

# Affordances-Oriented Planning Using Foundation Models for Continuous Vision-Language Navigation

Jiaqi Chen<sup>1</sup>, Bingqian Lin<sup>2</sup>, Xinmin Liu<sup>3</sup>, Lin Ma<sup>3</sup>, Xiaodan Liang<sup>2\*</sup>, Kwan-Yee K. Wong<sup>1\*</sup>

<sup>1</sup>The University of Hong Kong

<sup>2</sup>Shenzhen Campus of Sun Yat-sen University

<sup>3</sup>Meituan

## Abstract

LLM-based agents have demonstrated impressive zero-shot performance in vision-language navigation (VLN) task. However, existing LLM-based methods often focus only on solving high-level task planning by selecting nodes in predefined navigation graphs for movements, overlooking low-level control in navigation scenarios. To bridge this gap, we propose **AO-Planner**, a novel **Affordances-Oriented Planner** for continuous VLN task. Our AO-Planner integrates various foundation models to achieve affordances-oriented low-level motion planning and high-level decision-making, both performed in a zero-shot setting. Specifically, we employ a **Visual Affordances Prompting (VAP)** approach, where the visible ground is segmented by SAM to provide navigational affordances, based on which the LLM selects potential candidate waypoints and plans low-level paths towards selected waypoints. We further propose a high-level PathAgent which marks planned paths on the image input and reasons the most probable path by comprehending all environmental information. Finally, we convert the selected path into 3D coordinates using camera intrinsic parameters and depth information, avoiding challenging 3D predictions for LLMs. Experiments on the challenging R2R-CE and RxR-CE datasets show that AO-Planner achieves state-of-the-art zero-shot performance (8.8% improvement on SPL). Our method can also serve as a data annotator to obtain pseudo-labels, distilling its waypoint prediction ability into a learning-based predictor. This new predictor does not require any waypoint data from the simulator and achieves 47% SR competing with supervised methods. We establish an effective connection between LLM and 3D world, presenting novel prospects for employing foundation models in low-level motion control.

## Introduction

Large Language models (LLMs) (OpenAI 2023a; Team et al. 2023; Anil et al. 2023) have shown strong capabilities and potentials in embodied intelligence tasks. In particular, for vision-and-language navigation (VLN) task (Anderson et al. 2018; Qi et al. 2020b; Hong et al. 2021; Chen et al. 2022; Qiao et al. 2022; Wang et al. 2023b), where agents need to make decisions by comprehending language instructions and visual observations, LLMs have demonstrated impressive zero-shot navigation performance. Recent works (Zhou,

\*Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

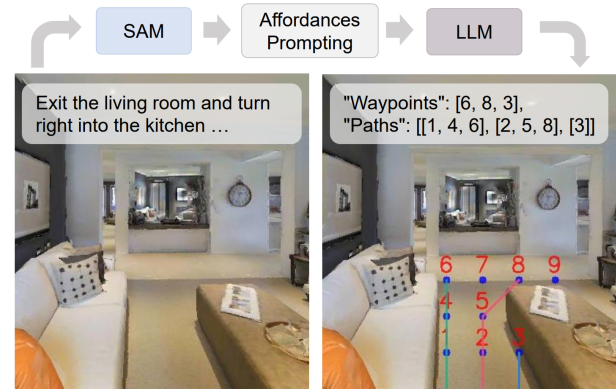


Figure 1: In discrete VLN, LLMs only need to perform high-level planning by selecting a view as the forward direction (left). For continuous environments, previous agents rely on collecting simulator data to train low-level policies. In this paper, we utilize multimodal foundation models and propose visual affordances prompting to predict low-level candidate waypoints and paths in a zero-shot setting (right).

Hong, and Wu 2023; Long et al. 2023; Chen et al. 2024b) have effectively leveraged the high-level capabilities of LLMs for task planning, decision-making, instruction parsing, etc.

However, there exist several challenges and gaps which hinder the application of LLM-based agents to more practical navigation tasks. For instance, agents in the classical discrete VLN setting only need to navigate between predefined viewpoints with high-level decision-making capabilities. In contrast, for more practical VLN in continuous environments (VLN-CE), learning-based models (Krantz et al. 2020; Krantz and Lee 2022; Hong et al. 2022; Wang et al. 2023a; An et al. 2024; Wang et al. 2024) must predict low-level actions (e.g., turn left  $15^\circ$ ) or positions of subgoals in the environment that can drive agents in the simulator. Recently proposed zero-shot agents for navigation in continuous environments (Cai et al. 2023; Chen et al. 2023) only utilize LLMs for high-level task planning, and they still require a large amount of sampled environmental data from the simulator to train low-level policies. Obviously, the gaps between discrete and continuous environments have made current re-

search inadequate in exploring and exploiting the practical navigation potentials of LLMs. Hence, we raise the following question: *can LLMs not only handle high-level tasks but also serve as low-level motion planners?*

In this paper, we propose a novel zero-shot affordances-oriented planning framework, named **AO-Planner**, which can navigate agents in continuous environments and is capable of handling both low-level path planning and high-level decision-making. Our key target is to bridge the gap between LLM’s RGB space predictions and 3D world navigation via pixel-level path planning in *affordances*. First, we utilize Grounded SAM (Ren et al. 2024) to segment navigable regions on the ground as navigational affordances and sample points for the LLM to choose from. We then propose a novel Visual Affordances Prompting (VAP) method to leverage LLM’s spatial understanding abilities and achieve low-level motion planning in a zero-shot manner. As shown in Figure 1, the LLM (Team et al. 2023; Reid et al. 2024) is tasked with selecting the most appropriate waypoints from affordances and connecting some of the points in affordances to plan navigable paths while avoiding obstacles. This process is conducted entirely in the RGB space, in which LLM excels, without the necessity to predict 3D spatial coordinates directly. These low-level candidates (including waypoints and the corresponding paths) are then visualized and fed to another zero-shot PathAgent in the second stage, which is responsible for selecting the path to follow. Finally, by combining depth information with the camera intrinsic parameters, pixel predictions in the RGB space are converted into a series of 3D coordinates that navigate the agent to the designated location.

