

# ActiveVLN: Towards Active Exploration via Multi-Turn RL in Vision-and-Language Navigation

Zekai Zhang<sup>1,\*</sup>, Weiye Zhu<sup>1,\*</sup>, Hewei Pan<sup>1</sup>, Xiangchen Wang<sup>1</sup>, Rongtao Xu<sup>2,3</sup>, Xing Sun<sup>4</sup>, Feng Zheng<sup>1,2,†</sup>

**Abstract**—The Vision-and-Language Navigation (VLN) task requires an agent to follow natural language instructions and navigate through complex environments. Existing MLLM-based VLN methods primarily rely on imitation learning (IL) and often use DAgger for post-training to mitigate covariate shift. While effective, these approaches incur substantial data collection and training costs. Reinforcement learning (RL) offers a promising alternative. However, prior VLN RL methods lack dynamic interaction with the environment and depend on expert trajectories for reward shaping, rather than engaging in open-ended active exploration. This restricts the agents ability to discover diverse and plausible navigation routes. To address these limitations, we propose ActiveVLN, a VLN framework that explicitly enables active exploration through multi-turn RL. In the first stage, a small fraction of expert trajectories is used for IL to bootstrap the agent. In the second stage, the agent iteratively predicts and executes actions, automatically collects diverse trajectories, and optimizes multiple rollouts via the GRPO objective. To further improve RL efficiency, we introduce a dynamic early-stopping strategy to prune long-tail or likely failed trajectories, along with additional engineering optimizations. Experiments show that ActiveVLN achieves the largest performance gains over IL baselines compared to both DAgger-based and prior RL-based post-training methods, while reaching competitive performance with state-of-the-art approaches despite using a smaller model.

## I. INTRODUCTION

Vision-and-Language Navigation (VLN) [1], [2] requires a navigation agent to interpret natural language instructions and navigate toward target locations. With the recent progress of multimodal large language models (MLLMs), VLN has entered a new stage. By leveraging their seamless integration of language understanding and visual perception, MLLM-based approaches have achieved substantial improvements over conventional VLN models. Currently, VLN methods using a monocular RGB camera [3]–[6] mostly adopt this paradigm, employing MLLMs as the backbone for action decision making.

Conventional MLLM-based VLN methods primarily adopt imitation learning (IL), where agents mimic expert demonstrations to map observations to actions. However, IL suffers from covariate shift, where errors accumulate when the agent encounters states unseen during training, leading to poor generalization. To mitigate this, many works [3]–[6]

employ DAgger [7] post-training, which augments training with expert corrections but still requires substantial expert data collection and retraining. Reinforcement learning (RL) offers a natural alternative by enabling agents to learn from self-generated trajectories. Yet, prior RL approaches [8], [9] lack dynamic interaction with the environment, relying on expert trajectories for reward shaping instead of engaging in open-ended active exploration, which limits their exploratory capacity and yields only marginal gains over IL.

To address these limitations, we propose ActiveVLN, a VLN framework that explicitly enable active exploration through multi-turn RL. ActiveVLN consists of two training stages. In the first stage, IL is used for bootstrapping, but with far less expert data compared with existing IL-based approaches. In the second stage, the agent performs multi-turn RL, iteratively refining its navigation policy through active exploration without relying on expert trajectories. Specifically, the agent iteratively predicts actions, executes them in the environment, and appends new observations to its input sequence until termination. Multiple rollouts are then optimized with a GRPO objective. To further enhance RL efficiency, we propose a dynamic early-stop strategy that prunes long-tail or likely failed trajectories, along with several engineering optimizations. Beyond the framework, we provide a comprehensive analysis of critical design choices in VLN RL, such as reward shaping and multi-turn paradigm, offering insights into why prior RL approaches often underperform.

We evaluate ActiveVLN on the R2R [1] and RxR [10] benchmarks in continuous environments. The active exploration process yields substantial gains (+11.6 SR on R2R and +9.7 SR on RxR over IL baselines), enabling ActiveVLN to surpass both DAgger-based and prior RL-based post-training methods, while also achieving competitive performance with state-of-the-art approaches using a smaller model, less training time, and lower data collection costs.

In summary, the main contributions are:

- We propose ActiveVLN, a VLN framework that leverages active exploration to refine navigation policies, allowing agents to learn from diverse, self-generated trajectories with minimal expert supervision.
- We develop several optimization techniques, including a dynamic early-stopping strategy, to improve RL training efficiency.
- We show that ActiveVLN achieves the largest performance gains over IL baselines compared to both DAgger-based and prior RL-based post-training methods and reaching competitive performance with state-

\* Equal contribution.

† Corresponding author.

This research is supported by the National Key Research and Development Program of China (2024YFE0203100).

<sup>1</sup>Southern University of Science and Technology, China.

<sup>2</sup>Spatialtemporal AI, China.

<sup>3</sup>MBZUAI, UAE.

<sup>4</sup>Tencent Youtu Lab.

of-the-art approaches despite using a smaller model.

## II. RELATED WORK

### A. Vision-and-Language Navigation

Vision-and-Language Navigation (VLN) tasks require a navigation agent to navigate through an environment to a target location by following natural language instructions [1]. Early VLN studies [11]–[14] primarily focused on discrete environments, where the navigation space is abstracted as a topological graph. In this setting, agents can only move between predefined waypoints on the graph [1], [15]–[17]. While this abstraction simplifies planning, it significantly diverges from real-world navigation, where both actions and observations are continuous. To address this gap, Vision-and-Language Navigation in Continuous Environments (VLN-CE) has been proposed [2], [18]–[20]. In VLN-CE, instead of jumping between discrete waypoints, the agent executes low-level actions such as moving forward or rotating. This setting makes it a more realistic and challenging task for embodied navigation.

