

# VPN: Visual Prompt Navigation

Shuo Feng<sup>1,2,\*†</sup>, Zihan Wang<sup>3\*</sup>, Yuchen Li<sup>5</sup>, Rui Kong<sup>5</sup>, Hengyi Cai<sup>5</sup>, Shuaiqiang Wang<sup>5</sup>,  
Gim Hee Lee<sup>3</sup>, Piji Li<sup>1,2,‡</sup>, Shuqiang Jiang<sup>4</sup>

<sup>1</sup>College of Artificial Intelligence, Nanjing University of Aeronautics and Astronautics

<sup>2</sup>The Key Laboratory of Brain-Machine Intelligence Technology, Ministry of Education

<sup>3</sup>School of Computing, National University of Singapore

<sup>4</sup>University of Chinese Academy of Sciences, Beijing

<sup>5</sup>Baidu Inc.

{fengshuo, pjli}@nuaa.edu.cn, zihan.wang@u.nus.edu

## Abstract

While natural language is commonly used to guide embodied agents, the inherent ambiguity and verbosity of language often hinder the effectiveness of language-guided navigation in complex environments. To this end, we propose Visual Prompt Navigation (VPN), a novel paradigm that guides agents to navigate using only user-provided visual prompts within 2D top-view maps. This visual prompt primarily focuses on marking the visual navigation trajectory on a top-down view of a scene, offering intuitive and spatially grounded guidance without relying on language instructions. It is more friendly for non-expert users and reduces interpretive ambiguity. We build VPN tasks in both discrete and continuous navigation settings, constructing two new datasets, R2R-VP and R2R-CE-VP, by extending existing R2R and R2R-CE episodes with corresponding visual prompts. Furthermore, we introduce VPNet, a dedicated baseline network to handle the VPN tasks, with two data augmentation strategies: view-level augmentation (altering initial headings and prompt orientations) and trajectory-level augmentation (incorporating diverse trajectories from large-scale 3D scenes), to enhance navigation performance. Extensive experiments evaluate how visual prompt forms, top-view map formats, and data augmentation strategies affect the performance of visual prompt navigation.

**Code** — <https://github.com/farlit/VPN>

## Introduction

Developing robots that can navigate based on human instructions has been a continuous research focus in artificial intelligence (AI) and robotics. Towards this goal, the research community has developed successive navigation paradigms spanning from PointGoal navigation (PointNav) (Gupta et al. 2017) and ImageGoal navigation (ImageNav) (Zhu et al. 2017) to more complex formulations including object navigation (ObjectNav) (Savva et al. 2017), vision-and-language navigation (VLN) (Anderson et al.

\*These authors contributed equally.

†Work done as an Intern at Baidu Inc.

‡Corresponding author.



Figure 1: Illustration of three types of visual navigation tasks. Compared to both language-based and image-goal instructions, visual prompt instructions provides clearer and more interpretable guidance.

2018b). Despite significant progress in these visual navigation paradigms, they still have some limitations.

As Figure 1 illustrates, natural language often falls short in precisely specifying navigation trajectories. This is particularly evident when describing object positions, directional transitions, or distance relationships (Chang et al. 2024). Furthermore, the pursuit of linguistic precision through exhaustive description inevitably leads to prohibitive verbosity. This exposes an inherent dilemma in human-robot interaction: linguistic precision often comes at the cost of excessive verbosity (Bonarini 2020; Wang et al. 2024a). However, overly verbose language instructions often impose high demands on users and result in poor user experiences. In contrast to language instructions, image goals offer more direct and precise navigation targets. However, they suffer from two key limitations: 1) Lack of intermediate cues. The absence of intermediate navigation cues makes it difficult for

agents to plan the correct route. 2) Impractical user demands. The requirement to pre-capture a large number of goal images as various destination candidates imposes impractical demands on usability.

Motivated by these findings, we propose a new navigation paradigm, called **Visual Prompt Navigation (VPN)**, which guides robots in 3D environments using user-provided visual prompts within 2D top-view maps for navigation. Visual prompt navigation aims to explore a more user-friendly interaction paradigm between users and navigation agents, providing navigation commands by simple and intuitive visual prompts. It offers several key benefits: 1) High user accessibility, allowing non-expert users to naturally specify navigation targets through clicking or drawing navigation trajectories on a visualized interface (e.g., top-view image). 2) Rich spatial information, as the top-view map with visual prompts inherently preserves complete spatial layouts and contains rich spatiotemporal trajectory descriptions, providing robots geometry-aware guidance, as illustrated in Figure 1 (c). 3) High reusability, since such top-view maps can be acquired through aerial photography by drone or multi-view 3D scene reconstruction, and only need to be constructed once for repeated use across different navigation episodes. We believe VPN offers an efficient and user-friendly solution for robot navigation.

To verify the feasibility of using visual prompts as agent navigation guidance, we explore two navigation settings: the discrete VPN setting and the continuous VPN setting (VPN-CE), transferred from the VLN benchmarks: Room-to-Room (R2R) (Anderson et al. 2018b) and R2R-CE (Krantz et al. 2020) datasets, respectively. In both settings, we design visual prompts within 2D top-view maps, where users can simply mark a sequence of key navigation waypoints, which are then connected by arrows to indicate the expected navigation trajectory. Based on such visual prompts, we construct two new datasets: R2R with visual prompts (R2R-VP) and R2R-CE with visual prompts (R2R-CE-VP), which serve as standardized benchmarks for VPN tasks in discrete and continuous environments.

