

Why Only Text: Empowering Vision-and-Language Navigation with Multi-modal Prompts

Haodong Hong^{1,2}, Sen Wang^{1*}, Zi Huang¹, Qi Wu³ and Jiajun Liu^{2,1*}

¹The University of Queensland

²CSIRO Data61

³The University of Adelaide

{haodong.hong, sen.wang}@uq.edu.au, huang@itee.uq.edu.au, qi.wu01@adelaide.edu.au, jiajun.liu@csiro.au

Abstract

Current Vision-and-Language Navigation (VLN) tasks mainly employ textual instructions to guide agents. However, being inherently abstract, the same textual instruction can be associated with different visual signals, causing severe ambiguity and limiting the transfer of prior knowledge in the vision domain from the user to the agent. To fill this gap, we propose Vision-and-Language Navigation with Multi-modal Prompts (VLN-MP), a novel task augmenting traditional VLN by integrating both natural language and images in instructions. VLN-MP not only maintains backward compatibility by effectively handling text-only prompts but also consistently shows advantages with different quantities and relevance of visual prompts. Possible forms of visual prompts include both exact and similar object images, providing adaptability and versatility in diverse navigation scenarios. To evaluate VLN-MP under a unified framework, we implement a new benchmark that offers: (1) a training-free pipeline to transform textual instructions into multi-modal forms with landmark images; (2) diverse datasets with multi-modal instructions for different downstream tasks; (3) a novel module designed to process various image prompts for seamless integration with state-of-the-art VLN models. Extensive experiments on four VLN benchmarks (R2R, RxR, REVERIE, CVDN) show that incorporating visual prompts significantly boosts navigation performance. While maintaining efficiency with text-only prompts, VLN-MP enables agents to navigate in the pre-explore setting and outperform text-based models, showing its broader applicability. Code is available at <https://github.com/honghd16/VLN-MP>.

1 Introduction

With the booming of large language models (LLMs) [Brown *et al.*, 2020; Touvron *et al.*, 2023], there is a growing interest in employing natural language to guide embodied agents

*Corresponding author



Figure 1: An example of the effectiveness of multi-modal prompts for VLN. Initially, the agent struggles to identify the “wicker chair” due to unfamiliarity. While detailed textual explanations can increase comprehension challenges, introducing a visual prompt from the web is convenient and simplifies the task significantly. In this case, a photo of a wicker chair provides a clear, direct reference to help the agent accurately identify the target object.

in executing various tasks [Huang *et al.*, 2022]. Vision-and-Language Navigation (VLN) [Anderson *et al.*, 2018b] is a representative example, demanding not only individual language and vision understanding but also the co-grounding between these two modalities. Numerous methods [Zhu *et al.*, 2020; Zhu *et al.*, 2023; Wang *et al.*, 2023a; Majumdar *et al.*, 2020; Wang *et al.*, 2023b; Zhou *et al.*, 2024] and datasets [Jain *et al.*, 2019; Krantz *et al.*, 2020; Qi *et al.*, 2020; Ku *et al.*, 2020] have been proposed to address different challenges in VLN.

Current VLN tasks mainly rely on natural language instructions. However, real-world navigation often involves scenarios where images are crucial as guidance, providing essential and convenient supplementary information to instructions, as shown in Fig. 1. These image-enhanced instructions reduce the reliance on perfect language instructions and provide more precise path or destination depiction, thus helping agent navigation. For example, a photo of the target object can effectively supplement the language descriptions that

might be ambiguous. Moreover, while detailed language instructions are challenging to formulate, combining simple environmental photos with brief descriptions can create efficient multi-modal guidance. Such multi-modal prompts are standard practices in human navigation, whether when using apps like Google Maps or capturing images for complex descriptions. The ubiquity of cameras and smartphones has made acquiring these images practical and accessible, yet this setting remains under-explored in existing VLN research.

To address this, we introduce a novel task, Vision-and-Language Navigation with Multi-modal Prompts (VLN-MP), in which agents navigate with instructions that combine visual signals and natural language. Considering the varying quantity and quality of visual prompts in applications, we propose three distinct settings—Aligned, Related, and Terminal—to accommodate different scenarios. The Aligned setting needs multiple actual images, while the other two settings only demand similar views or a single view of the destination. Agents in VLN-MP are expected to be versatile across diverse image prompt modes, covering these three settings and handling varying numbers of images. VLN-MP works as an advancement over traditional VLN, enabling models to still operate effectively with textual instructions alone, yet achieving enhanced performance when supplemented with various image prompts, even a single image.

To facilitate research in VLN-MP, we propose a new benchmark with several essential components. Recognizing the importance and prevalence of landmarks in navigation [Yesiltepe *et al.*, 2021; He *et al.*, 2021], we focus on landmarks within instructions as the primary source of visual prompts. To transform textual instructions into multi-modal forms, we introduce a universal data generation pipeline leveraging large pre-training models [OpenAI, 2023; Li *et al.*, 2022c] to realize landmark phrase extraction, landmark detection, image selection, and data augmentation. Applying this pipeline to existing VLN datasets, we generate their corresponding multi-modal instruction datasets with various settings to assess agent adaptability to different image prompt modes. Furthermore, we develop a novel module, Multi-modal Prompts Fusion (MPF), which integrates features from visual prompts and text tokens in a unified embedding space. MPF is compatible with state-of-the-art VLN models like HAMT [Chen *et al.*, 2021] and DUET [Chen *et al.*, 2022] for enhanced navigation performance.