### B. Imitation Learning in VLN

Early studies [21]–[25] explored training-free approaches that incorporate large language models (LLMs) or multimodal large language models (MLLMs) into navigation pipelines, leveraging their strong capabilities in language understanding and visual perception. However, such modular frameworks often suffer from misalignment between components, and their performance heavily depends on the pre-trained LLMs or MLLMs. To overcome these limitations, recent work [3]–[5] has investigated end-to-end training with LLM [26]–[28] and MLLM backbones [29], [30], where imitation learning (IL) has become the dominant paradigm. In this setting, navigation policies are trained on large-scale expert trajectories to map visual observations to low-level actions. Yet, naive IL suffers from covariate shift. Since the agent is only exposed to expert trajectories during training, small errors at test time can accumulate when encountering unseen states, leading to compounding deviations from the goal. To mitigate this, DAGger [7] has been widely adopted [3]–[6], where additional expert corrections are collected on visited states to improve generalization and robustness. Despite their effectiveness, IL-based methods remain heavily dependent on large-scale expert data, restricting the agents ability to explore diverse and plausible routes. They also incur substantial data collection and retraining costs, as each iteration requires new annotations and policy updates. In contrast, our approach empowers the navigation agent to explore actively, enabling it to discover plausible navigation behaviors beyond expert demonstrations, without requiring manual trajectory collection or retraining from scratch.

### C. Reinforcement Learning in VLN

Reinforcement learning (RL) has been successful in improving the reasoning of large language models (LLMs), with methods like Group Relative Policy Optimization (GRPO) being particularly effective [31], [32]. Unlike other methods

such as Proximal Policy Optimization (PPO) [33], GRPO simplifies the process by using group-wise reward normalization, which eliminates the need for a separate value model and reduces computational costs. Despite its success in LLMs, the use of RL in Vision-and-Language Navigation (VLN) is still limited. Current approaches often rely on expert trajectories for reward shaping and lack dynamic interaction with the environment [8], [9], which restricts the agent’s ability to explore different paths. To address these limitations, we use a multi-turn RL approach that treats VLN as a sequential decision process [34], [35]. This method enables agents to learn optimal policies through direct interaction with the simulator, without needing expert trajectories.

## III. TASK FORMULATION

Each navigation episode is defined by a tuple  $(s_0, \mathcal{I}, s^*)$ , where  $s_0$  is the agents initial state (position and orientation),  $\mathcal{I}$  is a natural language instruction to follow, and  $s^*$  is the target location. At each time step  $t$ , the agent receives a forward-facing monocular RGB image  $\mathbf{v}_t$  as its observation and maintains a history  $\mathbf{h}_t = \{\mathbf{v}_{1:t}, a_{1:t-1}\}$  of past observations and actions. A parameterized policy  $\pi_\theta$  selects a low-level action  $a_t \in \mathcal{A}$  based on the instruction and current history:

$$a_t \sim \pi_\theta(a_t | \mathcal{I}, \mathbf{h}_t). \quad (1)$$

The action space is defined as:

$$\mathcal{A} = \{ \text{FORWARD}, \text{LEFT}, \text{RIGHT}, \text{STOP} \}, \quad (2)$$

where FORWARD moves the agent forward by a fixed step size, and LEFT/RIGHT rotate the agent by a fixed angle, both determined by the task configuration. An episode terminates when the agent issues STOP. Success is defined as stopping within 3 meters of the goal  $s^*$ .

## IV. METHOD

In this section, we first present our formulation of VLN as a multi-turn process (Section IV-A). We then describe the ActiveVLN framework (Section IV-B), which enables learning from self-generated trajectories through active exploration (Section IV-B.2) and employs a dynamic early-stopping strategy to enhance RL training efficiency (Section IV-B.3). Finally, we provide engineering details for further acceleration of RL training (Section IV-C).

### A. Multi-Turn Paradigm

Following video-based MLLMs, most prior end-to-end VLN models adopt the *single-turn paradigm* (Figure 2a), where each action is predicted from the instruction and past observations:

$$a_t \sim \pi_\theta(a_t | \mathcal{I}, \mathcal{V}_{1:t}). \quad (3)$$

In contrast, we adopt the *multi-turn paradigm* (Figure 2b), where actions are modeled autoregressively from both past observations and actions:

$$a_t \sim \pi_\theta(a_t | \mathcal{I}, \{(\mathcal{V}_i, a_i)\}_{i=1}^{t-1}, \mathcal{V}_t). \quad (4)$$

