

MonoDream: Monocular Vision-Language Navigation with Panoramic Dreaming

Shuo Wang^{1,3*}, Yongcai Wang^{1†}, Zhaoxin Fan^{2†}, Yucheng Wang^{3‡}, Maiyue Chen³, Kaihui Wang³
Zhizhong Su³, Wanting Li¹, Xudong Cai¹, Yeying Jin⁴, Deying Li¹

¹Renmin University of China

²Innovation Center for Future Blockchain and Privacy Computing

³Horizon Robotics

⁴National University of Singapore

Abstract

Vision-Language Navigation (VLN) tasks often leverage panoramic RGB and depth inputs to provide rich spatial cues for action planning, but these sensors can be costly or less accessible in real-world deployments. Recent approaches based on Vision-Language Action (VLA) models achieve strong results with monocular input, yet they still lag behind methods using panoramic RGB-D information. We present MonoDream, a lightweight VLA framework that enables monocular agents to learn a Unified Navigation Representation (UNR). This shared feature representation jointly aligns navigation-relevant visual semantics (e.g., global layout, depth, and future cues) and language-grounded action intent, enabling more reliable action prediction. MonoDream further introduces Latent Panoramic Dreaming (LPD) tasks to supervise the UNR, which train the model to predict latent features of panoramic RGB and depth observations at both current and future steps based on only monocular input. Experiments on multiple VLN benchmarks show that MonoDream consistently improves monocular navigation performance and significantly narrows the gap with panoramic-based agents.

Project page —

https://horizonrobotics.github.io/robot_lab/monodream

Introduction

Vision-Language Navigation tasks (Wu et al. 2024; Anderson et al. 2018; Gu et al. 2022) require embodied agents to follow language instructions and navigate to specified targets in 3D environments. Early successful approaches often rely on global perceptual inputs, such as panoramic RGB-D images (Hong et al. 2022; Wang et al. 2023; An et al. 2022), which provide a wide field of view and explicit visual and geometric information. These inputs allow agents to build a more complete understanding of the environment and achieve high navigation success rates.

However, panoramic cameras and depth sensors introduce higher cost, power consumption, added weight and hardware integration complexity, making them less practical in

many real-world deployments. Recent research has therefore focused on more lightweight settings where the agent is equipped with only a single forward-facing RGB camera (Zhang et al. 2024b,a; Cheng et al. 2024).

While monocular agents are easier to deploy, their navigation performance still lags significantly behind systems with panoramic RGB-D inputs. This gap arises from the narrow egocentric view of monocular agents, which limits the ability to infer latent global spatial and geometric cues that are beneficial for navigation. While such cues can be more directly captured from panoramic or depth observations, they are harder to extract from a monocular view. For example, when navigating a building, the agent may see a long hallway but be unaware of a side door or staircase just outside its view. Such blind spots make it difficult to reason about local cues into a coherent global picture and plan several steps ahead, which are all critical for reliable navigation.

To tackle this issue, we propose **MonoDream**, a lightweight VLA framework that equips monocular VLN agents with a latent imagination capability. Our approach is built upon a key insight: navigation-relevant information, including the predicted actions and the understanding of the global scene, can be encoded into a shared representation space and could be inferred by the agent. This insight is inspired by neuroscientific findings that the human brain reasons the current panoramic scene from partial views (Robertson et al. 2016) and internally simulates upcoming scenes based on intention (Seeber et al. 2025). MonoDream enables monocular agents to learn to infer and complete a holistic understanding of the current and future environments from limited egocentric observations.

To achieve this, MonoDream jointly aligns the navigation-relevant information, including implicit action intent, panoramic scene layout, depth perception, and future dynamics, into a unified latent space called the **Unified Navigation Representation (UNR)**. The UNR can be decoded into navigation actions or directly as features of global information. We further design **Latent Panoramic Dreaming (LPD)** tasks that supervise the UNR by aligning the latent features of panoramic RGB-D observations at both the current and future steps. These tasks encourage the agent to develop a coherent, geometry-aware and future-aware internal model of the environment, enabling more informed navigation deci-

*This work was done while Shuo Wang was a Research Intern with Horizon Robotics.

