label. You must choose one location from this set as your next move. Some direction images may contain no selectable locations, but they still help you understand your current position, the surrounding environment, and whether the destination is visible.

VisitInfo classifies location IDs in 'Action Options' as either visited or unvisited. Normally, you should select an unvisited location unless you must backtrack.

Requirements For each provided image of the environment, combine the 'Instruction' with visible details such as scene layout, landmarks, and objects. Align 'Instruction' with 'History' to estimate your execution progress and use the 'Topological Graph' for planning.

Stop Condition If you can clearly see the destination in the images provided with 'Action Options' and it is close enough, choose 'Stop'. Otherwise, continue moving toward it.

Thought In 'Thought', explain your reasoning by combining 'Instruction', 'History', 'Trajectory', 'Action Options', and the provided images.

New Planning Based on 'Previous Planning' and your current 'Thought', update your multi-step plan in 'New Planning'.

Fig. 3. Definition of Each Term in our Prompting System

(3) Trajectory Tr_t , the sequence of visited location IDs; (4) Topological Graph G_t , the connectivity between nodes; (5) VisitInfo V_t , the visitation status of the current action options; (6) Supplementary Info S_t , locations that are unvisited and their corresponding images. (7) Action Options A_t , navigable location IDs with observed images from the front, back, left, and right views. The output of our prompting system are Thought T_t and Action a_t . The detailed definition of each term in our prompting system are listed in Fig. 3, and it serves as the system prompt. Unlike AO-Planner [15], which projects RGB pixels into 3D using depth images and camera intrinsics, we instead use predicted 3D waypoint positions with camera intrinsics to project them back onto the RGB images. To incorporate the topological graph into the prompting system, we follow the approach of MapGPT [4] by representing the connectivity of each visited node to others through natural language descriptions. VisitInfo can be directly obtained from the topological graph by the distinction between visited and unvisited nodes.

$$T_t, a_t = MLLM(I, H_t, Tr_t, G_t, V_t, S_t, A_t)$$

2) *Encourage Exploration and Equip MLLMs with Local Path Planning*: The Topograph-and-VisitInfo-aware prompting system alleviates two common challenges faced by zero-shot methods in the VLN-CE setting. First, prior approaches do not record visitation information, which often causes the agent to repeatedly explore the same region without awareness, ultimately becoming trapped in a confined loop. By incorporating visit information, our system encourages exploration by prioritizing unvisited locations as candidate targets, unless a navigation error occurs and the MLLM determines that backtracking to previously visited locations is

necessary. We observe that this exploration-oriented strategy effectively mitigates the issue of the agent getting stuck in a limited area.

Second, when navigation errors arise, earlier prompting systems fail to provide sufficient information for the MLLM to conduct path planning for error correction. Our system addresses this by supplying the topological graph as a source of spatial structural knowledge. This allows the MLLM to utilize connectivity information between locations for planning, such as backtracking to a previously visited node and then exploring other interested unvisited locations, or even retreating to earlier visited locations to reconsider exploration strategies. We find that this capability of local path planning enables the MLLM to timely correct its route when navigation errors occur, thereby preventing large deviations from the instruction-guided trajectory. Fig. 4 illustrates the complete prompting information at a single step, as well as how the MLLM leverages the topological graph to perform local path planning for error correction.

IV. EXPERIMENTS

A. Experimental Settings

Dataset We evaluate our navigator on R2R-CE [8] and RxR-CE [7]. R2R-CE extends the R2R dataset [2] into continuous environments within the Habitat simulator [3]. RxR-CE is more challenging, with longer instructions, longer trajectories, and no sliding along obstacles, increasing the chance of getting stuck. We compare our method against both learning-based and LLM/MLLM-based navigators. For R2R-CE, evaluation is conducted on the full val-unseen split (1,839 episodes across 11 scenes). For RxR-CE, following prior work [6], [13], [15], we use 300 sampled episodes from the English val-unseen split. To reduce API costs, we also sample 300 episodes for ablation studies. All experiments use the GPT-5-mini API [11] as the MLLM.

We train our waypoint predictor on the train split of ground-truth connectivity graph used in DC-VLN [24], which is adapted from the pre-defined graph for MP3D [1] to fit the continuous environments in Habitat. We use the AdamW [40] optimizer with a learning rate of 10^{-4} , a batch size of 32, and 30 training epochs. Evaluation is performed on the val-unseen split of the ground-truth connectivity graph.