## Stage 1: Learning to imitate expert trajectories

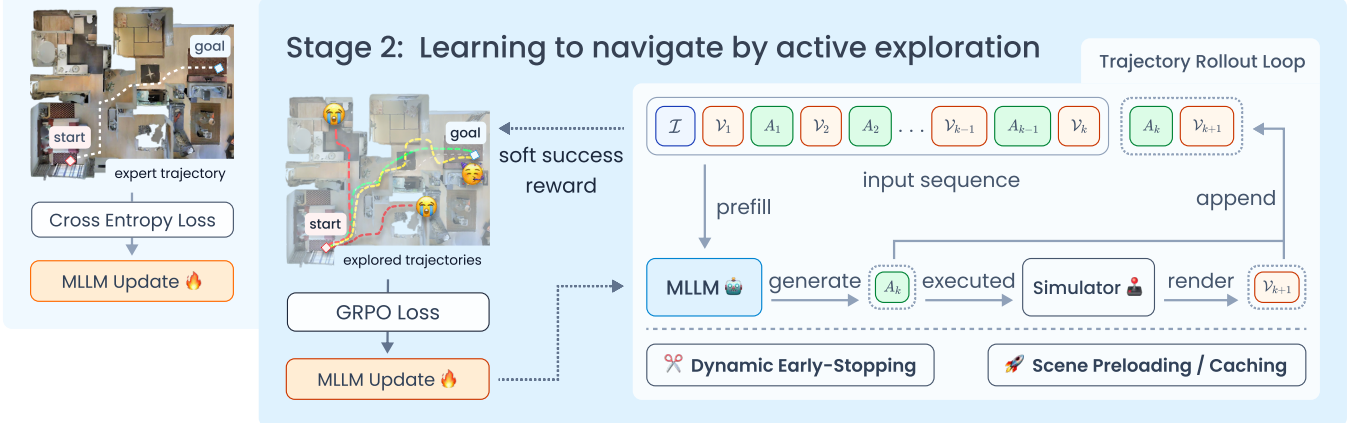
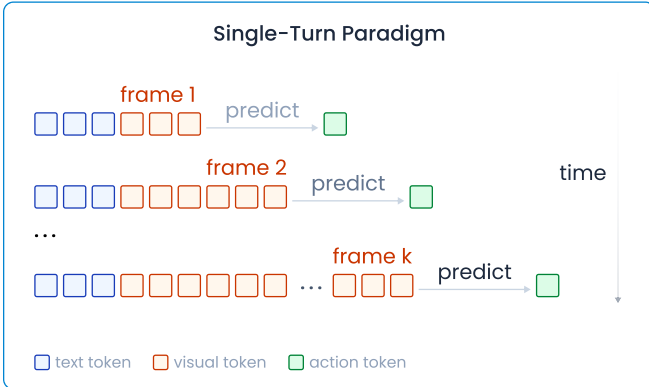



Fig. 1: Overview of ActiveVLN. In Stage 1, ActiveVLN performs imitation learning (IL) using expert trajectories. In Stage 2, it conducts multi-turn reinforcement learning (RL), autonomously collecting trajectories in the simulator, receiving rewards that encourage progress toward the goal, and updating the policy via GRPO. Key components, including the *dynamic early-stopping strategy*, *scene preloading*, and *scene caching*, are incorporated to ensure efficient training during RL.



(a) Single-turn paradigm: each action is predicted from the instruction and past observations only.



(b) Multi-turn paradigm: actions are generated autoregressively from the instruction, past observations, and past actions. This allows training signals from future steps to backpropagate and refine earlier predictions.

Fig. 2: Comparison between the single-turn and multi-turn paradigms.

This paradigm offers several advantages. First, it naturally enables KV-cache reuse for efficient inference. Second, while the single-turn paradigm breaks actions within the same episode into independent pieces, the multi-turn formulation allows actions to be packed together, accelerating training. Most importantly, it enables gradients associated with trajectory outcomes to be propagated to all preceding actions. We find this property to be crucial for the success of RL in VLN

(see Section V-F.2).

### B. ActiveVLN Framework

1) *Compact Action Prediction*: The raw VLN-CE action space comprises four primitive actions: FORWARD, TURN LEFT, TURN RIGHT, and STOP. Prior work augments this space by merging consecutive actions of the same type [3]–[5]. For instance, three FORWARD steps (each 25cm) can be merged into a single FORWARD 75cm action. This augmentation diversifies action granularity and shortens episode length, improving training efficiency. Building on this, we adopt action chunking to further reduce episode length, where the agent predicts up to three future actions at once rather than a single action per step.

2) *Active Exploration*: Since MLLMs are not pre-trained for VLN tasks, we start with imitation learning (IL) using a small number of expert demonstrations to initialize the navigation policy. The IL objective is:

$$\mathcal{L}_{\text{IL}} = - \sum_t \sum_{i=1}^{n_t} \log P(a_t^i | \mathcal{I}, \mathcal{H}_{<t}, a_t^{<i}), \quad (5)$$

where  $\mathcal{I}$  is the navigation instruction,  $\mathcal{H}_{<t}$  is the interaction history up to step  $t$ , and  $a_t^{<i}$  are the actions already generated in the current chunk. Formally, the history is:

$$\mathcal{H}_{<t} = \{\mathcal{V}_1, A_1, \mathcal{V}_2, A_2, \dots, \mathcal{V}_{t-1}, A_{t-1}, \mathcal{V}_t\}, \quad (6)$$

where  $\mathcal{V}_t$  is the observation at step  $t$ , and  $A_t = \{a_t^1, a_t^2, \dots, a_t^{n_t}\}$  is a sequence of low-level actions.

IL provides a good starting point, but it has a key limitation: the agent only learns to mimic expert trajectories. Once it encounters unfamiliar situations, it has no mechanism to recover or adapt. To overcome this, we introduce active exploration through reinforcement learning (RL). Here, the agent is no longer restricted to expert data. Instead, it can interact with the environment on its own: given an instruc-

tion, it predicts an action, executes it in the simulator, and observes the outcome. By repeating this loop until it issues a `stop` action (or an exception occurs), the agent actively generates diverse trajectories, learns from its successes and failures, and gradually improves its policy.

For optimization, we use Group Relative Policy Optimization (GRPO) [31]. GRPO samples  $G$  candidate trajectories for the same instruction and compares them within the group. Trajectories that perform better than the group average are reinforced, while weaker ones are suppressed. The RL objective is:

$$\mathcal{L}_{\text{RL}} = \mathbb{E}_{\{o_i\}} \left[ \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \min \left( \frac{\pi_{\theta}}{\pi_{\theta_{\text{old}}}} A_{i,t}, \text{clip} \left( \frac{\pi_{\theta}}{\pi_{\theta_{\text{old}}}}, 1 - \epsilon, 1 + \epsilon \right) A_{i,t} \right) \right], \quad (7)$$

where  $o_i$  is the  $i$ -th trajectory with length  $|o_i|$ ,  $\pi_{\theta}$  and  $\pi_{\theta_{\text{old}}}$  are the new and old policies,  $\epsilon$  is the clipping parameter, and  $A_{i,t}$  is the estimated advantage.