Experiments on the R2R-CE (Krantz et al. 2020) and RxR-CE (Ku et al. 2020) datasets demonstrate that our proposed framework achieves state-of-the-art zero-shot performance. The previous best-performing method,  $A^2$ Nav (Chen et al. 2023), still relies on sampling a large amount of data from the simulator to train low-level policies. In contrast, our approach has shown an improvement of 2.9% in SR on R2R-CE and 5.6% in SR on RxR-CE, without using any training data. For the SPL metric on RxR-CE, our method can even achieve an 8.8% performance improvement. To simulate the challenges of motion planning in the continuous Habitat simulator (Savva et al. 2019), we also adopt a setting that avoids using any additional data, such as the navigation graph from discrete MP3D simulator (Chang et al. 2017), to train low-level policies. We implement a variant of the AO-Planner where we utilize waypoints predicted by LLMs as pseudo-labels, replacing the navigation graph obtained from the MP3D simulator as pseudo ground truths (p-GTs), to train a new learning-based waypoint predictor at the same data scale. We combine this distilled waypoint predictor with the VLN agent used in ETPNav (An et al. 2024) and achieve competitive performance, with an SR of 47%. The tremendous potentials of multimodal LLM in low-level motion planning are revealed through AO-Planner’s impressive performance.

Our main contributions can be summarized as follows.

- We propose AO-Planner for the VLN-CE task. It leverages foundation models for affordances-oriented planning and converts predictions in the RGB space to 3D

coordinates, bridging the gap between LLM’s high-level decision-making and 3D world navigation.

- We present a novel visual affordances prompting (VAP) method to unleash spatial understanding and reasoning abilities of LLMs. This also uncovers previously unexplored motion planning abilities of LLMs.
- Our AO-Planner achieves state-of-the-art zero-shot performance without using any simulator data for training low-level policies. It can also provide reliable pseudo-labels for training supervised models to achieve competitive performance.

## Related Work

**Vision-and-Language Navigation (VLN)** VLN (Anderson et al. 2018) is a representative task in the field of embodied AI, requiring an agent to combine instructions and visual observations to move through complex environments and reach a target location. In this task, the most classic scenario is the discrete VLN setting based on the MP3D simulator (Chang et al. 2017), where an agent only needs to select a node from a predefined navigation graph and teleport to it. Therefore, previous works (Wang et al. 2019; Ma et al. 2019; Deng, Narasimhan, and Russakovsky 2020; Hong et al. 2021; Chen et al. 2021, 2022; Qiao et al. 2022; An et al. 2022; Wang et al. 2023b) have focused on endowing agents with high-level decision-making abilities without considering low-level motion.

In recent years, VLN in continuous environment (VLN-CE) built on the Habitat simulator (Savva et al. 2019) has received much attention. In this setting, the action space only consists of low-level actions, such as forward and rotate, instead of teleporting the agent directly. To overcome the challenge of directly predicting low-level actions, some works (Krantz et al. 2021; Hong et al. 2022; Krantz and Lee 2022; An et al. 2024) train a waypoint model to predict candidate waypoints in the surroundings as a replacement for the candidate waypoints provided by the MP3D simulator, bridging the gap between discrete and continuous VLN. However, such waypoint models rely on the simulator-specific navigation graph data from MP3D for training, and may suffer from generalization issues. Furthermore, the candidate positions predicted by these models are only some locations near the agent, which can be reached simply by moving forward without considering motion planning.

**Navigation with Large Language Models** In navigation tasks, there is often a wealth of visual and semantic information involved. Hence, many LLM-based navigation methods have emerged in recent years. Some methods (Zhou, Hong, and Wu 2023; Long et al. 2023; Chen et al. 2024b) directly utilize GPT as agents and achieve zero-shot navigation through appropriate prompting methods. Another group of methods (Pan et al. 2023; Lin et al. 2024; Zhang et al. 2024; Zheng et al. 2024), based on open-source LLMs such as Llama (Touvron et al. 2023), collect data and design tasks to fine-tune the LLMs, enabling them to have impressive navigation performance while maintaining interpretability and universality. There are also methods which incorporate LLMs as part of their frameworks (Cai et al. 2023; Chen et al.

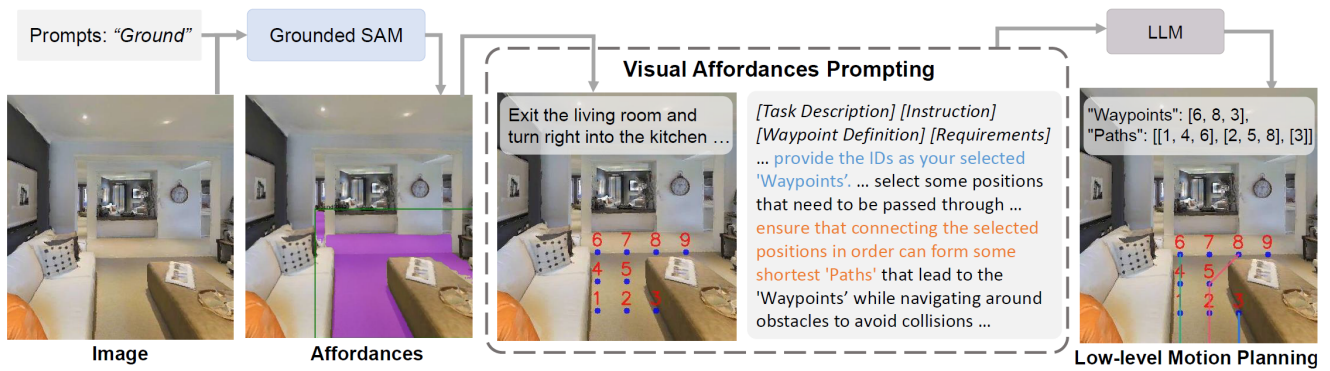


Figure 2: Our proposed low-level affordances-oriented planning framework with visual affordances prompting. First, we utilize Grounded SAM to segment the visible ground as affordances. We then introduce visual affordances prompting (VAP), where we uniformly scatter points with numeric labels within the affordances. After querying the LLM by combining the visualized new image with task definition, instruction, waypoint definition, and output requirements, we finally obtain potential waypoints and paths in this view.