We conduct extensive experiments on four datasets—R2R, RxR, REVERIE, and CVDN—to demonstrate that agents trained in VLN-MP improve navigation performance across different image prompt settings while maintaining robustness in traditional VLN tasks. Through these experiments, we also provide valuable insights on leveraging multi-modal instructions to boost agent generalizability. Additionally, training agents under the VLN-MP paradigm prepares them for application in the pre-explore setting of VLN, where they outperform agents trained under the original setup.

2 Related Work

Vision-and-Language Navigation. In VLN, agents navigate to a target destination following natural language in-

structions within simulated environments such as Matterport3D [Chang *et al.*, 2017]. Since the introduction of the Room-to-Room (R2R) dataset [Anderson *et al.*, 2018b], numerous tasks and datasets have been developed to address various challenges and scenarios for instruction following agents. Ku *et al.* [Ku *et al.*, 2020] proposed Room-across-Room (RxR), a multilingual dataset featuring fine-grained instructions and pose traces of annotators. REVERIE [Qi *et al.*, 2020] was presented in which instructions only describe the target location and object rather than providing step-by-step guidance. Thomason *et al.* [Thomason *et al.*, 2020] developed the first vision-and-dialogue navigation dataset, CVDN, which requires agents to learn from the dialogue history and ask for assistance during navigation. Liu *et al.* [Liu *et al.*, 2023b] extends the VLN setting from ground scenarios to the sky and proposes the AerialVLN task with UAV-based environments. However, these works are confined to text-only instructions, while our work serves as the first paradigm to facilitate the use of multi-modal instructions in VLN.

Landmarks in VLN. Previous research has highlighted the significance of landmarks and extensively employed them in navigation agents [Berg *et al.*, 2020; Yesiltepe *et al.*, 2021]. In VLN, landmarks have been used for navigation graph construction [Hong *et al.*, 2020], instruction and trajectory decomposition [He *et al.*, 2021], and instruction generation [Wang *et al.*, 2022]. GELA [Cui *et al.*, 2023] introduces the human-annotated entity-landmark dataset GEL-R2R to enhance cross-modal alignment of agents. Wang *et al.* [Wang *et al.*, 2022] also incorporated landmark images to generate multi-modal instructions. However, our work differs in two key ways. Firstly, while they aim at producing novel textual instructions, we emphasize multi-modal instructions as the direct prompt for navigation. Secondly, unlike their strict requirements, our work allows for a broader range of landmark image quality and quantity.

Multi-modal Prompts. Multi-modal prompts have gained increasing popularity for their superior expressive power over single-modality and uniform input interfaces for training Visual Language Models (VLM). Flamingo [Alayrac *et al.*, 2022] leverages large-scale multi-modal web corpora containing interleaved text and images to complete various textual tasks with few examples. VIMA [Jiang *et al.*, 2022] establishes the first robot learning benchmark for multi-modal-prompted tasks, providing corresponding models, datasets, and evaluation protocols. MaPLe [Khattak *et al.*, 2023] incorporates multi-modal prompt learning into CLIP [Radford *et al.*, 2021] to enhance model generalization, though its focus remains on classification tasks and can not manage complex instructions or detailed panoramas. Our work contributes to the VLN task, where multi-modal instructions are rarely studied but hold significant importance.

3 VLN-MP Task

3.1 Task Definition

In VLN, an agent needs to navigate towards a target destination following a natural language instruction $\hat{x} = (x_1, x_2, \dots, x_L)$ with L words. At each timestep t , the agent



Figure 2: An example of the R2R instruction and the landmark images in three settings of R2R-MP. The aligned setting (green) provides sufficient highly grounded landmark images. The related setting (yellow) presents similar images associated with the phrase but not necessarily the instruction, like the same image for two carpets. The Terminal setting (blue) includes only an image of the last landmark.

observes its surroundings through a panoramic representation comprising $N = 36$ views $O_t = \{o_t^i\}_{i=1}^N$. Each view o_t^i includes an RGB image I_t^i and directional details of heading θ_t^i and elevation ϕ_t^i angles of its current node V_t . Based on its policy π , the agent decides on an action a_t to transition to one of the neighboring nodes $\mathcal{N}(V_t)$ by selecting the view that aligns best with the target node. In this process, visual and textual signals are separate: the visual ones constantly change during navigation, while the textual input remains static, serving as the human-agent interface. However, in real-world navigation, images are often employed as part of the instructions, providing essential supplementary information. This problem is ignored by current VLN works, complicating human-agent communication and preventing agents from using image prompts to improve performance.

Therefore, we propose the Vision-and-Language Navigation with Multi-modal Prompts (VLN-MP) task first to consider multi-modal instructions in navigation. Assume we provide n images (I_1, I_2, \dots, I_n) corresponding to n phrases (z_1, z_2, \dots, z_n) within instruction \hat{x} . Each phrase z_i is represented as $z_i = x_{start_i:end_i}$, where $x_{i:j} = (x_i, x_{i+1}, \dots, x_j)$ denotes words from position i to j . By inserting images next to corresponding phrases, the textual instruction \hat{x} is transformed into the multi-modal form:

$$\bar{x} = (x_{1:end_1}, I_1, x_{end_1+1:end_2}, I_2, \dots, I_n, x_{end_n+1:L}) \quad (1)$$

The agent receives \bar{x} instead of \hat{x} as navigation guidance, with others the same as in VLN. Here we focus on scenarios where each phrase is paired with a single image.