To tackle the VPN task, we propose a new method called **Visual Prompt Navigation Network (VPNet)**. In particular, we implement a discrete navigation setting of VPNet based on DUET (Chen et al. 2022b), and a continuous navigation setting based on ETPNav (An et al. 2024), where the original instruction encoders in both models are replaced with a ViT-based (Dosovitskiy et al. 2020) encoder to encode top-view maps with visual prompts. Furthermore, we introduce two data augmentation strategies to enhance navigation performance: 1) View-level augmentation, which increases view diversity by randomly altering the agent’s initial heading and rotating the visual prompt maps. 2) Trajectory-level augmentation, which increases trajectory diversity by incorporating additional trajectories from both MP3D (Chang et al. 2017) scenes and HM3D (Ramakrishnan et al. 2021) scenes. Besides, we validate how visual prompt forms, top-view map formats, and data augmentation strategies affect VPN performance.

In summary, we make the following contributions:

- We propose a new navigation task, Visual Prompt Naviga-

tion (VPN), which explores guiding the navigation agent based on visual prompts, and construct two corresponding datasets, R2R-VP and R2R-CE-VP, as standardized benchmarks in discrete and continuous settings.

- We introduce VPNet, a baseline model designed for the VPN task, and incorporate two data augmentation strategies, view-level augmentation and trajectory-level augmentation, to enhance navigation performance.
- Extensive experiments evaluate how visual prompt forms, top-view map formats, and data augmentation strategies affect the performance of visual prompt navigation.

## Related Work

As VPN is a newly introduced task, we provide a brief overview of several closely related studies in the areas of visual navigation and visual prompt.

**Visual Navigation** requires the agent to utilize visual information to reach a specified destination. Based on the form of input instructions, it can be broadly categorized into five types: PointGoal navigation (PointNav) (Anderson et al. 2018a), object navigation (ObjNav) (Savva et al. 2017), vision-and-language navigation (VLN) (Anderson et al. 2018b), visual audio navigation (VAN) (Chen et al. 2020) and ImageGoal navigation (ImageNav) (Zhu et al. 2017). Besides, FloNa (Li et al. 2025a) leverages floor plan outlines to guide agent navigation, but it heavily relies on accurate location information of both the agent and the destination, similar to PointNav. In PointNav, the agent is provided with the relative direction and distance to the goal. In contrast, our proposed VPN guides navigation solely via visual prompts within 2D top-view maps. This design can be seen as a hybrid of VLN and ImageNav: it combines the spatiotemporal trajectory descriptions characteristic of VLN with the intuitive visual presentation style of ImageNav.

**Visual Prompt** (Wu et al. 2024) has recently emerged as a new paradigm that complements textual prompts by enabling more fine-grained, pixel-level instruction. Building on this idea, several studies have introduced symbolic elements, such as numbers, sketches, and arrows, to enhance natural language and improve the accuracy of Visual Question Answering (VQA) (Yang et al. 2023; Cai et al. 2024). Furthermore, recent works have explored the role of visual prompts in embodied AI (EAI). One line of work leverages generated visual prompts to improve agents’ planning capabilities, including robotic manipulation (Lee et al. 2024; Liu et al. 2024) and continuous VLN (Chen et al. 2025). However, they still rely primarily on natural language input, with visual prompts serving only a supplementary role. Another line of work, represented solely by RoVI (Li et al. 2025b), uses hand-drawn symbolic representations as the sole form of instruction for robotic manipulation. While the above work has demonstrated the role of visual prompts in EAI, using visual prompts as the sole instruction input in visual navigation remains unexplored. Unlike the relatively static scenes in robotic manipulation, the scenes in visual navigation change continuously as the agent moves. Therefore, in this work, we are the first to adopt a compact scene representation using 2D top-view maps, where visual prompts are




Figure 2: Illustration of the process of constructing visual prompts on 2D top-view map. The four subfigures are labeled (a)–(d) from left to right, top to bottom.

overlaid to serve as the sole guidance for agent navigation, without relying on any textual information.

### VPN Dataset

We transfer the VLN datasets Room-to-Room (R2R) and R2R-CE to VPN tasks by replacing language instructions with visual prompts in their navigation episodes, building two new datasets for VPN: R2R-VP and R2R-CE-VP, corresponding to the discrete and continuous environments, respectively. In this section, we detail the dataset construction process and provide an analysis of these datasets.

### Dataset Construction

As Figure 2 illustrates, we construct the R2R-VP and R2R-CE-VP datasets through a four-step process: a) For each episode in R2R and R2R-CE, we generate a 2D top-view map of the corresponding scene from 90 MP3D (Chang et al. 2017) scenes, following a procedure similar to LED (Hahn et al. 2020). b) We then compute the pixel coordinates of all viewpoints on the top-view map and connect them sequentially with arrows to represent the agent’s navigation trajectory. c) Next, we apply a center-cropping operation to extract a square region centered on the trajectory. The side length is determined by the greater of the trajectory’s width or height, plus an additional 60-pixel margin, ensuring the full trajectory is included while maintaining spatial focus. d) Finally, to remove unnecessary black borders and tightly fit the visual prompts, we detect the bounding box of all non-zero pixels and crop accordingly. This process yields our two new datasets: R2R-VP and R2R-CE-VP. To further improve generalization and coverage, we enrich the training datasets by incorporating additional in-domain episodes from PREVALENT (Hao et al. 2020) (based on MP3D scenes) and out-of-domain episodes from ScaleVLN (Wang et al. 2023a) (based on HM3D (Ramakrishnan et al. 2021) scenes).