2023), leveraging LLMs’ high-level capabilities for handling subtasks such as instruction parsing, scene description, and decision-making. However, these approaches merely exploit the high-level capabilities of LLMs and still require collecting data to train low-level policies. The potentials of LLMs in low-level motion planning remain unexplored.

**Visual Prompting** Prompting is an important technique for activating various capabilities of LLMs. Recently, with the development of multimodal LLMs, some visual prompting methods (Yang et al. 2023a,b; Nasiriany et al. 2024; Lei et al. 2024) have emerged. For example, SoM (Yang et al. 2023a) extracts and labels certain objects through object detection or segmentation to enhance the visual grounding capability of LLMs. Yang et al. (2023b) systematically summarize the results and potentials achieved using different visual markers on the multimodal GPT-4V (OpenAI 2023b) model. These methods have demonstrated the importance of visual prompting in unleashing the multimodal understanding capabilities of LLMs. Nonetheless, they do not address motion planning problem in navigation tasks.

## Method

### VLN-CE Task Definition

The task of vision-language navigation can be formulated as follows. For each episode, the agent needs to follow a fixed instruction  $I$  to move from the starting point to the target destination. At step  $t$ , the agent can obtain an observation  $O_t$ , which includes views  $\{V_t^i\}_{i=1}^N$  from different directions. In our framework, we set  $N = 4$  and collect non-overlapping views from the front, back, left, and right directions as observation, i.e.,  $O_t = \{V_t^i\}_{i=1}^4$ . For the action space, the VLN-CE task defines four parameterized low-level actions, namely FORWARD (0.25m), ROTATE LEFT/RIGHT ( $15^\circ$ ), and STOP. In this work, the agent needs to select a waypoint location and the corresponding path as an action  $a_t$ , and then convert the relative position into a series of low-level actions.

### Framework Overview

As shown in Figure 2 and Figure 3, we follow some popular methods (Hong et al. 2022; Krantz and Lee 2022) in the VLN-CE task and design a two-stage framework, with two LLM-based agents responsible for predicting candidate paths at the low level and making decisions based on all environmental information at the high level. This simplifies and reduces the difficulty of the task since there is no need to simultaneously address two tasks of completely different levels.

Specifically, we first utilize the Grounded SAM (Ren et al. 2024) model to obtain navigational affordances and sample some points from them, facilitating the selection by the LLM agent. We further design prompts for the low-level agent to search for potential waypoints and plan corresponding paths by connecting the sampled points within the affordances. These candidate results are visualized and fed to the high-level agent in the second stage, where we refer to the previous zero-shot method (Chen et al. 2024b) that combines instructions and historical information to make the final decision on the location to move to. Additionally, points predicted in RGB space in the first stage can be transformed into 3D world coordinates based on depth information and camera intrinsics, and further be converted into low-level actions, thereby connecting the LLM with the 3D world.

### Visual Affordances Prompting (VAP)

Prompting technique has been widely used to unleash the powerful capabilities of LLMs. For multimodal LLMs such as GPT-4 (OpenAI 2023a,b) and Gemini (Team et al. 2023; Reid et al. 2024), the importance of visual prompting is gradually being recognized (Yang et al. 2023a,b; Nasiriany et al. 2024). In this work, we propose a novel Visual Affordances Prompting (VAP) approach for low-level planning in continuous VLN task. As shown in Figure 2, we mark a set of navigable points within affordances, from which the LLM selects candidate waypoints and connects some points to form candidate paths, enabling low-level planning. VAP effectively achieves zero-shot path planning in the VLN-CE task.

**Navigational Affordances** “Affordances” is a widely used concept in the field of robotics, referring to the regions in the environment where an agent’s actions can be performed (Gibson 2014). Previous works in navigation tasks (Qi et al. 2020a; Luddecke and Worgotter 2017) have explored how to learn from supervised data and ultimately predict safe and navigable affordances in the environment.

Fortunately, we have found that current foundation models have the capability to solve this problem in a zero-shot setting. The Grounded SAM (Ren et al. 2024) method we adopt combines an open-set detector Grounding Dino (Liu et al. 2023) with Segment Anything Model (SAM) (Kirillov et al. 2023). By inputting the prompt “Ground”, the open-set detector can detect the visible ground and provide the corresponding bounding boxes to SAM for segmentation, resulting in high-quality navigational affordances. In the environment, we set the FOV of the agent’s camera to 90 degrees and collect observations from four directions in counterclockwise order, namely front, left, back, and right views. For each view  $V_t^i$  in direction  $i$ , the output of Grounded SAM is a set of  $k$  masks  $\{m_1, \dots, m_k\}$ , and we fuse these masks into one to represent the final affordances  $A_t^i$ .

**Visual Prompting** Inspired by previous Set-of-Mark (SoM) prompting (Yang et al. 2023a), we propose a novel visual affordances prompting method to assist LLM in unleashing its visual grounding and affordances-oriented spatial planning. Given a view  $V_t^i$  of size  $h \times w$ , we uniformly distribute a set of points within its affordances  $A_t^i$ . We then label these points in order and visualize them on the image, resulting in an augmented view  $V_t^{i*}$ . This allows us to easily query LLM to select suitable points and provide their corresponding IDs without directly predicting their coordinates in RGB space or the 3D world.

Compared to the previous waypoint models (Hong et al. 2022; Krantz and Lee 2022; An et al. 2024), which only take RGB and/or depth information as inputs and directly predict reachable locations near the current position, we incorporate semantic information and require path planning of the first-stage agent to navigate around obstacles. Specifically, we include the instructions  $I$  of the current episode as part of the prompts, asking the agent to examine key information such as scene descriptions, landmarks, and objects, and to select corresponding waypoints and paths. Additionally, we provide some low-level task descriptions  $D_L$ , such as specific definitions for waypoints and paths, like requiring waypoint predictions to maintain a distance from obstacles and be located in crucial regions, and connecting points to form paths that navigate around obstacles. The prompting process can be formulated as

$$W_t^i, P_t^i = LLM(D_L, I, V_t^{i*}), \quad (1)$$

where  $i$  represents one of the four view directions. We further merge the predicted waypoints  $W_t^i$  and paths  $P_t^i$  from four directions into two candidate lists  $W_t$  and  $P_t$  for waypoints and paths at the current position. Our proposed prompts help LLM better understand the task, enabling it to predict waypoints and plan reasonable paths based on affordances.