3.2 Image Prompt Settings

Due to varying difficulties in acquiring images, agents may be provided with visual prompts of different numbers and qualities, ranging from one similar image from the web to sufficient actual images from cameras. To enable agents to adapt to such variations, we propose three different settings for VLN-MP based on the number and relevance of the prompt images in the instructions: **Aligned**, **Related**, and **Terminal**. The Aligned setting aims to provide precise and abundant prompts with two criteria. Firstly, each image I_i should be aligned with the actual view that the phrase z_i describes. Secondly, the number of images n should be large

enough to cover the entire instruction. The Related setting relaxes the first requirement and only demands I_i related to the phrase, aiming to facilitate instruction understanding without requiring perfect alignment. The Terminal setting removes the second requirement, providing only one image depicting the view around the destination. This setting helps agents target the stop position and accomplish object-centric tasks such as object navigation [Zhu *et al.*, 2021]. Besides these settings, we also randomly vary the number of visual prompts to further evaluate agent adaptability.

4 Benchmark Implementation

In this section, we introduce the proposed benchmark of VLN-MP in detail, including the data generation pipeline, four novel datasets, and the MPF module.

4.1 VLN-MP Data Generation Pipeline

To utilize textual instructions in VLN, we propose a novel data generation pipeline to transform them into multi-modal forms, as shown in Fig. 3. Using large pre-training models, our pipeline can be applied to VLN datasets both generally and efficiently without extra training.

Extraction. Extracting landmark phrases from instructions is a crucial first step, as its results will impact all subsequent processes. Fortunately, large language models excel in this task. Previous work [Wang *et al.*, 2022] employed a BERT-based parser [Devlin *et al.*, 2019] but still misidentified irrelevant words as landmarks. Moreover, some landmark phrases are not suitable for detection within environments. To improve accuracy and accommodate subsequent stages, we utilize the advanced GPT-4 model [OpenAI, 2023] from OpenAI with specially crafted prompts to perform the extraction.

Detection. With the landmark phrases, we detect potential entities in the environment, a task requiring phrase grounding or open-vocabulary object detection. Our strategy involves selecting different models based on the instruction type. For object-centered instructions or short instructions, we use GroundingDINO [Liu *et al.*, 2023a], known for its high accuracy in detecting specific object names. For detailed and lengthy instructions, we choose the GLIP [Li *et al.*, 2022c] model, which excels in phrase grounding and is more

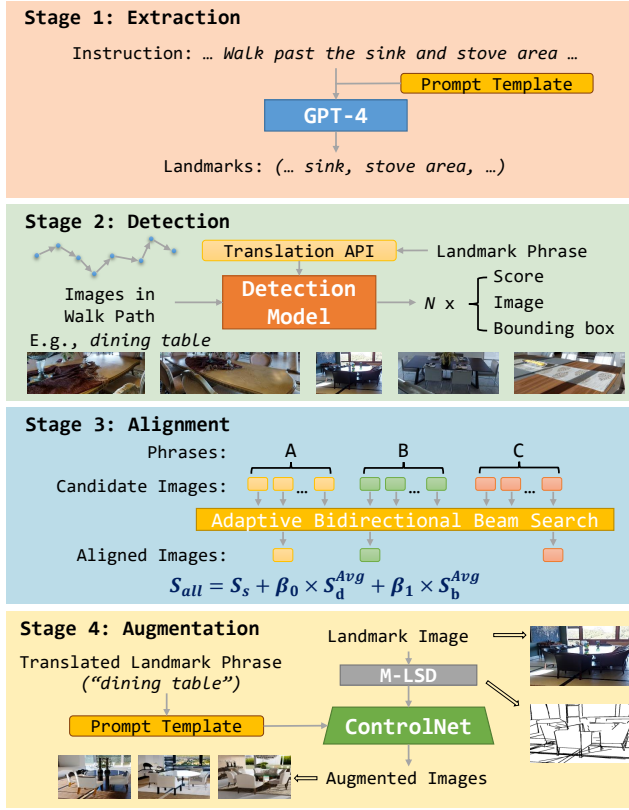


Figure 3: Four stages of our data generation pipeline for transforming textual instructions into multi-modal forms.

suitable for objects contextualized within phrases. Specifically, we use the chosen model in a zero-shot manner to scan all images along the path (N_1, N_2, \dots, N_m) consisting of m nodes, and generate multiple candidates for each landmark phrase z_i . The resulting candidates C are represented in triads $C = (S, I, B)$, where S is the detection score, I is the detected image, and B indicates the bounding box.

Alignment. In this stage, we select the most suitable image from numerous candidates for each phrase based on different settings. In the Related setting, we choose the image with the highest detection score by $i^* = \arg \max_{j \in [1, k_i]} S_j^i$ for z_i , corresponding to the scenario of choosing a similar picture from the surroundings or the web. For the Aligned setting, we use the order of landmarks in the instruction to guide our selection, prioritizing images that appear in the same sequence. We apply the Kendall rank correlation coefficient to measure this alignment as the sequence score $S_s = \frac{4 \times c}{n \times (n-1)} - 1$, where c is the number of concordant pairs ($I_i \leq I_j$ for any $1 \leq i < j \leq n$). Additionally, we introduce the bounding box score S_b to prioritize larger images for more details. The final score for selecting images is a combination of detection score S_d , sequence score S_s , and bounding box score S_b :

$$S_{all} = S_s + \beta_0 \times S_d^{Avg} + \beta_1 \times S_b^{Avg} \quad (2)$$

For lengthy instructions, we propose an adaptive bidirectional beam search to find the best image combinations from vast

search space efficiently. For the Terminal setting, we directly use the image of the last landmark from the Aligned setting.