Evaluation Metrics For the evaluation of our navigator in R2R-CE task, we use the standard metrics used in VLN-CE: navigation error (NE), oracle success rate (OSR), success rate (SR), success weighted by path length (SPL) and normalized dynamic-time warping (nDTW).

For the evaluation of our waypoint predictor, we adopt the same metrics as DC-VLN [24]. Specifically, $|\Delta|$ denotes the difference between the number of predicted waypoints and the ground-truth targets, %Open represents the proportion of predicted waypoints located in open space, and Avg_Score represents the score of the generated waypoints on the ground-truth heatmap, where waypoints closer to the ground-truth waypoint receive higher scores. In addition, d_C and d_H correspond to the Chamfer distance and Hausdorff

Equip MLLM with Local Path Planning

Instruction: Walk towards fireplace, walk on the left side of the brown couch, turn right after couch, turn left after four person table and stop in front of barstools.

Previous Planning: 'Step 6 (completed): Arrived at location 7.'
'Step 7 (now): Move from location 7 to location 10 to reach the four- person table.'
'Step 8: Turn left after the four-person table and proceed toward the barstools.'
'Step 9: Stop in front of the barstools (end).'

History: Step 0, turn sharp left and go straight to location 3 which is on the carpet in front of the fireplace, near the edge of the sofa in the image.

.....
Step 6, turn around and go straight to location 10 which is located on the carpet in front of two blue-cushioned chairs in the image.

Trajectory: 0 3 7 4 5 4 7 10

Topological Graph: Location 0 is connected with Locations 1, 2, 3, 4, 5, 10.
Location 3 is connected with Locations 0, 6, 7, 8, 2, 9.
Location 7 is connected with Locations 3, 4, 10, 11, 12. ...

Supplementary Info: Corresponding images of locations 12, 13, 2, 14, 11, 8, 9, 15 (not shown).

Action Options: (front view) Location 5: *visited*, turn slightly right.
(left view) no visible locations in the image.
(back view) Location 7: *visited*, turn around. Location 6: *unvisited*, turn around.
(right view) Location 18: *unvisited*, turn slightly right. Location 0: *visited*, turn right.



Thought: I am currently at location 10 (at the four-person table). The instruction says to turn left after the table and proceed toward the barstools. The barstools area appears to be reached via the seating area connected at location 7 (which leads to location 11 where the kitchen/bar is visible). To follow the instruction and approach the barstools efficiently, I should move back to location 7 now and then proceed to the barstools area (location 11) next.

New Planning: "Step 7 (completed): Arrived at location 10 (four-person table)."
"Step 8 (now): Move from location 10 to location 7 to head toward the barstools."
"Step 9: From location 7, go to location 11 (barstools/kitchen area)."
"Step 10: Stop in front of the barstools (end)."

Action: 7

Fig. 4. An example of the complete prompting information at a single step. In this scenario, a navigation error occurs, and the MLLM leverages the connectivity information in the topological graph to perform local path planning for correction: it first backtracks to a previously visited location and then selects an unvisited node that is better aligned with the instruction.

distance, respectively, both of which are standard measures for quantifying discrepancies between point clouds.

Because the planner relies on an API-based MLLM, overall latency is dominated by the external model response. The waypoint predictor runs in real time on a commodity GPU (NVIDIA L40S in our experiments), and prompt construction introduces negligible overhead. Future local multimodal models may further reduce deployment latency.

B. Experimental Results

Table I reports the results on R2R-CE. As expected, supervised methods generally outperform zero-shot ones due to task-specific training. Our method, however, establishes a

TABLE I

COMPARISON OF DIFFERENT METHODS ON THE VAL-UNSEEN SPLIT OF R2R-CE IN HABITAT SIMULATOR

Method	NE↓	OSR↑	SR↑	SPL↑
Supervised				
Seq2Seq [8]	7.77	37	25	22
CMA [8]	7.37	40	32	30
DC-VLN [24]	5.89	51	42	36
Navid [36]	5.47	49	37	36
ETPNav [5]	4.71	65	57	49
Zero-Shot				
OpenNav [14]	6.70	23	19	16.1
A2Nav [13]	–	–	23	11.1
AO-Planner [15]	6.95	38	25	16.6
CA-Nav [6]	7.58	48	25	10.8
SmartWay [10]	7.01	51	29	22.0
Ours	6.12	55	41	25.4