### Dataset Analysis

Table 1 summarizes the training dataset statistics for both VLN and VPN tasks under discrete and continuous navigation settings. It includes R2R, PREVALENT (constructed

Dataset	Setting	VLN		VPN	
		Scenes	Episodes	Scenes	Episodes
R2R	Discrete	61	4,675	61	4,638
	Continuous	61	3,603	60	3,597
PRE	Discrete	60	178,270	60	177,134
	Continuous	-	-	59	88,694
SCA	Discrete	1289	4,941,710	523	1,600,945

Table 1: Statistics of training datasets for VLN and VPN tasks under discrete and continuous navigation settings. “PRE” denotes the PREVALENT dataset. “SCA” denotes the ScaleVLN dataset.

from MP3D scenes), and ScaleVLN (built from both HM3D and Gibson (Xia et al. 2018) scenes). In the ScaleVLN dataset used for VPN, we retain 523 scenes from HM3D for training. In general, certain entries under VPN contain fewer scenes or episodes compared to VLN, mainly due to two reasons: 1) Regarding scenes, some scenes are excluded due to poor reconstruction quality in their corresponding top-down maps, which makes them unsuitable for annotating reliable visual prompts. For MP3D, where each scene consists of multiple floors with separate top-down views, we manually exclude floors with poor reconstruction. For the HM3D scenes, we retain only the first floor, which generally has the best reconstruction quality. We apply the above strategies to transfer the VLN datasets to visual prompt navigation, including R2R for both continuous and discrete settings, as well as ScaleVLN and PREVALENT for the discrete setting. 2) Regarding episodes, some episodes are excluded because certain viewpoints cannot be mapped to MP3D’s 3D meshes (Krantz et al. 2020) under the continuous navigation setting. As a result, only valid viewpoints from R2R-CE are retained, and any episodes with invalid viewpoints in the PREVALENT dataset are discarded.

### VPNet

**Problem Formulation.** In the discrete navigation setting of VPN, the environment is an undirected navigation connectivity graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  denotes navigable nodes and  $\mathcal{E}$  represents the edges indicating connectivity between nodes. Starting from an initial node and guided by visual prompts, the agent is required to explore the graph  $\mathcal{G}$  and navigate to the target node. The visual prompt embeddings  $\mathcal{P} = \{vp_i\}_{i=1}^m$  are encoded by a ViT-based (Dosovitskiy et al. 2020) encoder. At each time step  $t$ , the agent observes panoramic views  $V_t = \{v_{t,i}\}_{i=1}^n$ . The agent is aware of a few navigable views  $\mathcal{N}(V_t) \subseteq V_t$ , each corresponding to a neighboring node and its spatial coordinates.

VPN in continuous environments (VPN-CE) is established over Habitat (Savva et al. 2019) simulator. At each navigation step, we employ a pretrained waypoint predictor (Hong et al. 2022) to generate navigable waypoints in continuous environments, thereby aligning the task formulation with that of VPN in discrete environments.

**Overview.** We propose the Visual Prompt Navigation Network (VPNet) to address the newly introduced VPN task. Specifically, VPNet is based on DUET (Chen et al. 2022b) in

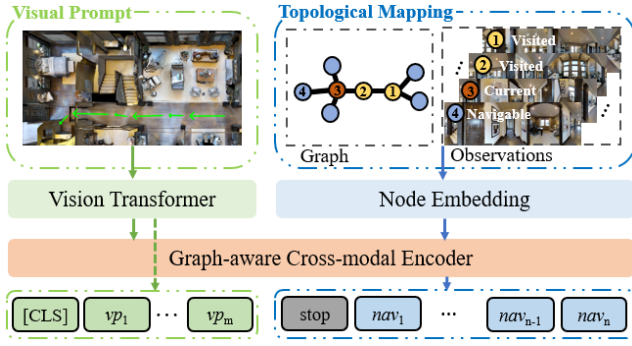


Figure 3: Illustration of VPNet. “ $vp_m$ ” denotes token corresponding to visual prompts. “stop” indicates the token for the stop action, and “ $nav_n$ ” represents token corresponding to navigable candidate.

the discrete navigation setting and ETPNav (An et al. 2024) in the continuous navigation setting. The architecture of VPNet, as illustrated in Figure 3, comprises three main components: a ViT-based Visual Prompt Encoder, a Node Embedding module, and a Graph-aware Cross-modal Encoder.

### Visual Prompt Processing

To ensure consistent input size, all 2D top-view maps with visual prompts are resized to  $224 \times 224$ . To account for floor transitions caused by actions like ascending or descending stairs, we introduce an Order-Aware Floor Concatenation (O AFC), which concatenates visual prompt features while preserving their traversal order. Among 90 MP3D scenes, 32 have one floor, while the rest span multiple: 32 (2 floors), 17 (3), 8 (4), and 1 (5). In contrast, all 523 HM3D scenes used for training are single-floor. The final prompt feature extraction is defined as:

$$\mathcal{P}_i^o = \text{ViT}(\mathcal{P}_i) + b_i \quad (1)$$

$$\mathcal{P} = [\mathcal{P}_1^o, \mathcal{P}_2^o, \dots, \mathcal{P}_k^o] \quad (2)$$

where  $\mathcal{P}_i$  represents the  $i$ -th input visual prompt map,  $b_i$  denotes the order embedding encoding the traversal sequence of floors, and  $k$  denotes the total number of input visual prompt maps corresponding to the navigation trajectory.

### Topological Mapping

Since the agent has no prior knowledge of the environment graph  $\mathcal{G}$ , it incrementally constructs a topological graph based on observations gathered during navigation. At each time step  $t$ , the graph  $\mathcal{G}_t = \{\mathcal{V}_t, \mathcal{E}_t \mid \mathcal{V}_t \subseteq \mathcal{V}, \mathcal{E}_t \subseteq \mathcal{E}\}$  comprises three types of nodes: the current node, visited nodes, and navigable nodes. The agent accesses panoramic views at current and visited nodes, and partial views of navigable nodes observed from current or previously visited locations. The current node  $V_t$  and its neighboring unvisited nodes  $\mathcal{N}(V_t)$  are added to  $\mathcal{V}_{t-1}$  to update  $\mathcal{V}_t$ .