We use a *soft success reward*:

$$R = 15 \cdot \mathbb{I}(\text{success}) \cdot \frac{d_{\text{goal}}}{3}, \quad (8)$$

where  $d_{\text{goal}}$  is the geodesic distance to the goal. The indicator  $\mathbb{I}(\text{success}) = 1$  if the agent issues a valid `stop` within 3 meters of the goal, and 0 otherwise. In this way, the agent is no longer just imitating experts but is encouraged to explore actively, discover different ways of reaching the goal, and improve by trial and error. This self-driven learning process is the key to stronger generalization in unseen environments.

3) *Dynamic Early-Stopping Strategy*: Trajectory rollout time can account for over half of total RL training time, making it the primary bottleneck. We observe that excessively long trajectories often dominate this time, and in most cases correspond to unsuccessful attempts in which the agent wanders aimlessly or explores irrelevant regions.

To address this issue, we introduce a *dynamic early-stopping strategy* that adaptively terminates unpromising rollouts. Specifically, a trajectory is stopped and marked as failed once its length exceeds a threshold  $T_{\text{max}}$ , defined as:

$$T_{\text{max}} = \alpha_{\text{roll}} \cdot |\tau^*|, \quad (9)$$

where  $|\tau^*|$  is the length of the oracle (expert trajectory), and  $\alpha_{\text{roll}} > 1$  is a tolerance factor that specifies how much deviation from the oracle is acceptable (we set  $\alpha_{\text{roll}} = 2$  in our experiments).

This strategy avoids wasting computation on hopeless rollouts, while maintaining a balance between preventing excessive exploration and avoiding overly strict cutoffs, ultimately leading to more efficient training.

### C. Engineering Details in RL

To further accelerate RL training, we adopt several techniques: *Scene caching* stores frequently accessed scene data in memory, enabling faster loading when the same scene is revisited. *Scene preloading* pipelines scene loading with

policy updates, reducing idle time during training. These techniques cut down scene-loading overhead and improve training efficiency. In addition, similar to [35], we decouple the simulator from the training server by deploying it as a standalone HTTP server, which allows scalable and parallel execution of multiple navigation environments.

## V. EXPERIMENTS

### A. Simulation Benchmark and Evaluation Metrics

We evaluate the R2R and RxR val-unseen splits on Matterport3D [36] scenes in continuous environments using the Habitat simulator [37]. Since these scenes are unseen during training, the task requires generalization to novel environments. Following prior work, we report several VLN evaluation metrics, including success rate (SR), oracle success rate (OSR), success weighted by path length (SPL), navigation error (NE) and and normalize dynamic time wrapping (nDTW) [38]. An episode is considered successful if the agent issues the `STOP` action within 3m of the goal.

### B. Real-World Robot Settings

We also validate our method in real world using a wheeled humanoid robot. RGB observations are captured using the Intel RealSense D435i camera mounted on the robots head, with the height adjusted to approximately 1.25 m to match the simulation setting. The trained navigation policy is served remotely and the robot continuously streams observations to the server and executes the control commands returned in response.

### C. Implementation Details

1) *Training Data*: We utilize a total of 167.6k expert trajectories to bootstrap the policy with IL, including 10.8k from R2R [1], 146.3k from R2R-EnvDrop [39], and 10.5k from RxR [10]. To ensure data quality, we filter out RxR samples with excessively long instructions or paths. For the subsequent RL post-training, we sample 4k navigation instructions from R2R and RxR, respectively. At this stage, ground-truth actions are not required. Instead, the agent actively engages in the exploration of diverse navigation routes.

2) *Hyper-Parameters*: We use Qwen2.5-VL-3B-Instruct [40] as the base model with the AdamW optimizer [41]. During the IL stage, only the connector and language model are updated while the vision encoder remains frozen. The peak learning rate is  $2 \times 10^{-5}$  with a global batch size of 64. In the RL post-training stage, all parameters are trainable. Each batch includes 8 prompts, each of which would generate 4 trajectory rollouts. The KL coefficient is set to 0.0, the peak learning rate to  $1 \times 10^{-6}$ , and the clipping parameter to 0.28. We employ dynamic sampling [42] with at most 2 attempts to ensure at least one successful rollout per group. Each action chunk contains up to three merged actions. All experiments are conducted on 4 NVIDIA L40S GPUs, requiring about 20 hours for IL and RL training, respectively.

Method (Model Size)	Stage	R2R Val-Unseen				
		NE↓	OS↑	SR↑	SPL↑	ΔSR
<b>DAgger-based post-training</b>						
NaVid (7B) [3]	IL	5.90	46.7	33.1	30.7	
	IL + DAgger	5.50	49.1	37.4	35.9	+4.3
UniNaVid (7B) [4]	IL	-	-	-	-	
	IL + DAgger	5.58	53.3	47.0	42.7	-
MapNav <sup>†</sup> (7B) [43]	IL	6.38	38.2	23.9	19.5	
	IL + DAgger	6.02	46.1	33.5	30.7	+9.6
StreamVLN <sup>†</sup> (7B) [6]	IL	6.05	53.8	45.5	41.6	
	IL + DAgger	5.47	57.8	<b>50.8</b>	<b>45.7</b>	+5.3
<b>RL-based post-training</b>						
VLN-R1 (7B) [8]	IL	7.92	37.1	24.9	17.5	
	IL + RL	7.00	41.2	30.2	21.8	+5.3
<b>ActiveVLN (3B, Ours)</b>	IL	6.83	42.7	38.5	33.7	
	IL + DAgger	6.11	50.2	45.4	41.5	+6.9
	IL + RL	<b>5.31</b>	<b>58.3</b>	50.1	43.7	<b>+11.6</b>

TABLE I: Comparison of state-of-the-art methods on the R2R Val-Unseen split before and after post-training (DAgger or RL). ΔSR indicates the improvement in success rate (SR) after post-training. † denotes results obtained using only VLN-specific data to ensure a fair comparison.