†Corresponding authors

‡Project leader

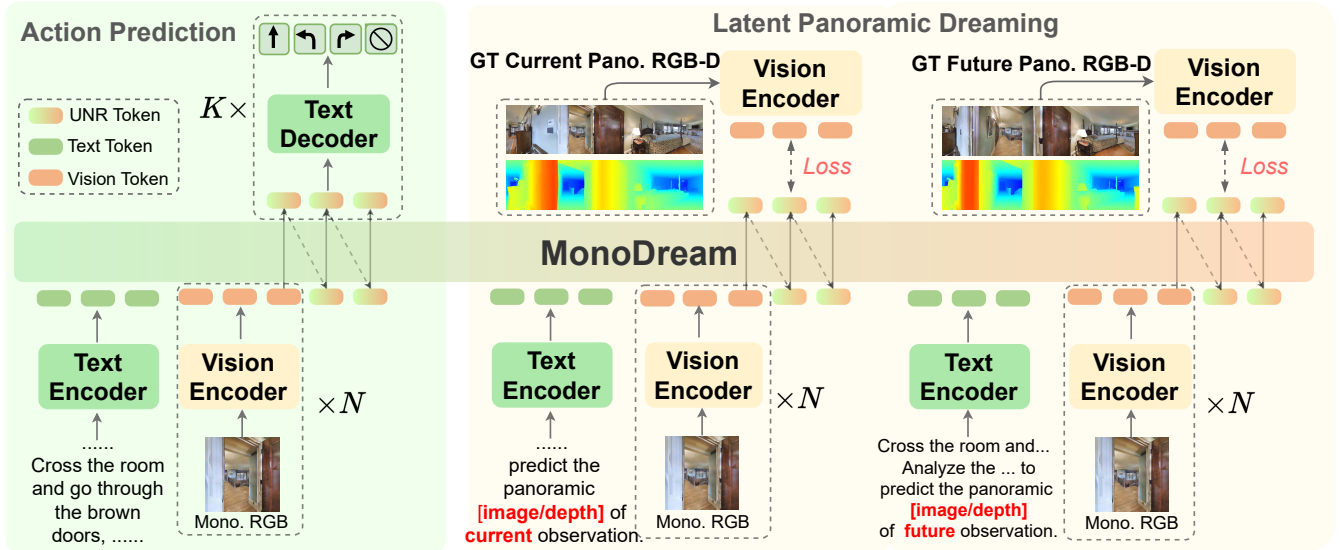


Figure 1: Overview of the MonoDream framework. MonoDream employs a Vision-Language Action (VLA) framework to encode both visual observations and textual instructions into a Unified Navigation Representation (UNR). The Action Prediction task generates the next action in natural language and is trained with action loss. The Latent Panoramic Dreaming (LPD) encourages the model to internally imagine the latent features of panoramic RGB-D images of current and future steps, providing global visual and geometric context via feature-based loss. This multi-task co-training enables monocular agents to reason beyond the limited field of view and make more informed navigation decisions.

sions from limited monocular input. Our contributions are as follows:

- We propose **MonoDream**, a novel monocular VLN framework that enhances the agent’s internal global-aware ability. MonoDream enables the agent to infer implicit global, geometric and temporal context with monocular images.
- We introduce two key components: a Unified Navigation Representation that jointly encodes navigation actions and latent global scene, and Latent Panoramic Dreaming tasks that supervise UNR learning from current and future panoramic RGB-D latent features.
- We demonstrate the effectiveness of MonoDream by achieving state-of-the-art performance on the monocular VLN-CE benchmark, including R2R-CE and RxR-CE, while using a smaller VLA model without external training data. Cross-dataset evaluations further validate the strong generalization ability of our approach.

Related Work

Vision Language Navigation

Vision-Language Navigation tasks (Nguyen et al. 2019; Wang et al. 2022; Wu et al. 2020; Schumann et al. 2024) challenge an embodied agent to follow natural language instructions and reach a target location within a previously unseen environment. Although early approaches largely focus on discrete navigation settings (Hong et al. 2021; Chen et al. 2021; Liu et al. 2023), where the agent moves along a predefined graph and typically has access to panoramic RGB-D observations, more recent studies (Hong et al. 2022; Wang et al. 2023; An

et al. 2022; Dai et al. 2024) have shifted toward more realistic, continuous environments to predict low-level actions, and some works also begin to investigate monocular RGB-D settings.

Recent works based on large models (Liu et al. 2025; Zhang et al. 2024a; Cheng et al. 2024; Wei et al. 2025) leverage large-scale, RGB-only video models to build monocular VLN systems with enhanced generalization and real-world applicability. However, agents in these works rely on a single forward-facing RGB camera, which poses additional challenges due to the limited field of view and the absence of explicit depth information. Some works (Wang et al. 2024b; Wang, Lee, and Lee 2025) attempt to use neural rendering techniques to recover observations from novel viewpoints from the reconstructed map. However, these methods rely on additional localization and mapping modules, and their performance still falls significantly short compared to approaches that directly utilize the panoramic input.

To tackle the limitation, MonoDream introduces LPD tasks, predicting the latent features of global and geometric information as auxiliary tasks, to enhance the performance of monocular VLN systems. We demonstrate that our monocular navigation model with the unified navigation representation, equipped with panoramic, depth, and future awareness through auxiliary supervision, achieves state-of-the-art performance in monocular settings.

Vision Imagination in Embodied Agents

Learning to imagine vision information has become an increasingly popular strategy for improving policy learning in embodied agents, particularly when the agent operates un-

der egocentric and limited monocular observations. In such settings, inferring spatial layouts, geometric structures, and future dynamics beyond the current observation becomes essential for decision-making. Early approaches (Wu et al. 2023; Wang et al. 2024a) often adopt a two-stage pipeline, where future states are predicted using off-the-shelf world models (Zheng et al. 2024; Bar et al. 2025) or neural rendering (Mildenhall et al. 2021; Kerbl et al. 2023) from reconstructed maps. Although simple and modular, these methods are inherently constrained by the accuracy and generalizability of the underlying world models. More recent works (Cen et al. 2025; Tian et al. 2024; Chen et al. 2024) advocate for end-to-end training paradigms that integrate forecasting and action planning within a unified framework. These methods allow the agent to jointly predict future observations and actions, leading to improved performance and generalization. Notably, DreamVLA (Zhang et al. 2025) introduces a future-aware policy that forecasts dynamics, depth, and semantics simultaneously from monocular RGB input, demonstrating the benefits of future modeling in visual language tasks.

Our proposed MonoDream is the first VLN framework to learn a unified latent representation supervised by global panoramic scene layouts, while using only monocular RGB images as input. This panoramic-aware latent supervision equips the agent with a holistic spatial understanding essential for navigation, and leads to state-of-the-art performance under the monocular setting.

Method

Overview. We study monocular Vision-and-Language Navigation in Continuous Environments (VLN-CE), where an agent navigates in realistic spaces based on natural language instructions. At each step t , the agent receives three types of inputs: the natural language instruction \mathcal{I} , the current egocentric RGB observation o_t , and a sampled history of past observations \mathcal{O}_t . The primary objective is to predict the next navigation action a_t based on these multimodal inputs.