**Node Embedding.** Each node in the current and visited set is represented by aggregating the features of its panoramic views. Each view  $W_{ij}^n$  is computed as the sum of three components: the visual features  $W_{ij}^v$ , a direction embedding  $W_{ij}^d$

that encodes the relative orientation between the view and the agent’s egocentric perspective, and a navigability type embedding  $W_{ij}^n$  indicating whether the view leads to at least one adjacent, traversable node. The panoramic view encoder (a two-layer transformer) is then applied over all views of the node to capture inter-view contextual information, resulting in the contextual view representations  $W_i$ :

$$W_{ij}' = \text{LN}(W_{ij}^v + W_{ij}^d + W_{ij}^n) \quad (3)$$

$$W_i = \text{SelfAttn}(W_{ij}') \quad (4)$$

where  $i$  indexes the node in  $\mathcal{V}_t$ ,  $j$  indexes the view within that node,  $\text{LN}(\cdot)$  and  $\text{SelfAttn}(\cdot)$  denote layer normalization and the self-attention mechanism. For current and visited nodes,  $W_i = (W_{i1}, W_{i2}, \dots, W_{in})$  denotes the contextual representations of all  $n$  panoramic views. For navigable nodes, which are only partially observed from visited or current locations,  $W_i$  is constructed using the available subset of views with valid visual features  $W_{ij}^v$ . The current node is represented by averaging its local view features. Visited nodes retain their previously computed features. Navigable nodes are represented by averaging all partially observed views collected throughout the trajectory. In addition to view-based features, each node is further enriched with a step embedding  $\mathcal{V}_{ti}^s$  encoding the agent’s visitation order, and a position embedding  $\mathcal{V}_{ti}^p$  capturing spatial distance and orientation relative to other nodes. For unexplored nodes, the step embedding is set to 0. A special “stop” node is added to the graph to represent the stop action. The final node embedding is computed as:

$$\mathcal{V}_{ti} = \mathcal{V}_{ti}^s + \mathcal{V}_{ti}^p + \frac{1}{n} \sum_{j=1}^n W_{ij} \quad (5)$$

where  $\mathcal{V}_{ti}$  is the  $i$ -th node in  $\mathcal{V}_t = (\mathcal{V}_{t1}, \mathcal{V}_{t2}, \dots, \mathcal{V}_{tv})$ , and  $v$  is the total number of nodes at time  $t$ .

### Graph-Aware Navigation Policy

**Graph-Aware Cross-Modal Encoder.** After obtaining the visual prompt embeddings  $\mathcal{P}$  and node embeddings  $\mathcal{V}_t$ , we feed them into a multi-layer cross-modal graph transformer to model their interactions. Specifically, each transformer layer comprises a cross-attention layer to model the relations between visual prompts  $\mathcal{P}$  and nodes  $\mathcal{V}_t$ , followed by a graph-aware self-attention (GASA) layer, which considers both visual similarity and spatial distances between nodes to obtain graph-aware node representations:

$$X = \text{CrosAttn}(\mathcal{V}_t, \mathcal{P}) \quad (6)$$

$$\text{GASA}(X) = \text{Softmax}\left(\frac{XW_q(XW_k)^T}{\sqrt{d}} + EW_d\right)XW_v \quad (7)$$

where  $\text{CrosAttn}(\cdot)$  denotes the cross-attention mechanism,  $X$  represents the stack of all node embeddings refined via cross-attention,  $E$  is the pair-wise distance matrix obtained from  $\mathcal{E}_t$ , and  $W_q, W_k, W_d, W_v$  are learnable matrices. The final visual-prompt aware representations of nodes are formulated as  $\hat{v} = (\text{stop}, nav_1, nav_2, \dots, nav_v)$ , where  $v$  indicates the total number of nodes in  $\mathcal{V}_t$ .



Figure 4: Illustration of prompt-view augmentation (left) and agent-view augmentation (right). The left side shows rotated 2D top-view maps with visual prompts at  $0$ ,  $\frac{\pi}{2}$ ,  $\pi$ , and  $\frac{3\pi}{2}$ . The right side shows the corresponding first-person observations when the agent starts with  $0$ ,  $\frac{\pi}{2}$ ,  $\pi$ , and  $\frac{3\pi}{2}$ .

**Global Action Prediction.** We apply a two-layer feed-forward network (FFN) to predict the navigation score  $s_i$  for the node representation  $\hat{v}_i$ :

$$s_i^g = \text{FFN}(\hat{v}_i) \quad (8)$$

where it is worth noting that  $s_0^g$  represents the stop score. To avoid redundant revisits and promote exploration, we mask the scores of previously visited nodes and the current node. The agent then chooses the highest-scoring candidate, either a navigable node or the “stop” node, and navigates to it via the shortest path in the topological graph  $\mathcal{G}_t$ .

### Data Augmentation

To further enhance navigation performance, we introduce two data augmentation strategies.

**Trajectory-Level Augmentation.** To enhance trajectory diversity, we introduce trajectory-level augmentation by incorporating additional navigation episodes from in-domain and out-of-domain environments. Here, in-domain refers to environments included in the R2R training set, while out-of-domain refers to those that are not. Specifically, we include in-domain trajectories from the PREVALENT dataset (Hao et al. 2020), which are based on 60 MP3D (Chang et al. 2017) scenes (Chang et al., 2017), as well as out-of-domain trajectories from the ScaleVLN dataset (Wang et al. 2023a), which are constructed using 523 HM3D (Ramakrishnan et al. 2021) scenes. This augmentation strategy enriches the training data with diverse navigation paths, thereby expanding coverage across different scene types and improving the agent’s robustness to domain shifts.