Method (Model Size)	Training data		RxR Val-Unseen			
	VLN	Others	NE↓	SR↑	SPL↑	nDTW↑
NaVILA (8B) [5]	-	✓	6.77	49.3	44.0	58.8
UniNaVid (7B) [4]	2.4M	✓	6.24	48.7	40.9	-
StreamVLN (7B) [6]	1M	✓	6.22	<b>52.9</b>	<b>46.0</b>	<b>61.9</b>
ActiveVLN (3B, w/o RL)	177K	✗	7.25	41.0	34.1	-
<b>ActiveVLN (3B, Ours)</b>	177K	✗	<b>5.84</b>	50.7	41.2	58.1

TABLE II: Comparison of state-of-the-art methods on the RxR Val-Unseen split. The columns under Training data indicate the amount of VLN data used and whether additional datasets (e.g. visual question-answering) were included.

#### D. Comparisons with SOTA Methods

Table I compares our method with SOTA approaches on the R2R Val-Unseen split, before and after post-training (DAgger or RL). For fairness, results are reported under the same dataset setting (using only VLN-specific data) whenever available in the original papers. DAgger consistently improves SR across models, but its effect diminishes once the IL baseline is already strong. In contrast, our active exploration delivers the largest performance gain while maintaining SOTA performance, despite requiring less training time and a smaller model. Specifically, SR improves from 38.5 (IL) to 50.1 (IL+RL), yielding a remarkable +11.6 gain that surpasses larger models. We also apply DAgger to our IL baseline, and it yields smaller SR gains than our active exploratory RL approach, while requiring additional expert trajectories and longer training time (26h vs. 20h). These results underscore the effectiveness and efficiency of active exploratory RL post-training, enabling policy refinement without reliance on expert corrections.

Table II reports results on the RxR benchmark. Unlike prior methods that are trained not only on VLN data but also on additional data sources such as visual question answer-

ing [44], [45] and human touring data [5], [46], ActiveVLN is trained solely on 177K VLN expert trajectories without any extra data. Despite this constraint, ActiveVLN achieves the lowest navigation error (NE=5.84) and a competitive success rate (SR=50.7) compared to StreamVLN, which relies on substantially more data. This highlights that our method achieves strong generalization while remaining data-efficient. Notably, after incorporating active exploration (w/ RL), the success rate (SR) improves by 9.7. The slightly lower SPL can be attributed to the nature of active exploration, which prioritizes autonomous trajectory discovery over strict imitation of expert demonstrations.

#### E. Qualitative Results

1) *Results in Simulation*: Figure 3 illustrates a representative example where the IL-only baseline fails to reach the goal, whereas the policy with RL post-training succeeds. Specifically, the IL policy does not execute the instructed left turn to navigate around the brown couch. In contrast, the RL-enhanced policy follows the correct path, passes the couch on the appropriate side, and reaches the target successfully.

2) *Results in Real-World Scenarios*: We deploy ActiveVLN in diverse real-world scenes, including offices, hallway and laboratories. Representative qualitative results are shown in Figure 4, where ActiveVLN successfully completes navigation tasks. Full demonstrations are available in the supplementary video.

#### F. Ablations

We analyze key design choices in RL post-training, including reward design (Section V-F.1) and the single-turn vs. multi-turn paradigm (Section V-F.2). We then assess the effect of dynamic early-stopping (Section V-F.3) and other optimization techniques (Section V-F.4) on training efficiency.

1) *Reward Design*: We explore alternative reward designs beyond the *soft success reward*. Note that we do not incorporate dynamic sampling to accelerate training.

The *hard success reward* is defined as:

$$R_{\text{hard}} = 15 \cdot \mathbb{I}(\text{success}), \quad (10)$$

which eliminates the distance-to-goal normalization used in the soft success reward.

The *path alignment reward* is formulated as:

$$R_{\text{align}} = w_{\text{succ}} \cdot \mathbb{I}(\text{success}) \cdot \frac{d_{\text{to-goal}}}{3} + w_{\text{nDTW}} \cdot \text{nDTW}, \quad (11)$$

where  $w_{\text{succ}} + w_{\text{nDTW}} = 15$ .

This introduces an additional nDTW term that captures the similarity between the predicted and ground-truth trajectories.