The framework of our proposed MonoDream (denoted as π_θ with the network parameters as θ) is illustrated in Figure 1, which is built based on a Vision Language Model (VLM), including a vision encoder, a text encoder and decoder, and an LLM-based backbone. MonoDream constructs the Unified Navigation Representation, a shared latent space produced by the VLM backbone, which is designed to jointly align navigation actions and the agent’s internal estimation of the implicit global information, including panoramic visual and structural cues for the current and future observations, based on only monocular inputs. To supervise the formation of such representations, we introduce a set of auxiliary tasks termed Latent Panoramic Dreaming. These tasks are only applied during training, where the agent is guided to predict the latent features of: (1) the panoramic RGB-D observation at the current step, and (2) the panoramic RGB-D of a future step. Through this joint design, MonoDream learns to align action decisions with imagined global and future context, thereby significantly enhancing monocular navigation performance.

Unified Navigation Representation

To enable the agent to internalize global and future scene awareness from limited monocular observations, we propose the Unified Navigation Representation that captures rich navigation-relevant information in a compact latent space. UNR aligns navigation-relevant information together, including implicit features of actions, panoramic scene layout, panoramic depth perception, and future dynamics, into a shared latent space. It can be decoded into navigation actions, instruction-like text, or directly as navigation-relevant features. In this section, we detail how language and visual observations are encoded into a shared feature space, which serves as the output of our MonoDream backbone.

For the input, the language instruction \mathcal{I} is tokenized and processed by a text encoder, $\Phi_{\text{text}}(\cdot)$, to produce text feature $E_{\text{text}} \in \mathbb{R}^{L \times d}$, where L is the text sequence length and d is the hidden dimension.

$$E_{\text{text}} = \Phi_{\text{text}}(\mathcal{I}) \quad (1)$$

The agent’s visual input contains the current view o_t and a sampled historical images. To maintain computational efficiency while preserving essential historical context, we uniformly sample N frames, denoted as $\mathcal{O}_t = \{o_{p_0}, \dots, o_{p_{N-1}}\}$, from the full sequence of past observations. Each of these images is independently processed by the vision encoder, $\Phi_{\text{vis}}(\cdot)$, which encodes each image into d -dimensional feature and these individual image features are collected to form the visual input sequence E_{vis} . E_{vis} represents the agent’s complete visual context at step t from monocular image sequence.

$$E_{\text{vis}} = \{\Phi_{\text{vis}}(o_p), \dots, \Phi_{\text{vis}}(o_{p_{N-1}}), \Phi_{\text{vis}}(o_t)\} \quad (2)$$

After encoding, the text feature E_{text} and the vision feature E_{vis} are combined into a single input sequence S_t for the LLM-based backbone of MonoDream:

$$S_t = [E_{\text{text}}, E_{\text{vis}}] \quad (3)$$

This sequence is then fed into the backbone and the backbone outputs the hidden states $h_t \in \mathbb{R}^{l_{\text{seq}} \times d}$, where l_{seq} is the length of the output sequence. This representation encapsulates the comprehensive state of the agent at step t :

$$h_t = \text{MonoDream-Backbone}(S_t) \quad (4)$$

We define hidden feature h_t as the Unified Navigation Representation. To ensure that h_t encodes all navigation-relevant information, even under limited monocular observations, we design a multitask training framework. In this framework, MonoDream learns h_t to capture both vision-based signals (global layout, geometric structure, and future dynamics) and language-based signals (actions and instructions). Specifically, we jointly train the model with multiple objectives: Latent Panoramic Dreaming (Section), action prediction and instruction reasoning (Section).

Latent Panoramic Dreaming

To supervise the learning of the proposed UNR h_t , we introduce a set of auxiliary tasks collectively referred to as Latent

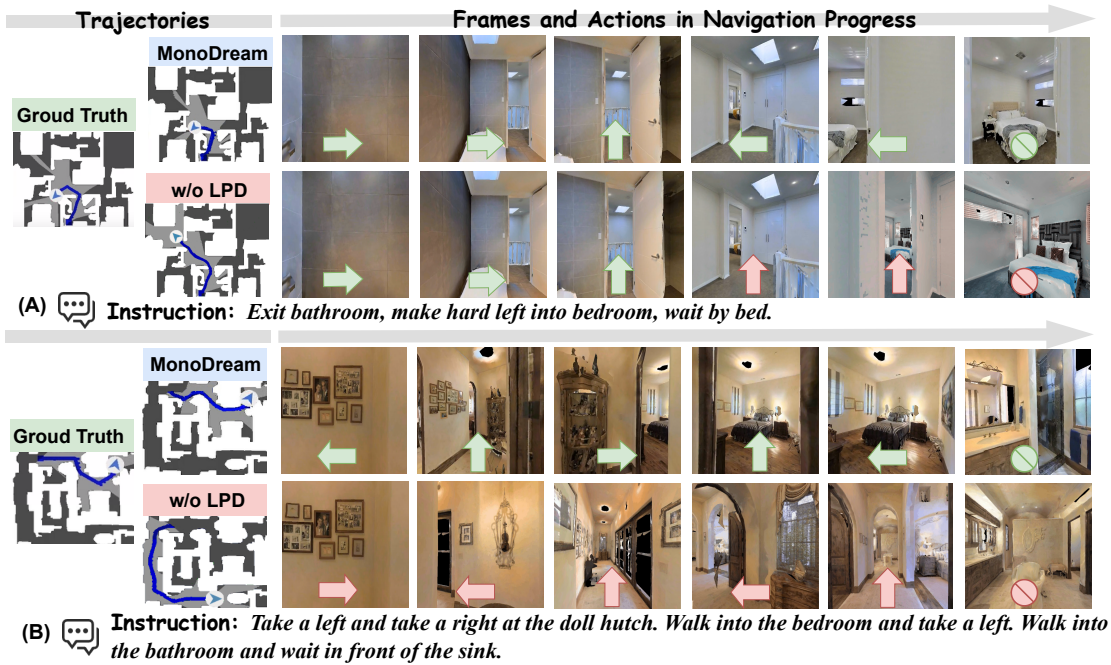


Figure 2: Quantitative results of MonoDream. We compare MonoDream with the ablated variant w/o LPD. Green arrows indicate correct actions, and red arrows indicate errors. (A) MonoDream correctly identifies the hard turning point at the fourth frame. In contrast, the w/o LPD baseline misreads the hallway layout, proceeds straight, and stops in the wrong room. (B) The w/o LPD model makes a critical mistake at the very first step, while MonoDream by leveraging internalized global features from LPD, correctly turns left even without explicit corner information in the initial monocular view.