**View-Level Augmentation.** As shown in Figure 4, we introduce two view-level augmentation strategies: 1) *Prompt-view augmentation* diversifies visual prompts by randomly rotating the 2D top-view maps by  $0$ ,  $\frac{\pi}{2}$ ,  $\pi$ , or  $\frac{3\pi}{2}$ . 2) *Agent-view augmentation* enhances the diversity of initial observations by randomly sampling the agent’s heading from  $[0, 2\pi)$ . Agent-view augmentation is particularly suitable for VPN tasks, where the initial heading is independent of the visual prompt. In contrast, for VLN tasks, instructions may implicitly specify the agent’s initial orientation, making such augmentation less applicable.

## Training

Behavior cloning (BC)  $\psi$  often suffers from distribution shifts between training and testing phases. To mitigate this issue, we fine-tune the VPNet using supervision from a pseudo interactive demonstrator  $\psi^*$ , following the DAGger algorithm (Ross, Gordon, and Bagnell 2011). The final loss combines Behavior cloning and DAGger supervision:

$$\mathcal{L}_{\text{BC}} = \sum_{t=1}^T -\log p(a_t^\psi | \mathcal{P}, \mathcal{G}_t) \quad (9)$$

$$\mathcal{L}_{\text{DAG}} = \sum_{t=1}^T -\log p(a_t^{\psi^*} | \mathcal{P}, \mathcal{G}_t^*) \quad (10)$$

where  $a_t^\psi$  denotes the ground truth action,  $a_t^{\psi^*}$  represents the pseudo label, defined as the next node on the shortest path to the destination, computed over the partial graph  $\mathcal{G}_t^*$  constructed by the agent via on-policy action sampling. The overall loss is defined as  $\mathcal{L} = \lambda \mathcal{L}_{\text{BC}} + (1 - \lambda) \mathcal{L}_{\text{DAG}}$ , where  $\lambda$  is a weighting factor that balances the two objectives.

## Experiments

### Datasets and Evaluation Metrics

We train our model using visual prompt maps corresponding to trajectories selected from the training split of R2R (Anderson et al. 2018b) and R2R-CE (Krantz et al. 2020), PREVALENT (Hao et al. 2020), and ScaleVLN (Wang et al. 2023a) datasets. For details on the validation and test splits, please refer to the appendix.

### Implementation Details

**Features.** We adopt the ViT-B/16 (Dosovitskiy et al. 2020), pretrained on ImageNet (Deng et al. 2009), to extract the panoramic view features  $W_{ij}^v$  for each view in panoramas and to encode the visual prompt map  $\mathcal{P}_i$  for each navigation episode across all datasets.

**Model Architecture.** We use 12, 2, 4, and 4 transformer layers for the visual prompt encoder (ViT-B/16), panoramic view encoder, graph-aware cross-modal encoder, and local cross-modal encoder (see the appendix for details), respectively, all with hidden size 768. VPNet includes both cross-modal encoders in discrete environments, but only the graph-aware encoder is used in continuous environments. We set  $\lambda$  to 0.5 for both VPN and VPN-CE.

**Training Details.** For the VPN task, we train VPNet on the R2R-VP dataset along with additional episodes extracted from PREVALENT and ScaleVLN. VPNet is trained for 400k iterations on a single NVIDIA RTX A5000 GPU using a batch size of 10 and a learning rate of  $1.5e-5$ . For the VPN-CE task, VPNet is trained on R2R-CE-VP and episodes extracted from PREVALENT, using a total batch size of 16 across two NVIDIA RTX A5000 GPUs for 400k iterations, with a learning rate of  $1e-5$ .

Methods	R2R/R2R-VP Val Seen				R2R/R2R-VP Val Unseen				R2R/R2R-VP Test Unseen			
	NE↓	OSR↑	SR↑	SPL↑	NE↓	OSR↑	SR↑	SPL↑	NE↓	OSR↑	SR↑	SPL↑
<i>Vision-and-Language Navigation Evaluated on R2R in Discrete Environments:</i>												
Seq2Seq (Anderson et al. 2018b)	6.01	53	39	–	7.81	28	21	–	7.85	27	20	–
RCM (Wang et al. 2019)	3.53	75	67	–	6.09	50	43	–	6.12	50	43	38
PREVALENT (Hao et al. 2020)	3.67	–	69	65	4.71	–	58	53	5.30	61	54	51
VLN $\odot$ BERT (Hong et al. 2020)	2.90	–	72	68	3.93	–	63	57	4.09	70	63	57
HAMT (Chen et al. 2021b)	2.51	–	76	72	2.29	–	66	61	3.93	72	65	60
DUET (Chen et al. 2022b)	2.28	86	79	73	3.31	81	72	60	3.65	76	69	59
BEVBert (An et al. 2022)	2.17	88	81	74	2.81	84	75	64	3.13	81	73	62
GridMM (Wang et al. 2023b)	-	-	-	-	2.83	-	75	64	3.35	-	73	62
Lily (Lin et al. 2023)	-	-	-	-	2.90	-	74	62	3.44	-	72	60
ScaleVLN (Wang et al. 2023a)	2.12	87	81	75	2.09	88	81	70	2.27	86	80	70
NaviLLM (Zheng et al. 2024)	-	-	-	-	3.51	-	67	59	3.71	-	68	60
NavGPT-2 (Zhou et al. 2024)	1.78	88	81	75	3.13	81	72	61	3.35	78	71	60
VER (Liu, Wang, and Yang 2024)	-	-	-	-	2.80	-	76	65	2.74	-	76	66
MAGIC (Wang et al. 2024b)	1.73	89	84	80	2.22	86	79	70	2.75	82	77	69
SRDF (Wang et al. 2024d)	-	-	-	-	1.62	90	86	79	1.82	89	85	78
DUET (R2R+PRE)	2.28	86	79	73	3.31	81	72	60	3.65	76	69	59
DUET (R2R+PRE+SCA)	2.12	87	81	75	2.09	88	81	70	2.27	86	80	70
<i>Visual Prompt Navigation Evaluated on R2R-VP in Discrete Environments:</i>												
VPNet (R2R)	3.37	72.41	66.69	62.13	5.43	59.53	51.23	43.47	5.11	62.33	52.40	42.87
VPNet (R2R+PRE)	<b>0.05</b>	<b>100</b>	<b>100</b>	<b>99.77</b>	<u>2.18</u>	<u>76.76</u>	<u>65.92</u>	<u>56.17</u>	<u>1.94</u>	<u>78.21</u>	<u>66.38</u>	<u>56.26</u>
VPNet (R2R+PRE+SCA)	<u>0.14</u>	<u>99.41</u>	<u>99.41</u>	<u>99.08</u>	<b>0.48</b>	<b>97.45</b>	<b>96.68</b>	<b>94.84</b>	<b>0.31</b>	<b>98.56</b>	<b>97.56</b>	<b>94.60</b>