As shown in Table III, the soft success reward yields the best overall performance. In contrast, the hard success reward, which removes distance-to-goal normalization, leads to a reduced SR, likely because the normalization encourages the agent to continue moving closer to the goal when it is

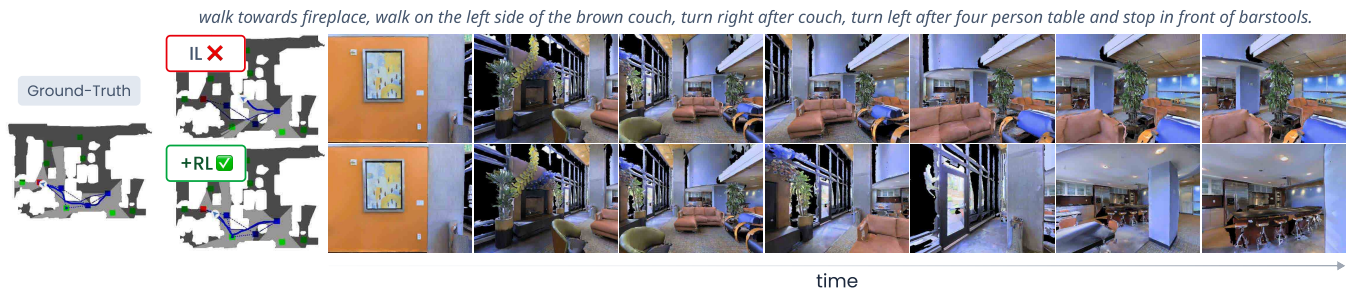


Fig. 3: Qualitative results of ActiveVLN in simulation. We compare ActiveVLN with its ablated variant without RL post-training. The IL-only baseline fails to execute the instructed left turn around the brown couch and eventually gets stuck. In contrast, ActiveVLN with RL post-training successfully follows the given instructions, passes on the correct side of the couch, and reaches the goal near the bar stools.

*In front of you there are several chairs. Move to the chair on the right with a black bag and stop nearby.*



*You are facing a table. Turn around and walk forward along the cabinet on your left. When you see the green plant, walk a few more steps and stop in front of it.*



Fig. 4: Qualitative results of ActiveVLN in real-world scenarios, including a laboratory environment and an office workspace. Additional examples are provided in the supplementary video.

Reward Type	R2R Val-Unseen			
	NE↓	OS↑	SR↑	SPL↑
soft success reward	5.56	<b>56.4</b>	<b>48.1</b>	<b>40.6</b>
hard success reward	5.68	53.6	46.2	39.9
path align. reward ( $w_{\text{succ}} : w_{\text{nDTW}} = 3:1$ )	<b>5.49</b>	51.9	46.3	40.1
path align. reward ( $w_{\text{succ}} : w_{\text{nDTW}} = 2:1$ )	5.64	48.6	44.4	40.0
path align. reward ( $w_{\text{succ}} : w_{\text{nDTW}} = 1:1$ )	5.61	49.6	44.4	40.3

TABLE III: Performance under different reward designs.

nearby. Incorporating the nDTW term via the path alignment reward also leads to performance degradation. Interestingly, the path alignment reward with  $w_{\text{succ}} : w_{\text{nDTW}} = 3:1$  achieves the lowest navigation error (NE), suggesting that the nDTW term helps the agent better mimic expert trajectories. However, this comes at the cost of reduced SR, indicating that overemphasizing trajectory similarity can hinder goal completion. This trend is further supported by the fact that increasing the weight on nDTW (from 3:1 to 1:1) leads to a consistent decline in both OS and SR. These results suggest that the soft success reward, while simple, is an effective training signal that encourages both goal reaching and diverse and effective pathfinding.

2) *Effect of the Multi-Turn Paradigm:* To investigate the effect of the multi-turn paradigm, we construct a single-turn baseline and apply RL under the same training data and key hyper-parameter settings as used in the multi-turn version of ActiveVLN. The results are summarized in Table IV. After IL, the single-turn baseline achieves a slightly higher

Stage	R2R Val-Unseen					
	IL	RL	NE↓	OS↑	SR↑	SPL↑
<b>Multi-turn Paradigm</b>						
✓		6.83	42.7	38.5	33.7	
✓	✓	<b>5.31</b> (-1.52)	<b>58.3</b> (+15.6)	<b>50.1</b> (+11.6)	<b>43.7</b> (+10.0)	
<b>Single-turn Paradigm</b>						
✓		6.44	48.1	39.7	35.3	
✓	✓	6.25 (-0.19)	48.1 (+0.0)	40.3 (+0.6)	37.0 (+1.7)	

TABLE IV: Performance under the multi-turn and single-turn paradigms. RL post-training provides substantial improvements across all metrics in the multi-turn paradigm, while yielding only marginal gains in the single-turn paradigm.

SR than the multi-turn variant. A plausible explanation is that the single-turn paradigm processes inputs as short video segments, which aligns more closely with the base models pretraining on large-scale video datasets. In contrast, the multi-turn paradigm involves longer-horizon, interleaved image-text sequences, which are less common in the pretrained distribution and may initially limit performance. However, the multi-turn paradigm demonstrates substantially larger performance improvements after RL. This indicates that RL particularly benefits from the multi-turn formulation, as it enables policy optimization over action histories. While single-turn may initially appear advantageous, multi-turn ultimately provides a stronger foundation for learning complex navigation strategies under RL.

3) *Effect of Dynamic Early-Stopping Strategy*: Figure 6d shows that about 10% of trajectory rollouts are forcefully terminated due to excessive length, reducing the average rollout time by about 10% (Table V). Figure 5 compares the episode length distributions of rollout and expert trajectories. As expected, rollout distributions concentrate around the green reference line (expert trajectories) and become sparser with increasing deviation. The blue line indicates the threshold set by our early-stopping strategy. These results demonstrate that the dynamic early-stopping strategy effectively prunes long-tail and failed trajectories while retaining most successful ones, improving both training efficiency and stability.




Fig. 5: Comparison of rollout and expert trajectory lengths.

Optimization Technique	Avg. Rollout Time (s)
ActiveVLN	<b>106.9</b>
w/o Dynamic Early-Stopping	118.6 (+10.9%)
w/o Scene Pre-Loading	111.9 (+4.68%)
w/o Scene Caching	109.3 (+2.24%)

TABLE V: Ablation study of optimization techniques on rollout time.