Panoramic Dreaming. These tasks guide the agent to enrich h_t with the latent feature of the global scene, depth geometry, and future information, all without requiring explicit scene reconstruction.

First, we supervise the agent to align h_t with the latent feature of the global and geometric context of the current scene from its limited forward-facing view. This encourages the learned UNR h_t to capture a more complete spatial understanding beyond the monocular observation.

- $\mathcal{H}_t^{\text{PI}}$: Panoramic RGB at step t .
- $\mathcal{H}_t^{\text{PD}}$: Panoramic depth at step t .

In addition, we guide the agent to anticipate how the scene evolves by predicting the latent features of the next-step panoramic RGB and depth, i.e. $\mathcal{H}_{t+1}^{\text{PI}}$ and $\mathcal{H}_{t+1}^{\text{PD}}$. This future-aware supervision encourages h_t to incorporate short-term temporal dynamics and improves the agent’s ability to plan ahead during navigation.

These latent supervision signals are obtained by encoding the corresponding panoramic RGB and depth using a vision encoder that shares weights with the visual encoder Φ_{vis} , which is also used to encode the monocular input images. This unified encoder design ensures that the supervision features and the learned UNR h_t align in the same feature space and benefit from joint training.

The LPD training objective minimizes the mean squared error (MSE) between h_t and the above latent targets. The

feature loss function is:

$$L_t^{\text{Fea}}(\theta) = \sum_{m \in \mathcal{M}} \|h_t - \mathcal{H}_t^m\|^2 + \sum_{m \in \mathcal{M}} \|h_t - \mathcal{H}_{t+1}^m\|^2 \quad (5)$$

where $\mathcal{M} = \{\text{PI}, \text{PD}\}$.

Through this auxiliary supervision, LPD encourages the Unified Navigation Representation h_t to internalize implicit global semantics, geometric layout, and short-term future cues. Our design improves the agent’s holistic understanding of the environment, leading to more informed and anticipatory navigation decisions.

Notably, the LPD module is used only during training. Panoramic RGB and depth signals are not required during inference, where the agent predicts actions based on monocular RGB images with the internalized global awareness.

Multi-Task Co-Training

In addition to the LPD task, which aligns the UNR h_t with visual features, we further supervise h_t using two language-based objectives: action prediction and instruction reasoning. These tasks align h_t with the linguistic representation of navigation behaviors and goals, ensuring the learned UNR effectively bridges both visual and textual modalities.

At each time step t , the model predicts navigation actions in the natural language form. Specially, the model predicts a sequence of the next three actions (a_t, a_{t+1}, a_{t+2}) based on the embedding h_t , generated from the instruction I , current observation o_t , and navigation history \mathcal{O}_t for the sample at

step t . This setup encourages short-term action forecasting while retaining reactivity to new observations. The training objective for each timestamp t is defined as:

$$L_t^{Act}(\theta) = - \sum_{k=0}^K \log \pi_{\theta}(a_{t+k}^* | h_t) \quad (6)$$

where a_{t+k}^* denotes the ground-truth action at step $t+k$.

To further enhance the agent’s understanding of instructions, we introduce instruction reasoning as another auxiliary task, which infers the underlying language instruction \mathcal{I} based on the agent’s visual trajectory context, effectively promoting multimodal alignment from vision to language.

For each training trajectory, we uniformly sample N visual observations to construct the input sequence O_t . These images are processed by the MonoDream backbone to produce the UNR h_t . A text decoder is then applied to generate the distribution over instruction tokens based on h_T , and the loss is computed accordingly:

$$L_{\tau}^{Ins}(\theta) = - \log \pi_{\theta}(\mathcal{I}_{\tau} | h_T) \quad (7)$$

During training, we co-train the action prediction and all the auxiliary tasks and switch between different tasks by changing the prompt (see our supplementary material for the prompts of different tasks). The final loss function is:

$$L = \sum_{\tau \in D} \left(\sum_t^{T_{\tau}} (L_t^{Act}(\theta) + \lambda L_t^{Fca}(\theta)) + L_{\tau}^{Ins}(\theta) \right) \quad (8)$$

where D is the set of training trajectories, T_{τ} is the step number of trajectory τ , and λ is the hyperparameter of weight.

Implementation Details

Action Design The action space of the agent is designed into four categories: move forward, turn left, turn right, and stop. The forward action includes step sizes of 25 cm, 50 cm, and 75 cm, while the turn actions are parameterized by rotation angles of 15°, 30°, and 45°. This fine-grained design allows for more precise and flexible control, which is critical in complex environments.

Training Datasets All the training data used in our work are collected from simulated environments, including the training splits of R2R-CE (Krantz et al. 2020) and RxR-CE (Ku et al. 2020), as well as additional data collected using the DAgger (Ross, Gordon, and Bagnell 2011) strategy. We first construct step-wise navigation data based on the action annotations provided in R2R-CE and RxR-CE, resulting in 320K and 600K samples respectively. In addition, we construct auxiliary supervision data based on R2R-CE by applying the aforementioned image preprocessing and instruction reasoning tasks. Moreover, following the DAgger strategy (Ross, Gordon, and Bagnell 2011), we further collect 500K step-wise samples from non-oracle trajectories generated in the R2R-CE training environments.

For the training panoramic images in LPD tasks, we adapt the cubemap (Trindade and Raposo 2011) format and split into four canonical directions (left, front, right, and back). For the depth image, we apply log scaling to the raw values to compress large-scale variations, followed by a colormap-based RGB rendering (Itseez 2015).