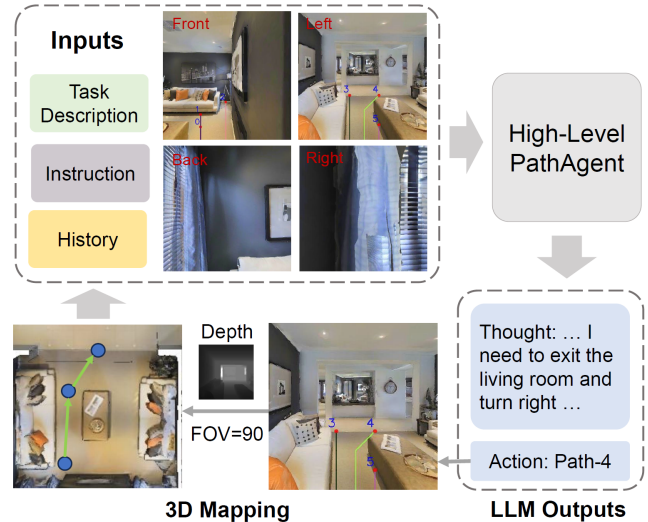


Figure 3: Our proposed high-level PathAgent. Different from previous zero-shot VLN agents, we utilize visual prompting by marking candidate waypoints and their corresponding paths (i.e., Path 0-5) in all four observation directions. This allows the PathAgent to make action decisions in the proficient RGB space and then map pixel-based paths to 3D coordinates using depth information and camera intrinsic parameters.

### High-level PathAgent

After obtaining potential low-level waypoint and path predictions, we introduce another agent, dubbed **PathAgent**, to perform high-level decision-making. As shown in Figure 3, we refer to MapGPT designed for discrete VLN tasks (Chen et al. 2024b) and make some improvements specifically for the VLN-CE task. Instead of directly providing some captions of the visual observations or providing originally observed images to a multimodal LLM, we also utilize the idea of visual prompting. For a given step  $t$ , we visualize the candidate waypoints  $W_t$  and paths  $P_t$  predicted in the first stage by rendering them on the visual observation  $O_t$  (containing four perspective directions) and marking them with IDs, resulting in an augmented observation  $O_t'$ . Through visual prompting, we only need four images to retain all observational information, without relying on the simulator to provide images corresponding to candidate waypoints like MapGPT (Chen et al. 2024b), which may neglect some observations. Additionally, we add high-level task description  $D_H$ , instruction  $I$ , and historical information  $H_t$  as inputs to the LLM. Following previous work, we require the LLM to output an interpretable thinking process  $T_t$  and select one path from  $P_t$  set as action  $a_t$ . The entire process is defined as

$$T_t, a_t = LLM(D_H, I, H_t, O_t', W_t, P_t). \quad (2)$$

### 3D Mapping and Motion Control

For our zero-shot framework, once the agent selects a waypoint as a subgoal, we need to convert the planned path into a sequence of actions to guide the agent along the path. During the low-level motion planning in the first stage, the LLM

agent plans paths in the RGB space by connecting points. A pixel-based path is actually a series of line segments. We can utilize the camera intrinsic parameters and depth information to map each point in the line segments from pixel coordinates to 3D world coordinates. For each line segment, we follow the low-level controller in ETPNav (An et al. 2024), which designs a rotate-then-forward control flow. Based on the world coordinates of two points in a line segment, we calculate their relative orientation and distance, and convert them into a series of low-level ROTATE and FORWARD actions within the action space. Through this approach, we bridge the gap between the LLM’s predicted 2D results and the 3D world, successfully applying the spatial planning abilities of LLM to low-level motion.

## Waypoint Distillation

In this work, we not only use LLMs with VAP as zero-shot waypoint predictors, but also explore transferring this capability to a learning-based waypoint predictor to obtain similar waypoint predictions. Specifically, we utilize our LLM-based waypoint predictor as a zero-shot data annotator to predict waypoints as pseudo ground-truths (p-GTs). We then leverage these p-GTs to train the waypoint predictor in ETPNav (An et al. 2024) and combine it with the high-level VLN agent used in ETPNav. By distilling the waypoint prediction capability of LLMs to a learning-based waypoint predictor, we can achieve more competitive performance as it can be combined with a learning-based high-level agent for fine-tuning, without relying on additional MP3D simulator data. This is also a cost-effective paradigm since the scale of the pseudo labels we collect is the same as the scale of the MP3D data used by the previous waypoint predictor. We sample around 10K locations with 40K images for training, which only requires 40K calls to Gemini.

## Experiments

In this section, we present the main experimental results. For implementation details and more results, please refer to the appendices in our arXiv version paper (Chen et al. 2024a).

### Experimental Setting

**Dataset** We conduct experiments on the challenging R2R-CE (Krantz et al. 2020) and RxR-CE (Ku et al. 2020) datasets. R2R-CE is derived from the discrete path annotations from the R2R dataset (Anderson et al. 2018) and is converted into continuous environments with the Habitat simulator (Savva et al. 2019). R2R-CE provides step-by-step high-level instructions which are corresponding to indoor navigation trajectories that should be followed by agents. RxR-CE is a larger and more challenging dataset, which has longer instructions, and contains multilingual descriptions. However, its scene splits are similar to R2R-CE, and experiments are also conducted using the Habitat simulator. We follow some previous zero-shot work (Zhou, Hong, and Wu 2023; Long et al. 2023; Chen et al. 2024b) in the discrete VLN task, evaluating AO-Planner on the entire validation unseen set of R2R-CE and a random sampling subset with 500 cases from the validation

unseen set of RxR-CE. To save API costs, we also additionally sample a subset containing 100 cases from the validation unseen set of R2R-CE for ablation study.

**Evaluation Metrics** Following previous works (Krantz et al. 2020; Hong et al. 2022), we adopt some evaluation metrics widely used in the VLN tasks, namely navigation error (NE), oracle success rate (OSR), success rate (SR), success weighted by the inverse of the path length (SPL), and normalized dynamic time wrapping (nDTW), respectively.