Table 2: Performance comparison between VPNet and existing VLN methods in discrete environments. “PRE” denotes the PREVALENT dataset. “SCA” denotes the ScaleVLN dataset. Best results are marked in bold and second best are underlined.

Methods	Val Seen			Val Unseen		
	NE↓	SR↑	SPL↑	NE↓	SR↑	SPL↑
<i>VLN-CE Evaluated on R2R-CE:</i>						
Seq2Seq (Krantz et al. 2020)	7.12	37	35	7.37	32	30
CMTP (Chen et al. 2021a)	7.10	36	31	7.90	26	23
WPN (Krantz et al. 2021)	5.48	46	43	6.31	36	34
LAW (Raychaudhuri et al. 2021)	6.35	40	37	6.83	35	31
CM2 (Georgakis et al. 2022)	6.10	43	35	7.02	34	28
WS-MGMap (Chen et al. 2022a)	5.65	47	43	6.28	39	34
Sim2Sim (Krantz and Lee 2022)	4.67	52	44	6.07	43	36
CMA (Hong et al. 2022)	5.20	51	45	6.20	41	36
GridMM (Wang et al. 2023b)	4.21	59	51	5.11	49	41
ETPNav (An et al. 2024)	3.95	66	59	4.71	57	49
ERG (Wang et al. 2024c)	5.04	46	42	6.20	39	35
NaVILA (Cheng et al. 2024)	-	-	-	5.22	54	49
ETPNav (R2R+PRE)	3.95	66	59	4.71	57	49
<i>VPN-CE Evaluated on R2R-CE-VP:</i>						
VPNet (R2R)	<u>4.75</u>	<u>61.24</u>	<u>55.88</u>	<u>6.14</u>	<u>42.09</u>	<u>36.16</u>
VPNet (R2R+PRE)	<b>2.30</b>	<b>84.11</b>	<b>78.15</b>	<b>6.10</b>	<b>47.96</b>	<b>41.51</b>

Table 3: Performance comparison between VPNet and existing VLN methods in continuous environments.

## Main Results

Tables 2 and 3 compare VPNet in VPN with prior VLN methods and highlight the benefits of trajectory-level data augmentation. Adding PREVALENT trajectories notably improves performance on both val seen and unseen splits,

with VPNet achieving 100% on SR in the discrete environment and 84.11% in the continuous environment (val seen). Further incorporating trajectories from ScaleVLN leads to 96.68% on SR in the discrete environment (val unseen). Notably, as shown in Table 2, VPNet in the VPN setting outperforms DUET in VLN while using only 1,600,945 trajectories from ScaleVLN, compared to DUET’s 4,941,710, as reported in Table 1. These results highlight the data efficiency of scaling within the VPN framework, which refers to its ability to achieve strong performance with less training data, and further demonstrate the effectiveness of visual prompts in guiding robot navigation.

## Ablation Study

**The Effect of Different Types of Visual Prompts.** As shown in row (e) of Table 4, the visual prompt consisting solely of lines outperforms that in row (d), which combines both lines and text. This is because the lines provide enough spatiotemporal guidance, while the text may obscure some details around the navigation waypoints. We attribute the poor performance of row (a) to the absence of trajectory-centered cropping, which causes episodes from the same scene to share identical top-view maps despite following different visual trajectories. This not only makes the VPNet prone to overfitting on the limited and unvarying top-view maps but also leads to underfitting of the crucial visual trajectory information. Interestingly, row (b), which only uses cropped top-view maps without including visual trajectories,




Figure 5: Illustration of different 2D top-view maps with visual prompts. The six subfigures are labeled (a)–(f) from left to right, top to bottom.

TVP	Val Seen			Val Unseen		
	NE↓	SR↑	SPL↑	NE↓	SR↑	SPL↑
(a)	6.86	31.68	28.63	7.07	33.94	29.36
(b)	1.47	83.56	78.64	5.51	45.83	36.78
(c)	1.04	88.29	83.37	5.13	50.18	40.54
(d)	<u>0.79</u>	<u>95.74</u>	<u>91.93</u>	<u>3.74</u>	<u>65.36</u>	<u>54.81</u>
(e)	<b>0.05</b>	<b>100</b>	<b>99.77</b>	<b>2.18</b>	<b>65.92</b>	<b>56.17</b>

Table 4: Ablation study results on different types of visual prompts in discrete environments. “TVP” denotes the types of visual prompts, corresponding to subfigures in Figure 5.