4) *Effect of Other Optimization Techniques*: To evaluate the effect of other optimization techniques on RL training efficiency, we measure the average rollout time per step over 25 training steps. The results are reported in Table V, where each technique is ablated in turn. It shows that removing any of them slows down training, with the dynamic early-stopping strategy providing the largest reduction in rollout time. These underscore the importance of our proposed optimizations in significantly improving RL training efficiency.

### G. Training Dynamics

To gain insights into policy behavior during RL post-training, we monitor several key metrics throughout training, as shown in Figure 6. Figure 6a shows that reward steadily increases, converging around step 300, reflecting improved performance. This trend aligns with the success rate in Figure 6b, evaluated on R2R Val-Unseen every 50 steps, which peaks around step 350 before slightly declining. Figure 6c reports episode length (number of actions per



Fig. 6: Training dynamic of ActiveVLN.

trajectory). After an initial decrease due to format errors causing premature terminations (steps 1-50), episode length gradually approaches the IL baseline as the policy learns correct action formatting. Additionally, Figure 7 presents top-down visualizations of a group of rollout trajectories during RL post-training, qualitatively illustrating the benefits of incorporating active exploration: the agent learns to follow diverse paths that deviate from expert demonstrations while still successfully reaching the goal.

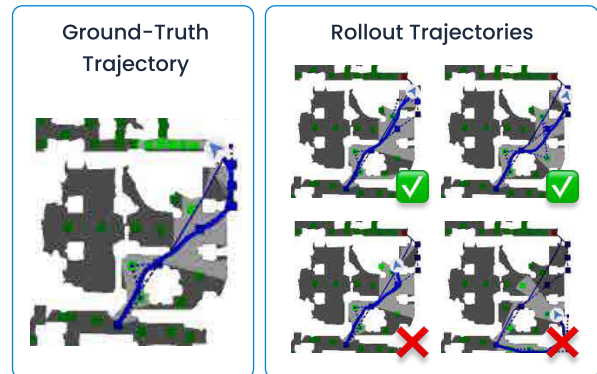


Fig. 7: Top-down views of ground-truth and rollout trajectories. The top two cases show successful navigation via alternative paths deviating from the expert trajectory. The bottom two illustrate failures: stopping too far from the goal (bottom left) and moving in the wrong direction from the start (bottom right).

## VI. CONCLUSION

In this work, we present ActiveVLN, a VLN framework that explicitly enables active exploration through multi-turn RL. This active exploration process allows agents to learn from diverse, self-generated trajectories without heavy reliance on expert supervision. We further introduce a dynamic early-stopping strategy and several engineering optimizations to improve RL efficiency. Experiments show that ActiveVLN

achieves the largest performance gains over IL baselines compared to both DAGger-based and prior RL-based post-training methods and reaching competitive performance with state-of-the-art approaches despite using a smaller model. These results highlight the critical role of active exploration in VLN and validate our RL post-training framework. Future work may explore more advanced exploration strategies and extend ActiveVLN to other embodied tasks.