Augmentation. A notable advantage of visual prompts is the ease and diversity of data augmentation compared to natural language. Unlike previous studies [Liu *et al.*, 2021; Li *et al.*, 2022b], we employ ControlNet [Zhang *et al.*, 2023] to generate N_{aug} novel images based on the line segments of the landmark images from M-LSD [Gu *et al.*, 2022]. This operation aims to maintain the coarse-grained features like shape and layout consistent with the visual observations while altering fine-grained details like texture and color to enrich visual prompts. These images serve as augmented multi-modal instructions during training to enhance the dataset diversity and improve the model’s generalizability.

4.2 Datasets for VLN-MP

We apply our pipeline to four popular VLN datasets—R2R [Anderson *et al.*, 2018b], RxR [Ku *et al.*, 2020], REVERIE [Qi *et al.*, 2020], and CVDN [Thomason *et al.*, 2020]—generating their corresponding multi-modal versions, denoted with the suffix “-MP”.

R2R-MP. R2R includes 17,409 instructions¹ for short paths in the Matterport3D dataset [Chang *et al.*, 2017]. Each instruction details the route from start to finish and contains several landmark phrases. Using GroundingDINO, our pipeline produces 17,328 multi-modal instructions, with only 0.4% of instructions failing to detect any landmark images. The average number of landmarks for the Aligned setting is 4.15, reaching up to 18 in some instructions. An example of the three settings in R2R-MP is shown in Fig. 2.

RxR-MP. RxR provides 101,932 instructions with three languages, longer paths, more visible entities, and time alignment between words and virtual poses. We use our pipeline with the GLIP model to generate 100,923 multi-modal instructions per setting. The Aligned setting has an average of 7.17 landmarks, with a maximum of 63. Additionally, we integrated the Marky-mT5 dataset [Wang *et al.*, 2022] as an alternative Aligned setting for comparison, which has an average of 9.57 landmarks with a maximum of 83 due to the forceful allocation of a frame for each phrase.

REVERIE-MP. Unlike R2R and RxR, REVERIE focuses on object-oriented instructions that describe the target object, typically involving a single landmark—the target object. We directly use its bounding box annotations to generate 15,410 multi-modal instructions, the same number as REVERIE, with each instruction including one visual prompt.

CVDN-MP. CVDN also provides object-oriented instructions but differs from REVERIE by employing ambiguous objects that can refer to multiple instances to make agents use the dialogue history. In CVDN-MP, we discard the dialogue to emphasize the necessity of visual prompts in ambiguous instructions like “Find the ⟨ object ⟩”, which is more applicable to real-world situations. We also employ BLIP2 [Li *et al.*, 2023] to generate captions for the provided visual prompts to offer an alternative, detailed instruction for comparison.

¹We did not include the test split with paths not open source.

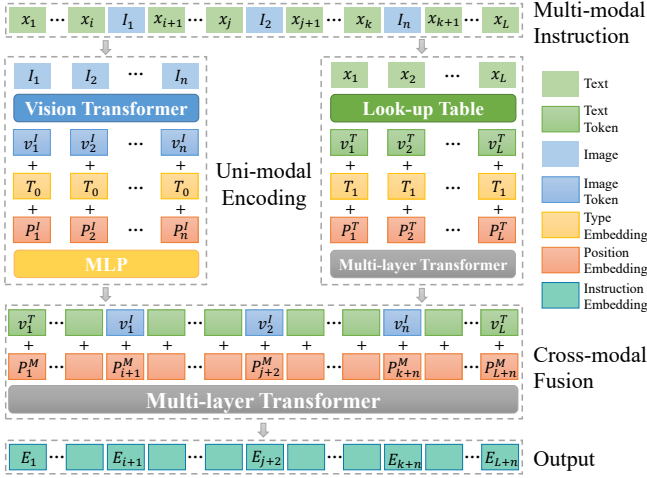


Figure 4: The architecture of the Multi-modal Prompts Fusion (MPF) module. Texts and images are processed separately and then combined to obtain cross-modal representations. Some symbols are omitted for simplicity.

CVDN-MP includes 6,031 instructions, each with a visual prompt and an alternate extended instruction.

4.3 Multi-modal Prompts Fusion Module

To effectively utilize multi-modal instructions, we propose a novel Multi-modal Prompts Fusion (MPF) module, as illustrated in Fig. 4. The MPF module includes a visual branch to process visual prompts individually, parallel to the text branch that handles language prompts, and a fusion layer to combine these two types of tokens. The visual prompts are initially processed by an image encoder to extract features and then enriched with type embedding and position embedding to yield image tokens. Concurrently, text tokens are processed and merged with visual tokens, further integrating another layer of position embedding to link them effectively. This combined multi-modal token sequence is input into a multi-layer Transformer [Vaswani *et al.*, 2017] to synthesize the final tokens for the multi-modal instruction. These tokens are then further processed to determine the next action according to the design of different VLN models. Crucially, position encoding is applied twice for each token: the first to mark the sequence of the images and texts separately and the second to match these two types of tokens. This dual application ensures the model recognizes the relationship between phrases and their corresponding images.

5 Experiments

5.1 Experimental Setup

Datasets. We conduct our experiments on the proposed four VLN-MP datasets. The original datasets have four splits: train, validation seen (val seen), validation unseen (val unseen), and test unseen. As the ground truth paths in the test split are not released, their multi-modal versions only contain the first three splits for training and evaluation. The distinction between val seen and val unseen is that houses in the former are also in the train split [Zhang *et al.*, 2021].