TABLE II

COMPARISON OF DIFFERENT METHODS ON THE VAL-UNSEEN SPLIT OF RxR-CE IN HABITAT SIMULATOR

Method	NE↓	SR↑	SPL↑	nDTW↑
Supervised				
Seq2Seq [8]	12.10	13.9	11.9	30.8
DC-VLN [24]	8.98	27.1	22.7	46.7
Navid [36]	8.41	23.8	21.2	–
ETPNav [5]	5.64	54.8	44.9	61.9
Zero-Shot				
A2Nav [13]	–	16.8	6.3	–
AO-Planner [15]	10.75	22.4	15.1	33.1
CA-Nav [6]	10.37	19.0	6.0	13.5
Ours	7.56	35.7	21.7	42.4

new state-of-the-art among zero-shot approaches, showing consistent improvements across all metrics. Notably, its performance is already comparable to the average of supervised methods, narrowing the gap to strong baselines such as DC-VLN and reducing the difference with the SOTA supervised method ETPNav.

Table II shows the results on RxR-CE. Our method establishes a new state-of-the-art in the zero-shot setting, with substantial gains over previous zero-shot SOTAs. Its performance is also comparable to several supervised methods, though a noticeable gap remains compared to ETP-Nav, the strongest supervised model, particularly on SR and nDTW.

Our abstract obstacle-based waypoint predictor is significantly more lightweight and easier to train than prior models, which often rely on complex cross-attention over RGB-D features and require hundreds of training epochs. In contrast, ours can be trained in only 30 epochs. As shown in Table III, it achieves the best results on %Open, Avg_Score, and d_C . Notably, %Open is improved by 3%, which indicates that the obstacle map effectively encourages the model to generate waypoints within free space. These results further suggest that semantic information from RGB

TABLE III

COMPARISON OF DIFFERENT WAYPOINT PREDICTORS ON VAL-UNSEEN SPLIT OF MP3D GROUND TRUTH GRAPH

Model	$ \Delta $	%Open \uparrow	Avg_Score \uparrow	d_c \downarrow	d_h \downarrow
DC-VLN [†] * [24]	1.40	79.86	1.30	1.07	2.00
ETPNav* [5]	1.39	84.05	1.40	1.04	2.01
SmartWay [†] * [10]	1.41	87.26	1.48	1.03	1.96
Ours[‡]	1.41	90.18	1.56	1.02	2.00

[†] RGB as input. * Depth as input. [‡] Obstacle map as input.

TABLE IV

ABLATION EXPERIMENTS OF OUR PROMPTING SYSTEM ON THE VAL-UNSEEN SPLIT OF R2R-CE

Prompting System	SR \uparrow	SPL \uparrow
Original	46	27.5
w/o VisitInfo	40	22.7
w/o TopoGraph	37	22.6

inputs is unnecessary—while redundant details in depth images can be compactly captured by the obstacle map.

C. Ablation Studies

MLLM-based-Planner To evaluate the effectiveness of the topological graph and visit information in our prompting system, we conduct ablation studies using three variants: the proposed TopoGraph-and-VisitInfo-Aware prompting system, a version without the topological graph, and a version without visit information. Table IV shows the results of ablation experiments of our prompting system. The absence of visit information leads to a 6% drop in SR, and a 4.8% drop in SPL. The absence of topological graph results in an 9% decrease in SR and a 4.9% decrease in SPL. These results highlight the effectiveness of the topological graph for capturing spatial structure and the value of visit information for identifying explored areas.

Waypoint Predictor We compare the performance of our waypoint predictor with DC-VLN under the R2R-CE task. For fairness, both methods are evaluated using our prompting system. Table V shows the results of ablation experiments of different waypoint predictors. If the waypoint predictor is changed to DC-VLN, SR decreases by 4%, SPL decreases by 2.1% and Collision Rate increases by 6.13%. This indicates that our waypoint predictor is more likely to generate waypoints in open space and achieves better alignment with the ground-truth connectivity graph.

Comparison Under the Same API For fairness, we adapt the SOTA model AO-Planner by replacing its high-level agent API with GPT-5-mini and substituting its low-level planner with our waypoint predictor. The results of comparative experiments under the same API are shown in Table VI. Comparing the first and second rows, we observe that with the same API, our pipeline still achieves a substantial improvement, with SR higher by 13% and SPL higher by 14.8% than AO-Planner. This gain comes from both our obstacle-based waypoint predictor and our TopoGraph-and-VisitInfo-

TABLE V

ABLATION EXPERIMENTS OF WAYPOINT PREDICTORS USING OUR PROMPTING SYSTEM ON THE VAL-UNSEEN SPLIT OF R2R-CE

Waypoint Predictor	SR \uparrow	SPL \uparrow	Collision \downarrow
Ours	46	27.5	3.85
DC-VLN	42	25.4	9.98