## REFERENCES

- [1] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. Van Den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *CVPR*, 2018.
- [2] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, "Beyond the nav-graph: Vision and language navigation in continuous environments," in *ECCV*, 2020.
- [3] J. Zhang, K. Wang, R. Xu, G. Zhou, Y. Hong, X. Fang, Q. Wu, Z. Zhang, and H. Wang, "Navid: Video-based vlm plans the next step for vision-and-language navigation," in *RSS*, 2024.
- [4] J. Zhang, K. Wang, S. Wang, M. Li, H. Liu, S. Wei, Z. Wang, Z. Zhang, and H. Wang, "Uni-navid: A video-based vision-language-action model for unifying embodied navigation tasks," in *RSS*, 2025.
- [5] A.-C. Cheng, Y. Ji, Z. Yang, X. Zou, J. Kautz, E. Biyik, H. Yin, S. Liu, and X. Wang, "Navila: Legged robot vision-language-action model for navigation," in *RSS*, 2025.
- [6] M. Wei, C. Wan, X. Yu, T. Wang, Y. Yang, X. Mao, C. Zhu, W. Cai, H. Wang, Y. Chen, *et al.*, "Streamvln: Streaming vision-and-language navigation via slowfast context modeling," *arXiv preprint arXiv:2507.05240*, 2025.
- [7] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *AISTATS*, 2011.
- [8] Z. Qi, Z. Zhang, Y. Yu, J. Wang, and H. Zhao, "Vln-r1: Vision-language navigation via reinforcement fine-tuning," *arXiv preprint arXiv:2506.17221*, 2025.
- [9] C. Gao, L. Jin, X. Peng, J. Zhang, Y. Deng, A. Li, H. Wang, and S. Liu, "Octonav: Towards generalist embodied navigation," *arXiv preprint arXiv:2506.09839*, 2025.
- [10] A. Ku, P. Anderson, R. Patel, E. Ie, and J. Baldridge, "Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding," in *EMNLP*, 2020.
- [11] Z. Wang, J. Li, Y. Hong, Y. Wang, Q. Wu, M. Bansal, S. Gould, H. Tan, and Y. Qiao, "Scaling data generation in vision-and-language navigation," in *CVPR*, 2023.
- [12] D. An, H. Wang, W. Wang, Z. Wang, Y. Huang, K. He, and L. Wang, "Etpnav: Evolving topological planning for vision-language navigation in continuous environments," *TPAMI*, 2024.
- [13] Y. Hong, Z. Wang, Q. Wu, and S. Gould, "Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation," in *CVPR*, 2022.
- [14] Y. Qiao, Z. Yu, and Q. Wu, "Vln-petl: parameter-efficient transfer learning for vision-and-language navigation," in *ICCV*, 2023.
- [15] S. Chen, P.-L. Goh, M. Tapaswi, C. Schmid, and I. Laptev, "Think global, act local: Dual-scale graph transformer for vision-and-language navigation," in *CVPR*, 2022.
- [16] H. Wang, W. Wang, W. Liang, C. Xiong, and J. Shen, "Structured scene memory for vision-language navigation," in *CVPR*, 2021.
- [17] Y. Long, X. Li, W. Cai, and H. Dong, "Discuss before moving: Visual language navigation via multi-expert discussions," in *ICRA*, 2024.
- [18] S. Raychaudhuri, S. Wani, S. Patel, U. Jain, and A. X. Chang, "Language-aligned waypoint (law) supervision for vision-and-language navigation in continuous environments," in *EMNLP*, 2021.
- [19] H. Wang, W. Liang, L. Van Gool, and W. Wang, "Dreamwalker: Mental planning for continuous vision-language navigation," in *CVPR*, 2023.
- [20] M. Z. Irshad, C.-Y. Ma, and Z. Kira, "Hierarchical cross-modal agent for robotics vision-and-language navigation," in *ICRA*, 2021.
- [21] D. Shah, B. Osinski, B. Ichter, and S. Levine, "LM-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *CoRL*, 2022.
- [22] Y. Long, W. Cai, H. Wang, G. Zhan, and H. Dong, "Instructnav: Zero-shot system for generic instruction navigation in unexplored environment," in *CoRL*, 2024.
- [23] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," in *ICRA*, 2023.
- [24] G. Zhou, Y. Hong, Z. Wang, X. E. Wang, and Q. Wu, "Navgpt-2: Unleashing navigational reasoning capability for large vision-language models," in *ECCV*, 2024.
- [25] N. Yokoyama, S. Ha, D. Batra, J. Wang, and B. Bucher, "Vlfm: Vision-language frontier maps for zero-shot semantic navigation," in *ICRA*, 2024.
- [26] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [27] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan, *et al.*, "The llama 3 herd of models," *arXiv preprint arXiv:2407.21783*, 2024.
- [28] Q. Team, "Qwen2 technical report," *arXiv preprint arXiv:2407.10671*, 2024.
- [29] Y. Zhang, J. Wu, W. Li, B. Li, Z. Ma, Z. Liu, and C. Li, "Video instruction tuning with synthetic data," *arXiv preprint arXiv:2410.02713*, 2024.
- [30] Z. Liu, L. Zhu, B. Shi, Z. Zhang, Y. Lou, S. Yang, H. Xi, S. Cao, Y. Gu, D. Li, *et al.*, "Nvila: Efficient frontier visual language models," in *CVPR*, 2025.
- [31] Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. Li, Y. Wu, *et al.*, "Deepseekmath: Pushing the limits of mathematical reasoning in open language models," *arXiv preprint arXiv:2402.03300*, 2024.
- [32] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi, *et al.*, "Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning," *arXiv preprint arXiv:2501.12948*, 2025.
- [33] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [34] Z. Zheng, M. Yang, J. Hong, C. Zhao, G. Xu, L. Yang, C. Shen, and X. Yu, "Deepeyes: Incentivizing thinking with images" via reinforcement learning," *arXiv preprint arXiv:2505.14362*, 2025.
- [35] Z. Wang, K. Wang, Q. Wang, P. Zhang, L. Li, Z. Yang, X. Jin, K. Yu, M. N. Nguyen, L. Liu, *et al.*, "Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning," *arXiv preprint arXiv:2504.20073*, 2025.
- [36] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," in *3DV*, 2017.
- [37] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, *et al.*, "Habitat: A platform for embodied ai research," in *CVPR*, 2019.
- [38] G. Ilharco, V. Jain, A. Ku, E. Ie, and J. Baldridge, "General evaluation for instruction conditioned navigation using dynamic time warping," in *NeurIPS*, 2019.
- [39] H. Tan, L. Yu, and M. Bansal, "Learning to navigate unseen environments: Back translation with environmental dropout," in *NAACL*, 2019.
- [40] S. Bai, K. Chen, X. Liu, J. Wang, W. Ge, S. Song, K. Dang, P. Wang, S. Wang, J. Tang, *et al.*, "Qwen2. 5-vl technical report," *arXiv preprint arXiv:2502.13923*, 2025.
- [41] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [42] Q. Yu, Z. Zhang, R. Zhu, Y. Yuan, X. Zuo, Y. Yue, W. Dai, T. Fan, G. Liu, L. Liu, *et al.*, "Dapo: An open-source llm reinforcement learning system at scale," *arXiv preprint arXiv:2503.14476*, 2025.
- [43] L. Zhang, X. Hao, Q. Xu, Q. Zhang, X. Zhang, P. Wang, J. Zhang, Z. Wang, S. Zhang, and R. M. Xu, "A novel memory representation via annotated semantic maps for vlm-based vision-and-language navigation," *arXiv preprint arXiv:2502.13451*, 2025.
- [44] Y. Zhang, J. Wu, W. Li, B. Li, Z. MA, Z. Liu, and C. Li, "LLaVA-video: Video instruction tuning with synthetic data," *TMLR*, 2025.
- [45] D. Azuma, T. Miyaniishi, S. Kurita, and M. Kawanabe, "Scanqa: 3d question answering for spatial scene understanding," in *CVPR*, 2022.
- [46] K. Lin, P. Chen, D. Huang, T. H. Li, M. Tan, and C. Gan, "Learning vision-and-language navigation from youtube videos," in *ICCV*, 2023.