Score	Precision	Recall	F1
Fuzzy Matching	0.95	0.89	0.91
ROGUE-L	0.87	0.80	0.81

Table 1: Phrase similarity between R2R-MP and GELR2R.

Method	Matching	Neighboring
w/o Alignment	0.40	0.62
w/ Alignment	0.69	0.91

Table 2: Accuracy of viewpoint matching and neighboring for R2R-MP and GELR2R images, with and without Alignment stage.

Baseline VLN Models. We adopt HAMT [Chen *et al.*, 2021] and DUET [Chen *et al.*, 2022] as baseline models, which are the mainstream architectures in VLN. HAMT employs a transformer-based network to encode instructions, visual observations, and history together for action prediction. DUET expands on HAMT by building a real-time topological map to enable global action decisions over local movement.

Evaluation Metrics. We evaluate the agent performance on the following metrics: (1) Success Rate (**SR**): the ratio of agents stopping within 3 meters of the target; (2) Success rate weighted by Path Length (**SPL**) [Anderson *et al.*, 2018a]: SR normalized by the ratio between the length of the shortest path and the predicted path; (3) the normalized Dynamic Time Warping (**nDTW**) [Ilharco *et al.*, 2019]: a measure of instruction fidelity by computing the similarity between the reference path and the predicted path; (4) Goal Progress (**GP**) [Thomason *et al.*, 2020]: metric for CVDN which measures the average difference between the length of the completed trajectory and the remaining distance to the goal.

Implementation Details. In the pipeline, we utilize GPT-4 from OpenAI’s official API, and the GLIP-L and GroundingDINO-T models for landmark detection. The weights for score balancing in Eq. (2) are $\beta_0 = 0.5$ and $\beta_1 = 0.1$ to prioritize the sequence score (S_s) over the others, reflecting their relative importance in our method. The small β_1 is used to ensure that the bounding box scores become significant only when multiple bounding boxes are generated for the same object, in which case S_d and S_s remain consistent. For non-English languages, we use the Google translate service² to translate them into English. We generate five novel images per visual prompt using the control_sd15_mlsd model for data augmentation. For baseline models, we follow the implementation details in their official repositories. Following their setting for processing visual observations, we employ the ViT-B/32 model as the feature extractor for landmark images in R2R, CVDN, and REVERIE, while CLIP-ViT-B/16 is used for RxR. During training, we select the augmented data with a probability of $\gamma = 0.2$ to replace original landmark images. All models are fine-tuned for 200K iterations with a learning rate of $1e-5$ and a batch size of 8 on a single NVIDIA A6000 GPU. The best model is selected based on performance in the val unseen split.

²<https://cloud.google.com/translate>

#	Model	Image Prompt		Setting	Validation Seen			Validation Unseen		
		Train	Eval		SR \uparrow	SPL \uparrow	nDTW \uparrow	SR \uparrow	SPL \uparrow	nDTW \uparrow
1	Multi Baseline [Ku <i>et al.</i> , 2020]	×	×	-	25.2	-	42.2	22.8	-	38.9
2	Mono Baseline [Ku <i>et al.</i> , 2020]	×	×	-	28.8	-	46.8	28.5	-	44.5
3	EnvDrop [Shen <i>et al.</i> , 2022]	×	×	-	-	-	42.6	-	55.7	-
4	CLEAR [Li <i>et al.</i> , 2022a]	×	×	-	-	-	44.4	39.3	57.0	-
5	HOP+ [Qiao <i>et al.</i> , 2023a]	×	×	-	53.6	47.9	59.0	45.7	38.4	52.0
6	VLN-PETL [Qiao <i>et al.</i> , 2023b]	×	×	-	60.5	56.8	65.7	57.8	54.2	64.9
7	HAMT [Chen <i>et al.</i> , 2021]	×	×	-	59.4	58.9	65.3	56.5	56.0	63.1
8	HAMT+MPF	✓	×	-	64.6	60.8	68.6	57.6	53.5	64.0
9	HAMT+MPF	✓	✓	T	65.8	62.0	69.2	58.1	54.0	64.6
10	HAMT+MPF	✓	✓	R	67.2	63.0	70.1	58.7	54.1	63.6
11	HAMT+MPF	✓	✓	A-P	68.8	63.7	71.1	59.3	55.1	65.0
12	HAMT+MPF	✓	✓	A-M	69.2	65.5	71.7	60.0	56.3	66.4

Table 3: Navigation performance on the RxR-MP dataset. “A” means Aligned, “R” is Related, and “T” refers to Terminal; “P” and “M” indicate images from our pipeline and Marky-mT5, respectively.

5.2 Dataset Evaluation

We compare the extracted phrases and landmark images with Marky-mT5 and GELR2R datasets for evaluation.

Phrases. Given the multiple errors produced by the BERT-based parser in Marky-mT5, we mainly compare our phrases with the human-annotated data from GELR2R. To thoroughly assess the similarity of phrases between our R2R-MP and GELR2R, we adopt Fuzzy Matching and ROGUE-L [Lin, 2004] scoring and calculate the Precision, Recall, and F1 scores for each method. As shown in Tab. 1, our results achieve high similarity with the gold data in GELR2R, especially with an impressive 95% precision in fuzzy matching. This result demonstrates the effectiveness of our method.