Model Training We adopt NVILA-lite-2B (Liu et al. 2024) as our base model, which includes a SigLIP vision encoder (Tschannen et al. 2025), a projection module, and a Qwen2-based language model (Bai et al. 2025). Starting from NVILA-lite-2B’s pre-trained weights, we perform supervised fine-tuning for our tasks. All components of the model are trainable during fine-tuning. Training is conducted on 8 NVIDIA H20 GPUs for 5 epochs with the learning rate of 1e-5, the warm-up as 0.03, and the batch size of 80. During training and inference, we set the number of future actions to predict (K) to 3, and the number of historical frames to sample (N) to 8.

Experimental Results

Experiment Setup

Simulated environments We evaluate our method on the VLN-CE benchmarks R2R-CE (Krantz et al. 2020) and RxR-CE (Ku et al. 2020) following the standard VLN-CE settings. All the methods are evaluated on the R2R val-unseen split and RxR val-unseen split. The quantitative results are in Figure 2.

Metrics We follow the standard VLN evaluation protocol (Krantz et al. 2020; Ku et al. 2020) to evaluate navigation performance for all methods, including success rate (SR), oracle success rate (OSR), success weighted by path length (SPL), and the navigation error from the goal (NE). Among them, SR and SPL are widely regarded as the primary metrics, reflecting the task completion and path efficiency respectively.

Comparison on VLN-CE Benchmarks

We evaluate our method on the VLN-CE benchmarks including R2R-CE and RxR-CE, which provide continuous environments for navigational actions in reconstructed photorealistic indoor scenes. We first focus on the val-unseen split in R2R-CE dataset in Table 1. To ensure a fair comparison, we group methods based on their sensor settings and mark those that do not rely on large language models (†).

As in Table 1, our proposed MonoDream achieves strong performance without relying on any external data beyond the simulation datasets. This demonstrates the data efficiency of our method, which learns effective navigation behaviors. We attribute this success primarily to the incorporation of LPD supervision. By jointly training on panoramic images, panoramic depth maps, and future panoramic views, the model learns to build a more comprehensive understanding of the scene structure and spatial layout. This enhances the model’s generalization ability, even when trained solely on limited simulator-generated data. The synergy between action prediction and LPD further amplifies the effectiveness of each training sample, enabling the model to extract richer visual cues and improve data efficiency.

To further evaluate MonoDream’s robustness and long-term ability, we conduct experiments on the RxR-CE Val-Unseen split (Ku et al. 2020), as in Table 2. Compared to R2R-CE, RxR-CE presents significantly more complex challenges due to its longer trajectories and more natural, diverse language instructions. Despite these challenges and the monocular setting, MonoDream achieves state-of-the-art

Method	Observation			R2R Val-Unseen				Training
	S.RGB	Depth	Pano.	NE ↓	OSR ↑	SR ↑	SPL ↑	External Data
BEVBert [†] (An et al. 2022)		✓	✓	4.57	67.0	59.0	50.0	-
ETPNav [†] (An et al. 2024)		✓	✓	4.71	65.0	57.0	49.0	-
ENP-ETPNav [†] (Liu, Wang, and Yang 2024)		✓	✓	4.69	65	58	50	-
Seq2Seq [†] (Krantz et al. 2020)	✓	✓		7.77	37.0	25.0	22.0	-
CMA [†] (Krantz et al. 2020)	✓	✓		7.37	40.0	32.0	30.0	-
LAW [†] (Raychaudhuri et al. 2021)	✓	✓		6.83	44.0	35.0	31.0	-
CM2 [†] (Georgakis et al. 2022)	✓	✓		7.02	41.0	34.0	27.0	-
WS-MGMap [†] (Chen et al. 2022)	✓	✓		6.28	47.0	38.0	34.0	-
sim2real [†] (Wang et al. 2024b)	✓	✓		5.95	55.8	44.9	30.4	-
NavMorph [†] (Yao, Gao, and Xu 2025)	✓	✓		5.75	56.9	47.9	33.2	-
NaVid-4D (Liu et al. 2025)	✓	✓		5.99	55.7	43.8	37.1	1500K
NaVid (Zhang et al. 2024b)	✓			5.47	49.1	37.4	35.9	0K
Uni-NaVid(Zhang et al. 2024a)	✓			5.58	53.3	47.0	42.7	2300K
NaVILA(Cheng et al. 2024)	✓			5.22	<u>62.5</u>	54.0	<u>49.0</u>	2215K
Aux-Think (Wang et al. 2025)	✓			5.49	62.9	<u>55.7</u>	48.7	1600K
MonoDream (Ours)	✓			<u>5.45</u>	61.5	55.8	49.1	0K

Table 1: Comparison of different methods on the R2R-CE Val-Unseen split. Observations used include single RGB camera (S.RGB), depth sensor (Depth) and panoramic view (Pano.). [†] indicates methods without using LLMs. External data refers to sources beyond the navigation simulator, such as real-world web data, general VQA datasets, and other similar resources.

results on the primary metrics (SR and SPL), outperforming strong baselines like Uni-NaVid (Zhang et al. 2024a) and NaVILA (Cheng et al. 2024) while using substantially less training data. These findings demonstrate the advantage of leveraging visual imagination as auxiliary supervision.

Cross-dataset Evaluation We assess the generalization capability of MonoDream on the RxR-CE Val-Unseen split (Table 3). Notably, all the methods are trained without any RxR-CE training data and our method achieves the state-of-the-art performance.

The effectiveness of our model in a cross-dataset setting highlights three key insights. First, the use of LPD strengthens the backbone’s ability to form a richer and more generalizable representation of the environment. Second, our LPD tasks exhibit strong generalization capabilities, with their benefits extending beyond the distribution of the training set. Third, for long-horizon navigation tasks in RxR-CE, having a global understanding or imagination of the environment significantly enhances the agent’s navigation capabilities.

