Habitat-Sim vs. Isaac Sim and Omniverse on GibsonEnv Benchmark Performance

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: How does the performance of Habitat-Sim compare to other 3D simulators like Isaac Sim or Omniverse when evaluated on the GibsonEnv benchmark in terms of FPS and agent training convergence. 12 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Causal-Paced Deep Reinforcement Learning. Research question: How does the performance of Habitat-Sim compare to other 3D simulators like Isaac Sim or Omniverse when evaluated on the GibsonEnv benchmark in terms of FPS and agent training convergence?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.5/10.

3 Results

10 papers retrieved. 12 claims extracted; 0 independently verified. Quality review score: 3.5/10.

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
CP-DRL was evaluated on two benchmark environments: Point Mass (PM) and BipedalWalker (BW).	×	0.09
For the PM environment, Proximal Policy Optimization (PPO) was used as the student algorithm.	×	0.03
For BW, Soft Actor-Critic (SAC) was employed as in Klink et al. (2022).	×	0.01
Experiments were conducted with multiple random seeds: 10 seeds for PM, 5 seeds for BW trivial tasks, and 3 seeds for BW	×	0.04
CP-DRL demonstrated superior performance in the Point Mass environment, reaching a return of 6.17 $\pm$ 0.08 at epoch 195.	×	0.07
CP-DRL outperformed the second-best method, CURROT (5.6 $\pm$ 0.34), with an improvement of approximately 10.2%.	×	0.05
CP-DRL maintained consistently low standard errors during training, indicating more stable learning compared to CURROT.	×	0.06
The PM environment requires controlling a point agent to navigate through a narrow gate to reach a target location on th	×	0.04
The task difficulty in the PM environment is modulated by the gate's width and position.	×	0.02
The target context distribution μ(c) in the PM environment is bimodal, corresponding to two gate positions on opposite s	×	0.02
Each training epoch in the PM environment consists of 4,096 rollouts, and all methods are trained for 200 epochs to ensu	×	0.02
CP-DRL quickly converges to tasks aligned with the target gate position and progressively focuses on narrower gate width	×	0.06

References

- http://arxiv.org/abs/2507.02910v1
- http://arxiv.org/abs/2604.12626v1

- http://arxiv.org/abs/2606.03551v1