VLDA	TL↓	NE↓	OSR↑	SR↑	SPL↑
None	11.47	1.69	90.75	86.33	82.92
Agent view	12.14	1.48	92.56	88.18	85.02
Prompt view	10.04	<u>0.62</u>	<u>97.28</u>	<u>96.41</u>	<u>94.37</u>
Both	9.92	<b>0.48</b>	<b>97.45</b>	<b>96.68</b>	<b>94.84</b>

Table 5: Ablation study results on different combinations of view-level data augmentation (VLDA).

still achieves decent performance. We speculate that the navigation model may still be able to infer the approximate destination area from the cropped top-view map (similar to the image goal in the ImageNav task). However, due to the absence of explicit guidance from visual trajectories, its overall performance remains limited.

**The Effect of View-Level Data Augmentation.** As shown in Table 5, it is evident that both Prompt-view and Agent-view augmentations (Section ) lead to improvements in performance. Notably, rotating the visual prompt yields significantly greater gains compared to randomly altering the agent’s initial heading. This phenomenon remains consistent with our findings in Section that increasing the scale and diversity of top-view map data contributes significantly to improved navigation performance.




Figure 6: Illustration of visual prompt maps with different types of noise. From left to right, subfigures (a)–(c) represent: (a) the original visual prompt map, (b) the map with 20% of its pixels replaced by salt-and-pepper noise, and (c) the map with the first-step visual cue removed.

Noise type	TL↓	NE↓	OSR↑	SR↑	SPL↑
(a)	9.92	<b>0.48</b>	<b>97.45</b>	<b>96.68</b>	<b>94.84</b>
(b)	11.73	<u>1.24</u>	<u>93.53</u>	<u>90.34</u>	<u>85.40</u>
(c)	12.48	2.13	89.31	85.97	74.98

Table 6: Evaluation of VPNet under different noise types corresponding to subfigures (a)–(c) of Figure 6.

Encoder setting	TL↓	NE↓	OSR↑	SR↑	SPL↑
ViT-1k&ViT-1k	9.78	0.50	96.93	96.17	<u>94.60</u>
ViT-21k&ViT-1k	9.75	<b>0.45</b>	<u>97.13</u>	<u>96.36</u>	94.37
ViT-21k&ViT-21k	9.92	<u>0.48</u>	<b>97.45</b>	<b>96.68</b>	<b>94.84</b>

Table 7: Ablation on encoder pairs: the left of “&” encodes visual prompts, the right encodes scene observations. “ViT-21k” is pretrained on ImageNet-21k, “ViT-1k” is further fine-tuned from “ViT-21k” on its subset, ImageNet-1k.

**Noise Robustness Analysis of VPNet.** As shown in Table 6, we evaluate VPNet under different types of noise corresponding to subfigures (a)–(c) in Figure 6, simulating challenges such as low-quality top-view map reconstructions and inaccuracies in user-provided visual prompts. Although performance drops under these perturbations, VPNet, trained on clean data, still maintains strong results. This demonstrates its robustness to reconstruction-induced noise and user-induced inaccuracies.

**The Effect of Different Encoder Settings.** Table 7 suggests that broader pretraining benefits visual understanding.

## Conclusion

We present Visual Prompt Navigation using 2D top-view prompts to guide agents. We build benchmarks, propose VPNet, and validate its effectiveness in discrete and continuous settings. Future work will explore extending VPN to real-world environments.

## Acknowledgements

This research is supported by the National Natural Science Foundation of China (92464204). We sincerely thank Zun Wang for providing the code used to generate discrete view images in HM3D.