Images. For landmark images, we first prove that our object-centric approach offers enhanced accuracy over Marky-mT5 in the RxR dataset. We employ the CLIP model [Radford *et al.*, 2021] to measure the alignment of images from both datasets with the landmark phrases in a binary classification manner. The results, as depicted in Fig. 5, reveal that RxR-MP outperforms Marky-mT5 with an average score of 0.649 compared to 0.351. Furthermore, we also quantify the number of instances where the phrase demonstrates a preference for either dataset. Additionally, 66.9% of phrases favored images from RxR-MP over 33.1% for Marky-mT5, indicating a stronger alignment of our images with the landmark phrases.

To demonstrate the effectiveness of our alignment phase, we calculate the matching accuracy of the viewpoints of landmark images in R2R-MP with the ground truth in GELR2R, including the exact matching and neighboring viewpoints. The comparison also included images with the highest detection scores without alignment in Tab. 2. Nearly 70% of landmark phrases in R2R-MP correctly identify the corresponding viewpoints, and 91% are adjacent to the actual viewpoint, which may also provide a different view for the same landmark, far surpassing those without alignment, underscoring the significance of our method.

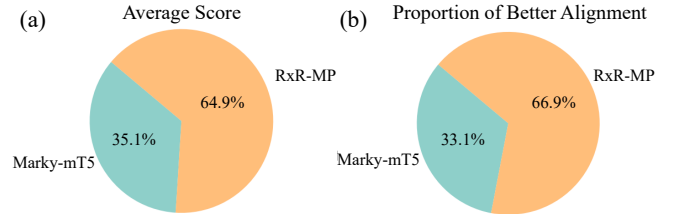


Figure 5: The average score and proportion of a better alignment of Marky-mT5 and our RxR-MP towards the landmark phrases.

5.3 Navigation Performance

We provide the navigation performance of RxR-MP and CVDN-MP here to represent the step-wise instructions and goal-oriented instructions.

RxR-MP. Tab. 3 presents the performance of different models on the val seen and unseen splits of the RxR-MP dataset. Agents in experiments #1-7 are trained on text-only instructions, and our HAMT+MPF agent in experiments #8-12 is trained with multi-modal instructions and evaluated under different settings. Experiments #7 and #8 demonstrate that VLN-MP-trained models maintain backward compatibility and achieve higher performance with textual instructions, especially in seen scenarios. Comparing experiments #9-12, we can see that our agent adapts well to different visual prompts and consistently outperforms text-only scenarios, even with one single image. It also reveals that the agent performance improves with the increase in the number and relevance of visual prompts, which is intuitive and outlines our alignment stage. The Marky-mT5 version achieves the best results, as it provides larger and more images.

CVDN-MP. In CVDN-MP, we investigate the situation when the visual prompts provide necessary information for navigation. Tab. 4 shows the Goal Progress (GP) of agents when presented with various types of prompts. Compared with single-modality prompts, utilizing multi-modal prompts maintains backward compatibility and significantly performs better with a target image. This phenomenon highlights the

Model	Image Prompt		Text Prompt	GP \uparrow	
	Train	Eval		Val Seen	Val Unseen
HAMT	×	×	✓	6.30	2.43
	×	×	✓*	9.80	2.52
	✓	×	✓	7.03	2.49
	✓	✓	×	8.83	2.29
	✓	✓	✓	10.92	3.11
DUET	×	×	✓	6.19	3.00
	×	×	✓*	9.61	3.12
	✓	×	✓	6.27	3.57
	✓	✓	×	7.89	3.12
	✓	✓	✓	10.09	4.13

Table 4: Navigation performance on the CVDN-MP. When visual prompts are used, MPF is integrated. * means with captions.

#	γ	Val Seen		Val Unseen	
		SR \uparrow	SPL \uparrow	SR \uparrow	SPL \uparrow
1	0.0	69.2	65.5	60.0	56.3
2	0.1	68.4	63.9	60.2	55.2
3	0.2	67.6	63.5	61.8	57.2
4	0.5	66.8	62.6	59.8	55.5
5	0.8	66.4	62.2	59.6	55.1
6	1.0	64.0	59.9	58.8	54.6

Table 5: Results with different proportions of augmented data.

significance of visual prompts in goal-oriented tasks and the effectiveness of our MPF module. The captions of the visual prompts enhance GP in seen scenarios but not in unseen ones, suggesting that visual prompts are more intuitive and effective than language.

5.4 Ablation Study

Proportion of Augmented Data. In Tab. 5, we evaluate the impact of the ratio γ of ControlNet-generated augmented data to original data during training. As γ increases, performance in the val seen split gradually declines. While in the val unseen split, it initially improves, peaking at $\gamma = 0.2$, before decreasing. This trend is intuitive, as augmented data mitigates over-fitting to familiar landmarks and improves adaptability to new environments. However, excessive augmented data can create a disparity between prompts and observations and confuse agents.

Position Encoding. Tab. 6 presents the navigation performance of the MPF module with different position encodings in the Aligned setting of RxR-MP. To distinguish them, we refer to the first as visual position encoding (VPE), as it is only applied to the visual prompts, and the second as multi-modal position encoding (MPE), as it is applied to the multi-modal tokens. Both types of position encoding are beneficial to agents when dealing with multi-modal instructions. Since MPE implies VPE, experiment #3 in Tab. 6 outperforms #2. Combining these two position encodings in experiment #4 results in the best performance, supporting our assertion of their distinct functions.