Ablation Study

In the ablation study, we separately evaluate the effectiveness of the two auxiliary tasks, LPD and Instruction Reasoning, as well as the individual contributions of the four sub-tasks in LPD. To ensure efficiency, all ablation experiments are conducted using only the R2R-CE training split and test on R2R-CE val-unseen split only.

Impact of Auxiliary Tasks We conduct ablation experiments to evaluate the contribution of the two auxiliary tasks: Instruction Reasoning and LPD. As shown in Table 4, removing both auxiliary tasks leads to the poorest performance.

Method	RxR Val-Unseen			
	NE↓	OSR↑	SR↑	SPL↑
CMA [†] (Hong et al. 2022)	8.76	-	26.5	22.1
VLNBERT [†] (Hong et al. 2022)	8.98	-	27.0	22.6
Seq2Seq [†] (Krantz et al. 2020)	11.8	-	13.9	11.9
sim2real [†] (Wang et al. 2024b)	8.79	36.7	25.5	18.1
Uni-NaVid(Zhang et al. 2024a)	6.24	<u>55.5</u>	48.7	<u>40.9</u>
NaVILA(Cheng et al. 2024)	6.77	-	<u>49.3</u>	44.0
MonoDream (Ours)	<u>6.38</u>	55.8	49.4	<u>40.9</u>

Table 2: Comparison of different methods on the RxR Val-Unseen split. [†] indicates methods without using LLMs.

Method	RxR Val-Unseen			
	NE↓	OSR↑	SR↑	SPL↑
Seq2Seq(Krantz et al. 2020)	11.8	5.02	3.51	3.43
CMA(Krantz et al. 2020)	11.7	10.7	4.41	2.47
LAW(Raychaudhuri et al. 2021)	10.87	21.0	8.0	8.0
CM2(Georgakis et al. 2022)	8.98	25.3	14.4	9.2
WS-MGMap(Chen et al. 2022)	9.83	29.8	15.0	12.1
A ² NAV(Chen et al. 2023)	-	-	16.8	6.3
NaVid(Zhang et al. 2024b)	8.41	<u>34.5</u>	<u>23.8</u>	<u>21.2</u>
MonoDream (Ours)	<u>8.57</u>	35.9	25.1	21.6

Table 3: Cross-dataset performance on the RxR-CE Val-Unseen split. All results are obtained without training on the RxR-CE training set.

Introducing IR alone yields moderate gains across all metrics, demonstrating that aligning the UNR h_t with language-based supervision (i.e., instruction prediction) enhances the model’s understanding of high-level navigation goals.

The most substantial improvement comes from LPD. Introducing LPD achieves the best performance across all metrics. These results confirm the complementary benefits of language- and vision-oriented supervision: IR improves semantic goal alignment, while LPD enriches UNR with global, geometric, and predictive spatial cues that are otherwise missing in monocular observations.

Auxiliary Tasks		Metrics			
IR	LPD	NE↓	OSR↑	SR↑	SPL↑
		7.78	43.7	35.1	30.2
✓		<u>7.67</u>	<u>45.8</u>	<u>37.7</u>	<u>32.1</u>
✓	✓	6.19	51.1	46.1	39.9

Table 4: Ablation study on auxiliary tasks. IR: Instruction Reasoning; LPD: Latent Panoramic Dreaming

Impact of Auxiliary Panoramic Dreaming Tasks Table 5 presents the ablation study of the four auxiliary tasks in LPD, including the latent feature of panoramic RGB image (PI), panoramic depth (PD), future panoramic RGB image (FPI), and future panoramic depth (FPD). We add each component on top of the baseline and evaluate their individual and cumulative contributions to navigation performance.

The results show that each LPD task brings consistent improvements, confirming their effectiveness as auxiliary supervision signals. In particular, latent panoramic and depth provide the most significant gains. We attribute this to their role in enhancing the agent’s spatial and structural understanding of the environment, while implicit depth introduces geometric awareness that complements RGB observations.

LPD Tasks				Metrics			
PI	PD	FPI	FPD	NE↓	OSR↑	SR↑	SPL↑
				7.67	45.8	37.7	32.1
✓				7.22	44.2	39.6	35.3
	✓			6.71	47.4	42.2	37.7
		✓		7.03	44.5	39.1	34.2
			✓	6.80	45.3	39.8	35.4
✓	✓	✓	✓	6.19	51.1	46.1	39.9

Table 5: Ablation study on four LPD tasks. PI: Panoramic RGB image; PD: Panoramic depth; FPI: Future panoramic RGB image; FPD: Future panoramic depth.

Impact of Prediction Steps in Latent Panoramic Dreaming We further investigate the effect of predicting the future panoramic feature of different future steps in LPD, and we take FPI as a representative example. As shown in Table 6, predicting a single future step delivers the best performance. Extending the prediction horizon to 2 or 3 future steps leads to performance degradation, due to compounded uncertainty in long-range forecasting under partial observations.

#step	NE↓	OSR↑	SR↑	SPL↑
1	7.03	44.5	39.1	34.2
2	<u>7.89</u>	<u>43.9</u>	<u>34.3</u>	<u>29.1</u>
3	8.77	40.0	29.2	26.2

Table 6: Ablation on the prediction steps in LPD.

Model Efficiency

In addition to navigation performance, we compare the efficiency of our method MonoDream with recent state-of-the-art VLN methods. Specifically, we report the size of model parameters and the average inference time per step on a single NVIDIA 4090 GPU with the same local hardware.

As in Table 7, our method achieves smaller model size and fastest inference speed among compared approaches. Despite introducing auxiliary latent imagination tasks during training, MonoDream is inference-efficient as auxiliary modules are disabled at test time. This makes MonoDream a practical solution for real-time embodied navigation applications.