TABLE VI

COMPARATIVE EXPERIMENTS UNDER THE SAME API ON THE VAL-UNSEEN SPLIT OF R2R-CE

Waypoint Predictor	Prompting System	SR \uparrow	SPL \uparrow
Ours	Ours-GPT5mini	46	27.5
AO-planner	AO-planner-GPT5mini	33	12.7
Ours	AO-planner-GPT5mini	38	19.9

aware prompting system. Comparing the first and third rows, we find that when using the same waypoint predictor, our prompting system improves SR by 8% and SPL by 7.6%. Finally, comparing the second and third rows, we observe that simply replacing AO-Planner’s waypoint predictor with ours yields an improvement of 5% in SR and 7.2% in SPL.

While our experiments are conducted in the Habitat simulator for controlled comparison with prior VLN-CE methods, real-world deployment introduces additional challenges such as sensor noise, depth estimation errors, and localization inaccuracies. Developing more robust obstacle-map construction and real-world validation remains important future work.

V. CONCLUSIONS

We proposed a zero-shot VLN-CE framework that combines a simplified waypoint predictor with an MLLM-based planner. The waypoint predictor operates on an abstract obstacle map, ensuring linearly reachable waypoints. We further introduced a topological graph with visitation records, incorporated into the MLLM prompt alongside visual inputs: the graph encodes spatial structure, while visitation records capture exploration history. This design enables the MLLM to reason more effectively about space and progress, encouraging exploration and enabling local path planning to backtrack and re-explore for error correction. Experiments on R2R-CE and RxR-CE show that our framework outperforms prior state-of-the-art zero-shot methods and achieves performance competitive with supervised approaches.

REFERENCES

- [1] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niebner, M. Savva, S. Song, A. Zeng, and Y. Zhang, “Matterport3d: Learning from rgb-d data in indoor environments,” in *2017 International Conference on 3D Vision (3DV)*. IEEE Computer Society, 2017, pp. 667–676.
- [2] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. van den Hengel, “Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [3] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, “Habitat: A platform for embodied ai research,” in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 9338–9346.