#	VPE	MPE	Val Seen		Val Unseen	
			SR \uparrow	SPL \uparrow	SR \uparrow	SPL \uparrow
1	×	×	67.6	63.3	58.7	53.8
2	✓	×	68.0	63.4	59.0	54.4
3	×	✓	67.4	62.9	59.3	55.2
4	✓	✓	69.2	65.5	60.0	56.3

Table 6: Navigation performance of MPF with different position encodings on the Aligned setting of RxR-MP.

Model	Pre-explore	Val Seen		Val Unseen	
		SR \uparrow	SPL \uparrow	SR \uparrow	SPL \uparrow
HAMT	×	59.4	58.9	56.5	56.0
Ours	×	64.6	60.8	57.6	53.5
Ours	✓	66.3	62.1	59.0	55.4

Table 7: Comparison of HAMT and our HAMT+MPF agents on RxR dataset under different settings.

5.5 Pre-explore Setting of VLN

Most VLN works focus on the single-run setting, where agents navigate in an environment once without prior knowledge or multiple trials. However, the pre-explore setting, where agents familiarize the environment before navigation, is also crucial, especially for household robots that operate persistently in an environment [Krantz *et al.*, 2023]. VLN-MP allows existing VLN models to be applied to this pre-explore setting with only one more trial and enhances performance without additional training. Specifically, the agent first follows the textual instruction to generate a pseudo path and then employs our pipeline to extract landmark images from the visited nodes and their neighbors to form a multi-modal instruction for navigation. Tab. 7 show that this method can significantly improve navigation performance in RxR, with a notable 6.9% and 2.5% SR increase in seen and unseen scenarios compared to the original HAMT, emphasizing the applicability of VLN-MP in traditional VLN tasks.

6 Conclusion

In this paper, we propose the novel VLN-MP task, which enhances agent navigation by integrating visual prompts into textual instructions. VLN-MP extends traditional VLN by ensuring full backward compatibility and demonstrating adaptability to various visual prompts. We establish the first benchmark for VLN-MP, including a pipeline for converting textual instructions into multi-modal forms, four datasets for different downstream tasks, and a novel MPF module to process multi-modal instructions efficiently. We conduct extensive experiments to evaluate our proposed components and provide empirical insights regarding the usage of visual prompts. Notably, VLN-MP allows agents to operate in a pre-explore setting, achieving enhanced results compared to traditional VLN. We believe that VLN-MP and our benchmark will broaden the scope of VLN in real-world applications and open up new research opportunities in the field.

Acknowledgements

This work was partially supported by ARC DECRA grant (DE200101610) and CSIRO’s Science Leader project R-91559.

References

- [Alayrac *et al.*, 2022] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 35:23716–23736, 2022.
- [Anderson *et al.*, 2018a] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018.
- [Anderson *et al.*, 2018b] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, pages 3674–3683, 2018.
- [Berg *et al.*, 2020] Matthew Berg, Deniz Bayazit, Rebecca Mathew, Ariel Rotter-Aboyoun, Ellie Pavlick, and Stefanie Tellex. Grounding language to landmarks in arbitrary outdoor environments. In *ICRA*, pages 208–215, 2020.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33:1877–1901, 2020.
- [Chang *et al.*, 2017] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niebner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *3DV*, pages 667–676, 2017.
- [Chen *et al.*, 2021] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. *NeurIPS*, 34:5834–5847, 2021.
- [Chen *et al.*, 2022] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Think global, act local: Dual-scale graph transformer for vision-and-language navigation. In *CVPR*, pages 16537–16547, 2022.
- [Cui *et al.*, 2023] Yibo Cui, Liang Xie, Yakun Zhang, Meishan Zhang, Ye Yan, and Erwei Yin. Grounded entity-landmark adaptive pre-training for vision-and-language navigation. In *ICCV*, pages 12043–12053, 2023.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [Gu *et al.*, 2022] Geonmo Gu, Byungsoo Ko, SeoungHyun Go, Sung-Hyun Lee, Jingeun Lee, and Minchul Shin. Towards light-weight and real-time line segment detection. In *AAAI*, volume 36, pages 726–734, 2022.
- [He *et al.*, 2021] Keji He, Yan Huang, Qi Wu, Jianhua Yang, Dong An, Shuanglin Sima, and Liang Wang. Landmark-rxr: Solving vision-and-language navigation with fine-grained alignment supervision. *NeurIPS*, 34:652–663, 2021.
- [Hong *et al.*, 2020] Yicong Hong, Cristian Rodriguez, Yuankai Qi, Qi Wu, and Stephen Gould. Language and visual entity relationship graph for agent navigation. *NeurIPS*, 33:7685–7696, 2020.
- [Huang *et al.*, 2022] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *ICML*, pages 9118–9147, 2022.
- [Ilharco *et al.*, 2019] Gabriel Ilharco, Vihan Jain, Alexander Ku, Eugene Ie, and Jason Baldridge. General evaluation for instruction conditioned navigation using dynamic time warping. *arXiv preprint arXiv:1907.05446*, 2019.
- [Jain *et al.*, 2019] Vihan Jain, Gabriel Magalhaes, Alexander Ku, Ashish Vaswani, Eugene Ie, and Jason Baldridge. Stay on the path: Instruction fidelity in vision-and-language navigation. In *ACL*, pages 1862–1872, 2019.
- [Jiang *et al.*, 2022] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. VIMA: General robot manipulation with multimodal prompts. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.
- [Khattak *et al.*, 2023] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *CVPR*, pages 19113–19122, 2023.
- [Krantz *et al.*, 2020] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the navigraph: Vision-and-language navigation in continuous environments. In *ECCV*, pages 104–120, 2020.
- [Krantz *et al.*, 2023] Jacob Krantz, Shurjo Banerjee, Wang Zhu, Jason Corso, Peter Anderson, Stefan Lee, and Jesse Thomason. Iterative vision-and-language navigation. In *CVPR*, pages 14921–14930, 2023.
- [Ku *et al.*, 2020] Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *EMNLP*, pages 4392–4412, 2020.
- [Li *et al.*, 2022a] Jialu Li, Hao Tan, and Mohit Bansal. Clear: Improving vision-language navigation with cross-lingual, environment-agnostic representations. In *NAACL*, pages 633–649, 2022.
- [Li *et al.*, 2022b] Jialu Li, Hao Tan, and Mohit Bansal. Envedit: Environment editing for vision-and-language navigation. In *CVPR*, pages 15407–15417, 2022.

