LongNav-R1 Turn Count Effects on GPU Memory and Inference Latency in RxR-CE

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: What is the impact of varying the number of turns in LongNav-R1 on GPU memory consumption and inference latency during long-horizon task execution on RxR-CE. This paper develops LongNav-R1, an end-to-end multi-turn reinforcement learning (RL) framework designed to optimize Visual-Language-Action (VLA) models for long-horizon navigation. Unlike existing single-turn paradigm, LongNav-R1 reformulates the navigation decision process as a. 12 claims were extracted from source literature; 3 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: LongNav-R1: Horizon-Adaptive Multi-Turn RL for Long-Horizon VLA Navigation. Research question: What is the impact of varying the number of turns in LongNav-R1 on GPU memory consumption and inference latency during long-horizon task execution on RxR-CE?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.7/10.

3 Results

11 papers retrieved. 12 claims extracted; 3 independently verified. Quality review score: 4.7/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
LongNav-R1 is an end-to-end framework that reformulates navigation as a multi-turn Reinforcement Learning (RL) process.	✓	0.27
Existing state-of-the-art navigation methods adopt a single-turn imitation learning paradigm.	×	0.11
Single-turn imitation learning treats navigation steps independently, overlooking sequential dependencies.	×	0.05
Single-turn imitation learning leads to behavioral rigidity by strictly imitating expert trajectories instead of optimiz	×	0.07
LongNav-R1 treats the navigation task as a continuous conversation between the VLA policy and the physical environment.	×	0.14
LongNav-R1 allows large VLA models to improve multi-step decision-making without the significant computational burden of	×	0.09
LongNav-R1 significantly outperforms existing methods in real-world and diverse navigation benchmarks.	✓	0.15
LongNav-R1 demonstrates zero-shot performance in long-horizon real-world navigation settings.	✓	0.22
All source code for LongNav-R1 will be open-sourced upon publication.	×	0.05
Early semantic navigation methods largely focused on acquiring task-specific skills via imitation learning or RL.	×	0.04
Early semantic navigation methods often suffer from poor generalization due to domain gaps.	×	0.03
Recent approaches leveraging LLMs and VLMs for navigation offer greater flexibility but often lack optimized task execut	×	0.04

References

- http://arxiv.org/abs/2412.09082v3
- http://arxiv.org/abs/2604.14552v2
- http://arxiv.org/abs/2602.12351v1