## References

- An, D.; Qi, Y.; Li, Y.; Huang, Y.; Wang, L.; Tan, T.; and Shao, J. 2022. Bevbert: Multimodal 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.; Chang, A.; Chaplot, D. S.; Dosovitskiy, A.; Gupta, S.; Koltun, V.; Kosecka, J.; Malik, J.; Mottaghi, R.; Savva, M.; et al. 2018a. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*.
- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and Van Den Hengel, A. 2018b. 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.
- Bonarini, A. 2020. Communication in human-robot interaction. *Current Robotics Reports*, 1(4): 279–285.
- Cai, M.; Liu, H.; Mustikovela, S. K.; Meyer, G. P.; Chai, Y.; Park, D.; and Lee, Y. J. 2024. Vip-llava: Making large multimodal models understand arbitrary visual prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12914–12923.
- Chang, A.; Dai, A.; Funkhouser, T.; Halber, M.; Niessner, M.; Savva, M.; Song, S.; Zeng, A.; and Zhang, Y. 2017. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*.
- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; et al. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3): 1–45.
- Chen, C.; Jain, U.; Schissler, C.; Gari, S. V. A.; Al-Halah, Z.; Ithapu, V. K.; Robinson, P.; and Grauman, K. 2020. Soundspaces: Audio-visual navigation in 3d environments. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, 17–36. Springer.
- Chen, J.; Lin, B.; Liu, X.; Ma, L.; Liang, X.; and Wong, K.-Y. K. 2025. Affordances-oriented planning using foundation models for continuous vision-language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 22, 23568–23576.
- Chen, K.; Chen, J. K.; Chuang, J.; Vázquez, M.; and Savarese, S. 2021a. Topological planning with transformers for vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11276–11286.
- Chen, P.; Ji, D.; Lin, K.; Zeng, R.; Li, T.; Tan, M.; and Gan, C. 2022a. Weakly-supervised multi-granularity map learning for vision-and-language navigation. *Advances in Neural Information Processing Systems*, 35: 38149–38161.
- Chen, S.; Guhur, P.-L.; Schmid, C.; and Laptev, I. 2021b. 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. 2022b. 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.
- Cheng, A.-C.; Ji, Y.; Yang, Z.; Gongye, Z.; Zou, X.; Kautz, J.; Bıyık, E.; Yin, H.; Liu, S.; and Wang, X. 2024. Navila: Legged robot vision-language-action model for navigation. *arXiv preprint arXiv:2412.04453*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Georgakis, G.; Schmeckpeper, K.; Wanchoo, K.; Dan, S.; Miltsakaki, E.; Roth, D.; and Daniilidis, K. 2022. Cross-modal map learning for vision and language navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15460–15470.
- Gupta, S.; Davidson, J.; Levine, S.; Sukthankar, R.; and Malik, J. 2017. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2616–2625.
- Hahn, M.; Krantz, J.; Batra, D.; Parikh, D.; Rehg, J. M.; Lee, S.; and Anderson, P. 2020. Where are you? localization from embodied dialog. *arXiv preprint arXiv:2011.08277*.
- Hao, W.; Li, C.; Li, X.; Carin, L.; and Gao, J. 2020. Towards learning a generic agent for vision-and-language navigation via pre-training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13137–13146.
- 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. 2020. A recurrent vision-and-language bert for navigation. *arXiv preprint arXiv:2011.13922*.
- 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.
- Lee, O. Y.; Xie, A.; Fang, K.; Pertsch, K.; and Finn, C. 2024. Affordance-guided reinforcement learning via visual prompting. *arXiv preprint arXiv:2407.10341*.
- Li, J.; Huang, W.; Wang, Z.; Liang, W.; Di, H.; and Liu, F. 2025a. FloNa: Floor Plan Guided Embodied Visual Navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 14, 14610–14618.
- Li, Y.; Gong, Z.; Li, H.; Huang, X.; Kang, H.; Bai, G.; and Ma, X. 2025b. Robotic visual instruction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12155–12165.
- Lin, K.; Chen, P.; Huang, D.; Li, T. H.; Tan, M.; and Gan, C. 2023. Learning vision-and-language navigation from youtube videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8317–8326.
- Liu, F.; Fang, K.; Abbeel, P.; and Levine, S. 2024. Moka: Open-vocabulary robotic manipulation through mark-based visual prompting. In *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*.
- Liu, R.; Wang, W.; and Yang, Y. 2024. Volumetric environment representation for vision-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16317–16328.
- Ramakrishnan, S. K.; Gokaslan, A.; Wijmans, E.; Maksymets, O.; Clegg, A.; Turner, J.; Undersander, E.; Galuba, W.; Westbury, A.; Chang, A. X.; et al. 2021. Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai. *arXiv preprint arXiv:2109.08238*.
- 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*, 4018–4028.
- Ross, S.; Gordon, G.; and Bagnell, D. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 627–635. JMLR Workshop and Conference Proceedings.
- Savva, M.; Chang, A. X.; Dosovitskiy, A.; Funkhouser, T.; and Koltun, V. 2017. MINOS: Multimodal indoor simulator for navigation in complex environments. *arXiv preprint arXiv:1712.03931*.
- 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.
- Wang, J.; Shi, E.; Hu, H.; Ma, C.; Liu, Y.; Wang, X.; Yao, Y.; Liu, X.; Ge, B.; and Zhang, S. 2024a. Large language models for robotics: Opportunities, challenges, and perspectives. *Journal of Automation and Intelligence*.
- Wang, L.; He, Z.; Shen, M.; Yang, J.; Liu, C.; and Chen, Q. 2024b. Magic: Meta-ability guided interactive chain-of-distillation for effective-and-efficient vision-and-language navigation. *arXiv preprint arXiv:2406.17960*.
- Wang, T.; Wu, Z.; Yao, F.; and Wang, D. 2024c. Graph-Based Environment Representation for Vision-and-Language Navigation in Continuous Environments. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8331–8335. IEEE.
- 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.; Li, S.; Li, K.; Yu, S.; Wang, Y.; Qiao, Y.; Wang, Y.; Bansal, M.; et al. 2024d. Bootstrapping Language-Guided Navigation Learning with Self-Refining Data Flywheel. *arXiv preprint arXiv:2412.08467*.
- Wang, Z.; Li, J.; Hong, Y.; Wang, Y.; Wu, Q.; Bansal, M.; Gould, S.; Tan, H.; and Qiao, Y. 2023a. 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.; and Jiang, S. 2023b. Gridmm: Grid memory map for vision-and-language navigation. In *Proceedings of the IEEE/CVF International conference on computer vision*, 15625–15636.
- Wu, J.; Zhang, Z.; Xia, Y.; Li, X.; Xia, Z.; Chang, A.; Yu, T.; Kim, S.; Rossi, R. A.; Zhang, R.; et al. 2024. Visual prompting in multimodal large language models: A survey. *arXiv preprint arXiv:2409.15310*.
- Xia, F.; Zamir, A. R.; He, Z.; Sax, A.; Malik, J.; and Savarese, S. 2018. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 9068–9079.
- Yang, Z.; Li, L.; Lin, K.; Wang, J.; Lin, C.-C.; Liu, Z.; and Wang, L. 2023. The dawn of Imms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9(1): 1.
- 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.; Wang, Z.; Wang, X. E.; and Wu, Q. 2024. Navgpt-2: Unleashing navigational reasoning capability for large vision-language models. In *European Conference on Computer Vision*, 260–278. Springer.
- Zhu, Y.; Mottaghi, R.; Kolve, E.; Lim, J. J.; Gupta, A.; Fei-Fei, L.; and Farhadi, A. 2017. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *2017 IEEE international conference on robotics and automation (ICRA)*, 3357–3364. IEEE.