## Experimental Results

**Zero-Shot Performance** As shown in Table 1, our method achieves state-of-the-art performance under zero-shot setting. It shows a 5.5% SPL improvement on R2R-CE and an 8.8% SPL improvement on RxR-CE over the previous best-performing method  $A^2$ Nav (Chen et al. 2023). Furthermore, while  $A^2$ Nav does not require any instruction-goal pair demonstrations, it still needs to sample a large number of image-path pairs from the Habitat simulator to train low-level policies, which may suffer from potential generalization issues. Some other methods that do not rely on simulator data show worse performance. For example, Cow (Gadre et al. 2022) only achieves a success rate of 7.8% and an SPL of 5.8% on R2R-CE. In contrast, our proposed method, while not requiring any data for training, significantly outperforms previous zero-shot methods.

Our zero-shot performance is comparable to three supervised methods, SASRA (Zubair Irshad et al. 2021), Seq2Seq (Krantz et al. 2020), and NaVid (Zhang et al. 2024) on R2R-CE or RxR-CE. However, we also observe that the current zero-shot performance still has a significant gap compared to supervised performance, which aligns with the performance comparison of zero-shot and supervised methods in the discrete VLN task (Zhou, Hong, and Wu 2023). This may be due to LLMs’ unfamiliarity with this VLN-CE task, lacking direct learning from the Habitat simulator. Additionally, the image quality of continuous scenes reconstructed by the Habitat simulator is not as high as images directly captured in MP3D. These factors affect the performance of LLMs.

It is important to note that we do not modify any prompts when applying AO-Planner to both R2R-CE and RxR-CE, which also demonstrates that our zero-shot method is more general and does not require additional fine-tuning like dataset-specific supervised models.

**Distilling Waypoint Predictor** We aim to transfer the capability of LLM-based waypoint prediction to a learning-based waypoint predictor to avoid reliance on discrete MP3D data. For the waypoint model selection, we adopt the Chamfer distance metric used by ETPNav to evaluate the alignment between the model’s waypoint predictions and the ground truths. Our waypoint predictor achieves a result of 1.08, which is slightly worse than the 1.04 achieved in the original ETPNav. This further demonstrates that using predictions from LLM as pseudo GTs can achieve high-quality waypoint predictions.

As shown in Table 1, the experiments on R2R-CE and RxR-CE datasets demonstrate the potential of our proposed AO-Planner. By using the waypoint predictor to provide sub-goals, AO-Planner significantly outperforms NaVid, which

Settings	Methods	Simulator Data	R2R-CE				RxR-CE			
			NE↓	OSR↑	SR↑	SPL↑	NE↓	SR↑	SPL↑	nDTW↑
Supervised	SASRA (Zubair Irshad et al. 2021)	✓	8.32	-	24	22	-	-	-	-
	Seq2Seq (Krantz et al. 2020)	✓	7.77	37	25	22	12.1	13.9	11.9	30.8
	LAW (Raychaudhuri et al. 2021)	✓	-	-	35	31	10.9	8.0	8.0	38.0
	DC-VLN (Hong et al. 2022)	✓	6.20	52	41	36	8.98	27.1	22.7	46.7
	NaVid (Zhang et al. 2024)	✓	5.47	49	37	36	8.41	23.8	21.2	-
	ETPNav (An et al. 2024)	✓	4.71	65	57	49	5.64	54.8	44.9	61.9
	AO-Planner (Ours)		5.55	59	47	33	7.06	43.3	30.5	50.1
Zero-Shot	CLIP-Nav (Dorbala et al. 2022)		-	-	5.6	2.9	-	9.8	3.2	-
	Seq CLIP-Nav (Dorbala et al. 2022)		-	-	7.1	3.7	-	9.1	3.3	-
	Cow (Gadre et al. 2022)		-	-	7.8	5.8	-	7.9	6.1	-
	ZSON (Majumdar et al. 2022)	✓	-	-	19.3	9.3	-	14.2	4.8	-
	$A^2$ Nav (Chen et al. 2023)	✓	-	-	22.6	11.1	-	16.8	6.3	-
	AO-Planner (Ours)		6.95	38.3	<b>25.5</b>	<b>16.6</b>	10.75	<b>22.4</b>	<b>15.1</b>	33.1

Table 1: Comparison with supervised and zero-shot methods on validation unseen split of R2R-CE and RxR-CE. Our approach achieves state-of-the-art zero-shot performance on two benchmarks (zero-shot AO-Planner) and also obtains competitive supervised performance via distilling LLM’s waypoint prediction ability into a learning-based predictor (supervised AO-Planner), without relying on simulator data for training low-level models. In contrast,  $A^2$ Nav requires sampling simulator data for training policies and ETPNav utilizes the navigation graph from discrete MP3D simulator for training a waypoint predictor.

Prompting Methods	OSR↑	SR↑	SPL↑
Visual Affordances Prompting	39	27	16.9
w/o Instructions	36	24	14.6
w/o Waypoint Definitions	35	23	14.7
w/o Affordances	27	15	7.2

Table 2: Ablation study on different prompting settings.

does not use any waypoint predictor, in both OSR and SR. We notice that the performance of AO-Planner is still behind ETPNav trained on MP3D ground truth. However, our method does not rely on any manual annotation or simulator data for training low-level policies. In this work, we point out the limitations of existing methods and re-demonstrate the effectiveness of waypoint prediction in a more practical setting. Our AO-Planner provides a meaningful reference by leveraging foundation model in not only serving as zero-shot predictors but also providing reliable waypoint data for training, free from the constraints of simulator data.

## Ablation Study

**Different Prompting Methods** As shown in Table 2, we design three representative prompting methods to compare with our proposed Visual Affordances Prompting (VAP). The waypoint model designed for low-level motion in previous work (Hong et al. 2022; Krantz and Lee 2022; An et al. 2024) typically only use RGB and/or depth information as input. Given LLM’s semantic understanding capabilities, we have incorporated instructions as part of the prompts in VAP to help LLMs check scene descriptions, landmarks, and objects information. Our method with instructions obtains a 3% improvement in SR and a 2.3% improvement in SPL.