- [4] J. Chen, B. Lin, R. Xu, Z. Chai, X. Liang, and K.-Y. K. Wong, "Mapgpt: Map-guided prompting with adaptive path planning for vision-and-language navigation," in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, 2024.
- [5] D. An, H. Wang, W. Wang, Z. Wang, Y. Huang, K. He, and L. Wang, "Etpnav: Evolving topological planning for vision-language navigation in continuous environments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 47, no. 7, pp. 5130–5145, 2025.
- [6] K. Chen, D. An, Y. Huang, R. Xu, Y. Su, Y. Ling, I. Reid, and L. Wang, "Constraint-aware zero-shot vision-language navigation in continuous environments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [7] A. Ku, P. Anderson, R. Patel, E. Ie, and J. Baldrige, "Room-Across-Room: Multilingual vision-and-language navigation with dense spatiotemporal grounding," in *Conference on Empirical Methods for Natural Language Processing (EMNLP)*, 2020.
- [8] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, "Beyond the nav-graph: Vision-and-language navigation in continuous environments," in *European Conference on Computer Vision*. Springer, 2020.
- [9] J. Krantz, A. Gokaslan, D. Batra, S. Lee, and O. Maksymets, "Waypoint models for instruction-guided navigation in continuous environments," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 15 162–15 171.
- [10] X. Shi, Z. Li, W. Lyu, J. Xia, F. Dayoub, Y. Qiao, and Q. Wu, "Smartway: Enhanced waypoint prediction and backtracking for zero-shot vision-and-language navigation," *arXiv preprint arXiv:2503.10069*, 2025.
- [11] OpenAI, "Introducing GPT-5," <https://openai.com/index/introducing-gpt-5/>, 2025, accessed: 2025-09-07.
- [12] G. Zhou, Y. Hong, and Q. Wu, "Navgpt: Explicit reasoning in vision-and-language navigation with large language models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.
- [13] P. Chen, X. Sun, H. Zhi, R. Zeng, T. H. Li, G. Liu, M. Tan, and C. Gan, " α^2 -nav: Action-aware zero-shot robot navigation by exploiting vision-and-language ability of foundation models," *arXiv preprint arXiv:2308.07997*, 2023.
- [14] Y. Qiao, W. Lyu, H. Wang, Z. Wang, Z. Li, Y. Zhang, M. Tan, and Q. Wu, "Open-nav: Exploring zero-shot vision-and-language navigation in continuous environment with open-source llms," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [15] J. Chen, B. Lin, X. Liu, L. Ma, X. Liang, and K.-Y. K. Wong, "Affordances-oriented planning using foundation models for continuous vision-language navigation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025.
- [16] R. Liu, X. Wang, W. Wang, and Y. Yang, "Bird's-eye-view scene graph for vision-language navigation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [17] L. Wang, Z. He, R. Dang, M. Shen, C. Liu, and Q. Chen, "Vision-and-language navigation via causal learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 13 139–13 150.
- [18] J. Wang, T. Wang, L. Xu, Z. He, and C. Sun, "Discovering intrinsic subgoals for vision-and-language navigation via hierarchical reinforcement learning," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 36, no. 4, pp. 6516–6528, 2024.
- [19] S. Chen, P.-L. Guhur, M. Tapaswi, C. Schmid, and I. Laptev, "Think global, act local: Dual-scale graph transformer for vision-and-language navigation," in *CVPR*, 2022.
- [20] Z. Wang, J. Li, Y. Hong, Y. Wang, Q. Wu, M. Bansal, S. Gould, H. Tan, and Y. Qiao, "Scaling data generation in vision-and-language navigation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 12 009–12 020.
- [21] Y. Qiao, Y. Qi, Z. Yu, J. Liu, and Q. Wu, "March in chat: Interactive prompting for remote embodied referring expression," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 15 758–15 767.
- [22] Y. Long, X. Li, W. Cai, and H. Dong, "Discuss before moving: Visual language navigation via multi-expert discussions," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 17 380–17 387.
- [23] Z. Zhan, L. Yu, S. Yu, and G. Tan, "Mc-gpt: Empowering vision-and-language navigation with memory map and reasoning chains," *arXiv preprint arXiv:2405.10620*, 2024.
- [24] Y. Hong, Z. Wang, Q. Wu, and S. Gould, "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*, 2022, pp. 15 439–15 449.
- [25] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, "Beyond the nav-graph: Vision-and-language navigation in continuous environments," in *European Conference on Computer Vision*. Springer, 2020.
- [26] M. Z. Irshad, C.-Y. Ma, and Z. Kira, "Hierarchical cross-modal agent for robotics vision-and-language navigation," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 13 238–13 246.
- [27] G. Georgakis, K. Schmeckpeper, K. Wanchoo, S. Dan, E. Mitsakaki, D. Roth, and K. Daniilidis, "Cross-modal map learning for vision and language navigation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [28] P. Chen, D. Ji, K. Lin, R. Zeng, T. Li, M. Tan, and C. Gan, "Weakly-supervised multi-granularity map learning for vision-and-language navigation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 38 149–38 161, 2022.
- [29] S. Sanyal and K. Roy, "Asma: An adaptive safety margin algorithm for vision-language drone navigation via scene-aware control barrier functions," *IEEE Robotics and Automation Letters*, 2025.
- [30] J. Krantz and S. Lee, "Sim-2-sim transfer for vision-and-language navigation in continuous environments," in *European conference on computer vision*. Springer, 2022, pp. 588–603.
- [31] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [32] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [33] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *European conference on computer vision*. Springer, 2024, pp. 38–55.
- [34] B. Pan, R. Panda, S. Jin, R. Feris, A. Oliva, P. Isola, and Y. Kim, "Langnav: Language as a perceptual representation for navigation," in *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2024.
- [35] D. Zheng, S. Huang, L. Zhao, Y. Zhong, and L. Wang, "Towards learning a generalist model for embodied navigation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 13 624–13 634.
- [36] J. Zhang, K. Wang, R. Xu, G. Zhou, Y. Hong, X. Fang, Q. Wu, Z. Zhang, and H. Wang, "Navid: Video-based vlm plans the next step for vision-and-language navigation," *Robotics: Science and Systems*, 2024.
- [37] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan *et al.*, "Grounded sam: Assembling open-world models for diverse visual tasks," *arXiv preprint arXiv:2401.14159*, 2024.
- [38] Y. Long, W. Cai, H. Wang, G. Zhan, and H. Dong, "Instructnav: Zero-shot system for generic instruction navigation in unexplored environment," in *Conference on Robot Learning*. PMLR, 2025.
- [39] J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *ICML*, 2022.
- [40] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," 2019. [Online]. Available: <https://arxiv.org/abs/1711.05101>