Method	Params.	Time / Step
NaVILA (Cheng et al. 2024)	8B	1.2s
Aux-Think (Wang et al. 2025)	8B	1.2s
MonoDream (Ours)	2B	0.8s

Table 7: Comparison of model efficiency.

Conclusion

In this work, we present MonoDream, a unified framework for vision-language navigation (VLN) that leverages a vision-language model as the backbone to integrate action prediction and latent imagination. We propose the Unified Navigation Representation to align all navigation-relevant information and use LPD task to co-train the UNR. MonoDream achieves SOTA performance on the monocular VLN-CE benchmark by implicitly modeling panoramic understanding, depth perception, and future view prediction by latent supervision.

Our model training is conducted using only simulator-based data from the R2R-CE and RxR-CE datasets, without relying on external data. Our experiments demonstrate that MonoDream can achieve strong performance using only monocular inputs, narrowing the gap with methods that explicitly consume panoramic or depth data. This work highlights the potential of implicit multimodal learning as a scalable and efficient solution for embodied navigation agents using only monocular inputs.

Limitations. MonoDream imagines the current scene and immediate future from past monocular observations, without explicitly reconstructing panoramic history. Incorporating richer temporal modeling, such as memory-based reasoning, may further improve planning and robustness.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant No. 62441617. It was

supported by the Postdoctoral Fellowship Program and China Postdoctoral Science Foundation under Grant No. 2024M764093 and Grant No. BX20250485, the Beijing Natural Science Foundation under Grant No. 4254100, the Fundamental Research Funds for the Central Universities under Grant No. KG16336301, and by Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing. Shuo Wang was supported in part by the Outstanding Innovative Talents Cultivation Funded Programs 2024 of Renmin University of China.

References

- An, D.; Qi, Y.; Li, Y.; Huang, Y.; Wang, L.; Tan, T.; and Shao, J. 2022. Bevbort: 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.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and Van Den Hengel, A. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3674–3683.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Bar, A.; Zhou, G.; Tran, D.; Darrell, T.; and LeCun, Y. 2025. Navigation world models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15791–15801.
- Cen, J.; Yu, C.; Yuan, H.; Jiang, Y.; Huang, S.; Guo, J.; Li, X.; Song, Y.; Luo, H.; Wang, F.; et al. 2025. WorldVLA: Towards Autoregressive Action World Model. *arXiv preprint arXiv:2506.21539*.
- Chen, P.; Ji, D.; Lin, K.; Zeng, R.; Li, T.; Tan, M.; and Gan, C. 2022. Weakly-supervised multi-granularity map learning for vision-and-language navigation. *Advances in Neural Information Processing Systems*, 35: 38149–38161.
- Chen, P.; Sun, X.; Zhi, H.; Zeng, R.; Li, T. H.; Liu, G.; Tan, M.; and Gan, C. 2023. A2 Nav: Action-Aware Zero-Shot Robot Navigation by Exploiting Vision-and-Language Ability of Foundation Models. *arXiv preprint arXiv:2308.07997*.
- Chen, S.; Guhur, P.-L.; Schmid, C.; and Laptev, I. 2021. History aware multimodal transformer for vision-and-language navigation. *Advances in neural information processing systems*, 34: 5834–5847.
- Chen, Y.; Ge, Y.; Li, Y.; Ge, Y.; Ding, M.; Shan, Y.; and Liu, X. 2024. Moto: Latent motion token as the bridging language for robot manipulation. *arXiv preprint arXiv:2412.04445*, 8.
- 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*.
- Dai, G.; Zhao, J.; Chen, Y.; Qin, Y.; Zhao, H.; Xie, G.; Yao, Y.; Shu, X.; and Li, X. 2024. Unitedvln: Generalizable gaussian splatting for continuous vision-language navigation. *arXiv preprint arXiv:2411.16053*.
- 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.
- Gu, J.; Stefani, E.; Wu, Q.; Thomason, J.; and Wang, X. E. 2022. Vision-and-language navigation: A survey of tasks, methods, and future directions. *arXiv preprint arXiv:2203.12667*.
- Hong, Y.; Wang, Z.; Wu, Q.; and Gould, S. 2022. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15439–15449.
- Hong, Y.; Wu, Q.; Qi, Y.; Rodriguez-Opazo, C.; and Gould, S. 2021. Vln bert: A recurrent vision-and-language bert for navigation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 1643–1653.
- Itseez. 2015. Open Source Computer Vision Library. <https://github.com/opencv/opencv>.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Krantz, J.; Wijmans, E.; Majumdar, A.; Batra, D.; and Lee, S. 2020. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*, 104–120. Springer.
- Ku, A.; Anderson, P.; Patel, R.; Ie, E.; and Baldrige, J. 2020. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. *arXiv preprint arXiv:2010.07954*.
- Liu, H.; Wan, W.; Yu, X.; Li, M.; Zhang, J.; Zhao, B.; Chen, Z.; Wang, Z.; Zhang, Z.; and Wang, H. 2025. NaVid-4D: Unleashing Spatial Intelligence in Egocentric RGB-D Videos for Vision-and-Language Navigation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 10607–10615. IEEE.
- Liu, R.; Wang, W.; and Yang, Y. 2024. Vision-Language Navigation with Energy-Based Policy. *arXiv preprint arXiv:2410.14250*.
- Liu, R.; Wang, X.; Wang, W.; and Yang, Y. 2023. Bird’s-eye-view scene graph for vision-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10968–10980.
- Liu, Z.; Zhu, L.; Shi, B.; Zhang, Z.; Lou, Y.; Yang, S.; Xi, H.; Cao, S.; Gu, Y.; Li, D.; et al. 2024. NVILA: Efficient frontier visual language models. *arXiv preprint arXiv:2412.04468*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.