In VAP, we have also made some simple definitions for waypoints, including obstacle avoidance and finding crucial

Methods	OSR↑	SR↑	SPL↑
Zero-Shot Predictor + Baseline Agent	31	20	12.4
Supervised Predictor + Baseline Agent	34	23	13.1
Zero-Shot Predictor + PathAgent	39	27	16.9
Supervised Predictor + PathAgent	42	29	17.6

Table 3: Ablation study on different combinations of waypoint predictors and high-level VLN agents.

regions. We conduct an additional ablation without these definitions, but instead using a data style similar to MP3D, selecting waypoints that are 2-3 meters away. This results in a 4% reduction in SR and a 2.2% reduction in SPL.

Another key factor in our VAP is the use of navigational affordances. We provide an experiment where we do not use SAM to segment the affordances, but directly scatter some points for the LLMs to select. The experiment shows that this greatly increases the difficulty for Gemini to perform low-level motion planning, as there are more options, and some points are located on obstacles, causing interference. This lowers SR from 27% to 15%.

**Different Combinations of Low-Level Predictors and High-Level Agents** Our AO-Planner defaults to using a zero-shot waypoint predictor in the first stage and our PathAgent in the second stage. In this experiment, we compare different implementations of the two stages. In Table 3, “Baseline Agent” refers to transferring MapGPT (Chen et al. 2024b) from a discrete environment to the VLN-CE task. It does not visualize the waypoint predictions in the observation or calculate the position of the waypoint via 3D mapping, but simply selects a view and moves a fixed distance each time. Our PathAgent shows a 6-7% performance improvement in SR compared to the “Baseline Agent”.

Besides, We explore replacing our zero-shot waypoint predictor with a supervised waypoint predictor. Our proposed

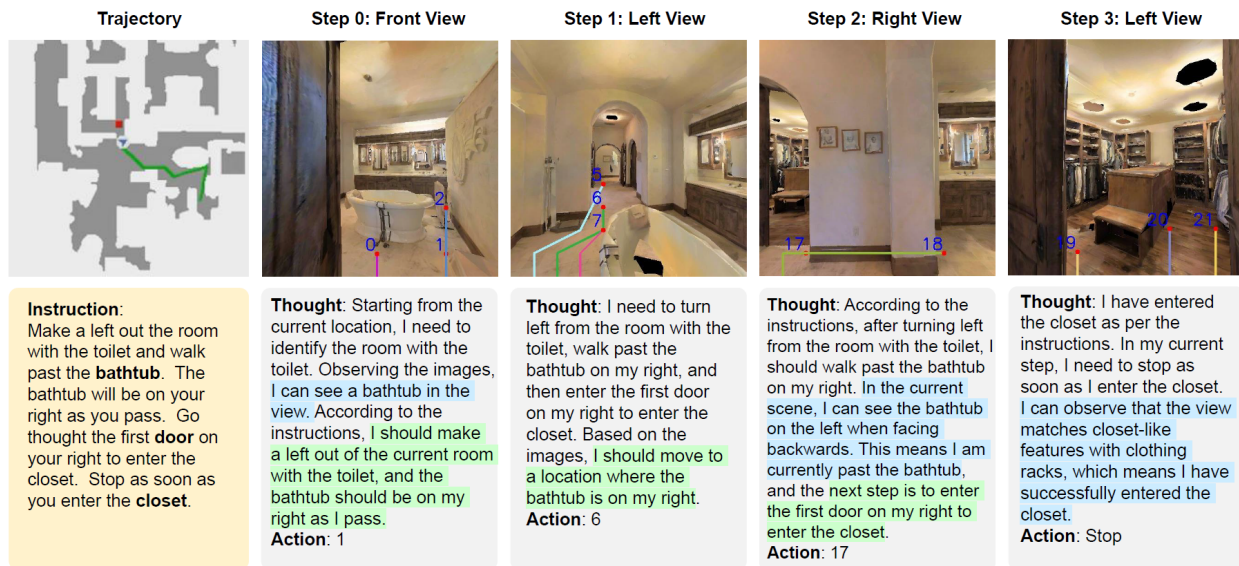


Figure 4: An example of successful navigation in a continuous environment. We present visualizations of the motion planning results from the low-level agent (upper) and the thinking process of the high-level PathAgent based on visualized candidate paths (only the selected directions are shown in this figure) and the selection of a path ID as an action (bottom). The agent ultimately decides to stop after observing the target “closet”.

LLMs		Metrics		
Low-Level	High-Level	OSR $\uparrow$	SR $\uparrow$	SPL $\uparrow$
GPT-4o	Gemini-1.5-Pro	23	11	6.9
Gemini-1.5-Pro	Gemini-1.5-Pro	25	14	9.1
GPT-4o	GPT-4o	34	16	9.4
Gemini-1.5-Pro	GPT-4o	39	27	16.9

Table 4: Ablation study on using different LLMs.

zero-shot predictor does not rely on any ground-truth from the MP3D simulator for training and performs on par with (16.9% vs. 17.6% in SPL) the supervised waypoint predictor used by DC-VLN (Hong et al. 2022). This demonstrates that our visual affordances prompting is a promising new approach in zero-shot continuous VLN.

**Agents using Different LLMs** Our zero-shot AO-Planner is a two-stage framework. As shown in Table 4, we use different LLMs for low-level and high-level agents, respectively, to observe their impact on performance. Gemini-1.5 has demonstrated impressive performance on some embodied tasks (Reid et al. 2024), such as OpenEQA (Majumdar et al. 2024). In our low-level motion planning task, we find that Gemini-1.5-Pro outperforms OpenAI’s GPT-4o model. However, contrary to the results of low-level ablation, Gemini-1.5 performs poorly in high-level decision-making. This could be attributed to differences in the training data of the two models, which have resulted in distinct distributions of capabilities. Investigating this underlying cause further will be our future work. Based on the aforementioned experimental results, we default to using Gemini-1.5 as the LLM for the low-level agent and GPT-4o for the high-level agent.