- [Li *et al.*, 2022c] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *CVPR*, pages 10965–10975, 2022.
- [Li *et al.*, 2023] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, pages 19730–19742, 2023.
- [Lin, 2004] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- [Liu *et al.*, 2021] Chong Liu, Fengda Zhu, Xiaojun Chang, Xiaodan Liang, Zongyuan Ge, and Yi-Dong Shen. Vision-language navigation with random environmental mixup. In *ICCV*, pages 1644–1654, 2021.
- [Liu *et al.*, 2023a] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023.
- [Liu *et al.*, 2023b] Shubo Liu, Hongsheng Zhang, Yuankai Qi, Peng Wang, Yanning Zhang, and Qi Wu. Aerialvln: Vision-and-language navigation for uavs. In *ICCV*, pages 15384–15394, 2023.
- [Majumdar *et al.*, 2020] Arjun Majumdar, Ayush Srivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra. Improving vision-and-language navigation with image-text pairs from the web. In *ECCV*, pages 259–274, 2020.
- [OpenAI, 2023] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [Qi *et al.*, 2020] Yuankai Qi, Qi Wu, Peter Anderson, Xin Wang, William Yang Wang, Chunhua Shen, and Anton van den Hengel. Reverie: Remote embodied visual referring expression in real indoor environments. In *CVPR*, pages 9982–9991, 2020.
- [Qiao *et al.*, 2023a] Yanyuan Qiao, Yuankai Qi, Yicong Hong, Zheng Yu, Peng Wang, and Qi Wu. Hop+: History-enhanced and order-aware pre-training for vision-and-language navigation. *TPAMI*, 2023.
- [Qiao *et al.*, 2023b] Yanyuan Qiao, Zheng Yu, and Qi Wu. Vln-petl: Parameter-efficient transfer learning for vision-and-language navigation. In *ICCV*, pages 15443–15452, 2023.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021.
- [Shen *et al.*, 2022] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can clip benefit vision-and-language tasks? *ICLR*, 2022.
- [Thomason *et al.*, 2020] Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. Vision-and-dialog navigation. In *CoRL*, pages 394–406, 2020.
- [Touvron *et al.*, 2023] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017.
- [Wang *et al.*, 2022] Su Wang, Ceslee Montgomery, Jordi Orbay, Vighnesh Birodkar, Aleksandra Faust, Izzeddin Gur, Natasha Jaques, Austin Waters, Jason Baldridge, and Peter Anderson. Less is more: Generating grounded navigation instructions from landmarks. In *CVPR*, pages 15428–15438, 2022.
- [Wang *et al.*, 2023a] Liuyi Wang, Zongtao He, Jiagui Tang, Ronghao Dang, Naijia Wang, Chengju Liu, and Qijun Chen. A dual semantic-aware recurrent global-adaptive network for vision-and-language navigation. In *IJCAI*, pages 1479–1487, 2023.
- [Wang *et al.*, 2023b] Zun Wang, Jialu Li, Yicong Hong, Yi Wang, Qi Wu, Mohit Bansal, Stephen Gould, Hao Tan, and Yu Qiao. Scaling data generation in vision-and-language navigation. In *ICCV*, pages 12009–12020, 2023.
- [Yesiltepe *et al.*, 2021] Demet Yesiltepe, Ruth Conroy Dalton, and Ayse Ozbil Torun. Landmarks in wayfinding: a review of the existing literature. *Cognitive processing*, 22:369–410, 2021.
- [Zhang *et al.*, 2021] Yubo Zhang, Hao Tan, and Mohit Bansal. Diagnosing the environment bias in vision-and-language navigation. In *IJCAI*, pages 890–897, 2021.
- [Zhang *et al.*, 2023] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, pages 3836–3847, 2023.
- [Zhou *et al.*, 2024] Gengze Zhou, Yicong Hong, and Qi Wu. Navgpt: Explicit reasoning in vision-and-language navigation with large language models. In *AAAI*, pages 7641–7649, 2024.
- [Zhu *et al.*, 2020] Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. Vision-language navigation with self-supervised auxiliary reasoning tasks. In *CVPR*, pages 10012–10022, 2020.
- [Zhu *et al.*, 2021] Fengda Zhu, Xiwen Liang, Yi Zhu, Qizhi Yu, Xiaojun Chang, and Xiaodan Liang. Soon: Scenario oriented object navigation with graph-based exploration. In *CVPR*, pages 12689–12699, 2021.
- [Zhu *et al.*, 2023] Fengda Zhu, Vincent CS Lee, Xiaojun Chang, and Xiaodan Liang. Vision language navigation with knowledge-driven environmental dreamer. In *IJCAI*, pages 1840–1848, 2023.