- Nguyen, K.; Dey, D.; Brockett, C.; and Dolan, B. 2019. Vision-based navigation with language-based assistance via imitation learning with indirect intervention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12527–12537.
- 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 Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 4018–4028. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Robertson, C. E.; Hermann, K. L.; Mynick, A.; Kravitz, D. J.; and Kanwisher, N. 2016. Neural representations integrate the current field of view with the remembered 360 panorama in scene-selective cortex. *Current Biology*, 26(18): 2463–2468.
- 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.
- Schumann, R.; Zhu, W.; Feng, W.; Fu, T.-J.; Riezler, S.; and Wang, W. Y. 2024. Velma: Verbalization embodiment of llm agents for vision and language navigation in street view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18924–18933.
- Seeber, M.; Stangl, M.; Vallejo Martelo, M.; Topalovic, U.; Hiller, S.; Halpern, C. H.; Langevin, J.-P.; Rao, V. R.; Fried, I.; Eliashiv, D.; et al. 2025. Human neural dynamics of real-world and imagined navigation. *Nature Human Behaviour*, 9(4): 781–793.
- Tian, Y.; Yang, S.; Zeng, J.; Wang, P.; Lin, D.; Dong, H.; and Pang, J. 2024. Predictive inverse dynamics models are scalable learners for robotic manipulation. *arXiv preprint arXiv:2412.15109*.
- Trindade, D. R.; and Raposo, A. B. 2011. Improving 3D navigation in multiscale environments using cubemap-based techniques. In *Proceedings of the 2011 ACM Symposium on Applied Computing*, 1215–1221.
- Tschannen, M.; Gritsenko, A.; Wang, X.; Naeem, M. F.; Alabdulmohsin, I.; Parthasarathy, N.; Evans, T.; Beyer, L.; Xia, Y.; Mustafa, B.; et al. 2025. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*.
- Wang, H.; Liang, W.; Gool, L. V.; and Wang, W. 2022. Towards versatile embodied navigation. *Advances in neural information processing systems*, 35: 36858–36874.
- Wang, J.; Zhang, Q.; Sun, J.; Cao, J.; Han, G.; Zhao, W.; Zhang, W.; Shao, Y.; Guo, Y.; and Xu, R. 2024a. Reinforcement learning with generalizable gaussian splatting. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 435–441. IEEE.
- Wang, S.; Wang, Y.; Li, W.; Cai, X.; Wang, Y.; Chen, M.; Wang, K.; Su, Z.; Li, D.; and Fan, Z. 2025. Aux-Think: Exploring Reasoning Strategies for Data-Efficient Vision-Language Navigation. *arXiv preprint arXiv:2505.11886*.
- Wang, Z.; Lee, S.; and Lee, G. H. 2025. Dynam3D: Dynamic Layered 3D Tokens Empower VLM for Vision-and-Language Navigation. *arXiv preprint arXiv:2505.11383*.
- Wang, Z.; Li, X.; Yang, J.; Liu, Y.; and Jiang, S. 2023. Gridmm: Grid memory map for vision-and-language navigation. In *Proceedings of the IEEE/CVF International conference on computer vision*, 15625–15636.
- Wang, Z.; Li, X.; Yang, J.; Liu, Y.; and Jiang, S. 2024b. Sim-to-real transfer via 3d feature fields for vision-and-language navigation. *arXiv preprint arXiv:2406.09798*.
- Wei, M.; Wan, C.; Yu, X.; Wang, T.; Yang, Y.; Mao, X.; Zhu, C.; Cai, W.; Wang, H.; Chen, Y.; et al. 2025. StreamVLN: Streaming Vision-and-Language Navigation via SlowFast Context Modeling. *arXiv preprint arXiv:2507.05240*.
- Wu, H.; Jing, Y.; Cheang, C.; Chen, G.; Xu, J.; Li, X.; Liu, M.; Li, H.; and Kong, T. 2023. Unleashing large-scale video generative pre-training for visual robot manipulation. *arXiv preprint arXiv:2312.13139*.
- Wu, Q.; Gong, X.; Xu, K.; Manocha, D.; Dong, J.; and Wang, J. 2020. Towards target-driven visual navigation in indoor scenes via generative imitation learning. *IEEE Robotics and Automation Letters*, 6(1): 175–182.
- Wu, W.; Chang, T.; Li, X.; Yin, Q.; and Hu, Y. 2024. Vision-language navigation: a survey and taxonomy. *Neural Computing and Applications*, 36(7): 3291–3316.
- Yao, X.; Gao, J.; and Xu, C. 2025. NavMorph: A Self-Evolving World Model for Vision-and-Language Navigation in Continuous Environments. *arXiv preprint arXiv:2506.23468*.
- Zhang, J.; Wang, K.; Wang, S.; Li, M.; Liu, H.; Wei, S.; Wang, Z.; Zhang, Z.; and Wang, H. 2024a. Uni-NaVid: A Video-based Vision-Language-Action Model for Unifying Embodied Navigation Tasks. *arXiv preprint arXiv:2412.06224*.
- Zhang, J.; Wang, K.; Xu, R.; Zhou, G.; Hong, Y.; Fang, X.; Wu, Q.; Zhang, Z.; and Wang, H. 2024b. Navid: Video-based vlm plans the next step for vision-and-language navigation. *arXiv preprint arXiv:2402.15852*.
- Zhang, W.; Liu, H.; Qi, Z.; Wang, Y.; Yu, X.; Zhang, J.; Dong, R.; He, J.; Wang, H.; Zhang, Z.; et al. 2025. DreamVLA: A Vision-Language-Action Model Dreamed with Comprehensive World Knowledge. *arXiv preprint arXiv:2507.04447*.
- Zheng, Z.; Peng, X.; Yang, T.; Shen, C.; Li, S.; Liu, H.; Zhou, Y.; Li, T.; and You, Y. 2024. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*.