## Qualitative Results

As shown in Figure 4, we present a comprehensive example of navigation in a continuous environment. It can be observed that the low-level motion planning agent provides some high-quality candidates. For instance, in the left view at step 1, the bathtub blocks most of the forward direction, but the low-level agent finds three paths that navigate around the obstacle and reach the distant waypoints. The trajectory in the top-down map also demonstrates that the agent successfully avoids the bathtub obstacle and prevents collisions.

The high-level PathAgent consistently selects the correct directions at each step, and the output thoughts also demonstrate its understanding of the task. Upon encountering landmarks that exist in the instruction, the agent promptly assesses the progress of instruction execution and makes accurate decisions until reaching the goal.

## Conclusion

This paper proposes AO-Planner, where we explore the previously under-researched low-level motion planning capability of LLM for addressing the challenging VLN-CE task, effectively bridging the gap between LLM and 3D world. We introduce a novel visual affordances prompting strategy, which provides navigational affordances to enable LLM to plan low-level motions. A high-level PathAgent is further proposed to perform path selection in the RGB space by comprehending all environmental information. Finally, we map the selected pixel-based path to 3D coordinates for navigation. Experiments on R2R-CE and RxR-CE demonstrate the effectiveness of our approach that achieves state-of-the-art zero-shot and competitive supervised performance.

## Acknowledgements

This work was supported in part by National Natural Science Foundation of China (NSFC) under Grants No.62476293, National Science and Technology Ministry Youth Talent Funding No. 2022WRQB002, Shenzhen Science and Technology Program (Grant No. GJHZ20220913142600001).

## References

- An, D.; Qi, Y.; Li, Y.; Huang, Y.; Wang, L.; Tan, T.; and Shao, J. 2022. BEVBert: Topo-Metric Map Pre-training for Language-guided Navigation. *arXiv preprint arXiv:2212.04385*.
- An, D.; Wang, H.; Wang, W.; Wang, Z.; Huang, Y.; He, K.; and Wang, L. 2024. Etpnav: Evolving topological planning for vision-language navigation in continuous environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and Van Den Hengel, A. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3674–3683.
- Anil, R.; Dai, A. M.; Firat, O.; Johnson, M.; Lepikhin, D.; Passos, A.; Shakeri, S.; Taropa, E.; Bailey, P.; Chen, Z.; et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.
- Cai, W.; Huang, S.; Cheng, G.; Long, Y.; Gao, P.; Sun, C.; and Dong, H. 2023. Bridging zero-shot object navigation and foundation models through pixel-guided navigation skill. *arXiv preprint arXiv:2309.10309*.
- Chang, A.; Dai, A.; Funkhouser, T.; Halber, M.; Niebner, M.; Savva, M.; Song, S.; Zeng, A.; and Zhang, Y. 2017. Matterport3D: Learning from RGB-D Data in Indoor Environments. In *2017 International Conference on 3D Vision (3DV)*, 667–676. IEEE.
- Chen, J.; Lin, B.; Liu, X.; Ma, L.; Liang, X.; and Wong, K.-Y. K. 2024a. Affordances-Oriented Planning using Foundation Models for Continuous Vision-Language Navigation. *arXiv preprint arXiv:2407.05890*.
- Chen, J.; Lin, B.; Xu, R.; Chai, Z.; Liang, X.; and Wong, K.-Y. K. 2024b. 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*.
- Chen, P.; Sun, X.; Zhi, H.; Zeng, R.; Li, T. H.; Liu, G.; Tan, M.; and Gan, C. 2023. A2Nav: Action-Aware Zero-Shot Robot Navigation by Exploiting Vision-and-Language Ability of Foundation Models. *arXiv preprint arXiv:2308.07997*.
- Chen, S.; Guhur, P.-L.; Schmid, C.; and Laptev, I. 2021. History aware multimodal transformer for vision-and-language navigation. *Advances in Neural Information Processing Systems*, 34: 5834–5847.
- Chen, S.; Guhur, P.-L.; Tapaswi, M.; Schmid, C.; and Laptev, I. 2022. Think Global, Act Local: Dual-scale Graph Transformer for Vision-and-Language Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16537–16547.
- Deng, Z.; Narasimhan, K.; and Russakovsky, O. 2020. Evolving graphical planner: Contextual global planning for vision-and-language navigation. *Advances in Neural Information Processing Systems*, 33: 20660–20672.
- Dorbala, V. S.; Sigurdsson, G.; Piramuthu, R.; Thomason, J.; and Sukhatme, G. S. 2022. Clip-nav: Using clip for zero-shot vision-and-language navigation. *arXiv preprint arXiv:2211.16649*.
- Gadre, S. Y.; Wortsman, M.; Ilharco, G.; Schmidt, L.; and Song, S. 2022. CLIP on Wheels: Zero-Shot Object Navigation as Object Localization and Exploration. *ArXiv*, abs/2203.10421.
- Gibson, J. J. 2014. *The ecological approach to visual perception: classic edition*. Psychology press.
- Hong, Y.; Wang, Z.; Wu, Q.; and Gould, S. 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*, 15439–15449.
- Hong, Y.; Wu, Q.; Qi, Y.; Rodriguez-Opazo, C.; and Gould, S. 2021. vlnbert: A recurrent vision-and-language bert for navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1643–1653.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. *arXiv:2304.02643*.
- Krantz, J.; Gokaslan, A.; Batra, D.; Lee, S.; and Maksymets, O. 2021. Waypoint models for instruction-guided navigation in continuous environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15162–15171.
- Krantz, J.; and Lee, S. 2022. Sim-2-sim transfer for vision-and-language navigation in continuous environments. In *European Conference on Computer Vision*, 588–603. Springer.
- Krantz, J.; Wijmans, E.; Majumdar, A.; Batra, D.; and Lee, S. 2020. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*, 104–120. Springer.
- Ku, A.; Anderson, P.; Patel, R.; Ie, E.; and Baldrige, J. 2020. Room-Across-Room: Multilingual Vision-and-Language Navigation with Dense Spatiotemporal Grounding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 4392–4412.
- Lei, X.; Yang, Z.; Chen, X.; Li, P.; and Liu, Y. 2024. Scaffolding Coordinates to Promote Vision-Language Coordination in Large Multi-Modal Models. *arXiv preprint arXiv:2402.12058*.
- Lin, B.; Nie, Y.; Wei, Z.; Chen, J.; Ma, S.; Han, J.; Xu, H.; Chang, X.; and Liang, X. 2024. NavCoT: Boosting LLM-Based Vision-and-Language Navigation via Learning Disentangled Reasoning. *arXiv preprint arXiv:2403.07376*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023. Grounding dino:

- Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Long, Y.; Li, X.; Cai, W.; and Dong, H. 2023. Discuss Before Moving: Visual Language Navigation via Multi-expert Discussions. *arXiv preprint arXiv:2309.11382*.
- Luddecke, T.; and Worgotter, F. 2017. Learning to segment affordances. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 769–776.
- Ma, C.-Y.; Lu, J.; Wu, Z.; AlRegib, G.; Kira, Z.; Socher, R.; and Xiong, C. 2019. Self-monitoring navigation agent via auxiliary progress estimation. *arXiv preprint arXiv:1901.03035*.
- Majumdar, A.; Aggarwal, G.; Devnani, B.; Hoffman, J.; and Batra, D. 2022. Zson: Zero-shot object-goal navigation using multimodal goal embeddings. *arXiv preprint arXiv:2206.12403*.
- Majumdar, A.; Ajay, A.; Zhang, X.; Putta, P.; Yenamandra, S.; Henaff, M.; Silwal, S.; Mcvay, P.; Maksymets, O.; Arnaud, S.; et al. 2024. Openeqa: Embodied question answering in the era of foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16488–16498.
- Nasiriany, S.; Xia, F.; Yu, W.; Xiao, T.; Liang, J.; Dasgupta, I.; Xie, A.; Driess, D.; Wahid, A.; Xu, Z.; et al. 2024. PIVOT: Iterative Visual Prompting Elicits Actionable Knowledge for VLMs. *arXiv preprint arXiv:2402.07872*.
- OpenAI. 2023a. GPT-4 Technical Report.
- OpenAI. 2023b. Gpt-4v(ision) technical work and authors.
- Pan, B.; Panda, R.; Jin, S.; Feris, R.; Oliva, A.; Isola, P.; and Kim, Y. 2023. LangNav: Language as a Perceptual Representation for Navigation. *arXiv preprint arXiv:2310.07889*.
- Qi, W.; Mullanpudi, R. T.; Gupta, S.; and Ramanan, D. 2020a. Learning to move with affordance maps. *arXiv preprint arXiv:2001.02364*.
- Qi, Y.; Pan, Z.; Zhang, S.; Hengel, A. v. d.; and Wu, Q. 2020b. Object-and-action aware model for visual language navigation. In *European Conference on Computer Vision*, 303–317. Springer.
- Qiao, Y.; Qi, Y.; Hong, Y.; Yu, Z.; Wang, P.; and Wu, Q. 2022. HOP: History-and-Order Aware Pre-training for Vision-and-Language Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15418–15427.
- Raychaudhuri, S.; Wani, S.; Patel, S.; Jain, U.; and Chang, A. 2021. Language-Aligned Waypoint (LAW) Supervision for Vision-and-Language Navigation in Continuous Environments. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Reid, M.; Savinov, N.; Teplyashin, D.; Lepikhin, D.; Lillcrap, T.; Alayrac, J.-b.; Soricut, R.; Lazaridou, A.; Firat, O.; Schrittwieser, J.; et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Ren, T.; Liu, S.; Zeng, A.; Lin, J.; Li, K.; Cao, H.; Chen, J.; Huang, X.; Chen, Y.; Yan, F.; Zeng, Z.; Zhang, H.; Li, F.; Yang, J.; Li, H.; Jiang, Q.; and Zhang, L. 2024. Grounded SAM: Assembling Open-World Models for Diverse Visual Tasks. *arXiv preprint arXiv:2401.14159*.
- Savva, M.; Kadian, A.; Maksymets, O.; Zhao, Y.; Wijmans, E.; Jain, B.; Straub, J.; Liu, J.; Koltun, V.; Malik, J.; et al. 2019. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9339–9347.
- Team, G.; Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, H.; Liang, W.; Van Gool, L.; and Wang, W. 2023a. Dreamwalker: Mental planning for continuous vision-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10873–10883.
- Wang, X.; Huang, Q.; Celikyilmaz, A.; Gao, J.; Shen, D.; Wang, Y.-F.; Wang, W. Y.; and Zhang, L. 2019. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6629–6638.
- Wang, Z.; Li, J.; Hong, Y.; Wang, Y.; Wu, Q.; Bansal, M.; Gould, S.; Tan, H.; and Qiao, Y. 2023b. Scaling data generation in vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12009–12020.
- Wang, Z.; Li, X.; Yang, J.; Liu, Y.; Hu, J.; Jiang, M.; and Jiang, S. 2024. Lookahead Exploration with Neural Radiance Representation for Continuous Vision-Language Navigation. *arXiv preprint arXiv:2404.01943*.
- Yang, J.; Zhang, H.; Li, F.; Zou, X.; Li, C.; and Gao, J. 2023a. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*.
- Yang, Z.; Li, L.; Lin, K.; Wang, J.; Lin, C.-C.; Liu, Z.; and Wang, L. 2023b. The dawn of llms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9(1): 1.
- Zhang, J.; Wang, K.; Xu, R.; Zhou, G.; Hong, Y.; Fang, X.; Wu, Q.; Zhang, Z.; and He, W. 2024. NaVid: Video-based VLM Plans the Next Step for Vision-and-Language Navigation. *arXiv preprint arXiv:2402.15852*.
- Zheng, D.; Huang, S.; Zhao, L.; Zhong, Y.; and Wang, L. 2024. Towards learning a generalist model for embodied navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13624–13634.
- Zhou, G.; Hong, Y.; and Wu, Q. 2023. NavGPT: Explicit Reasoning in Vision-and-Language Navigation with Large Language Models. *arXiv preprint arXiv:2305.16986*.
- Zubair Irshad, M.; Chowdhury Mithun, N.; Seymour, Z.; Chiu, H.-P.; Samarasekera, S.; and Kumar, R. 2021. SASRA: Semantically-aware Spatio-temporal Reasoning Agent for Vision-and-Language Navigation in Continuous Environments. *arXiv e-prints*, arXiv–2108.