

# Photorealistic Rendering Quality Effects on Agent Training Convergence in RoboStack

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: What is the impact of photorealistic rendering quality in Habitat-Sim on agent training convergence in the RoboStack benchmark compared to stylized or low-fidelity simulations. 13 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.6/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Habitat-GS: A High-Fidelity Navigation Simulator with Dynamic Gaussian Splatting. Research question: What is the impact of photorealistic rendering quality in Habitat-Sim on agent training convergence in the RoboStack benchmark compared to stylized or low-fidelity simulations?.

## 2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.6/10.

## 3 Results

13 papers retrieved. 13 claims extracted; 2 independently verified. Quality review score: 4.6/10.

## 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
Habitat-GS combines the InteriorGS dataset with additional real-world reconstructed GS scenes in a 4:1 ratio, yielding 1	×	0.08
For mesh scenes, Habitat-GS uses the Habitat-Matterport 3D (HM3D) dataset, similarly split into 100 training and 20 test	×	0.08
The GS and mesh test sets are drawn from disjoint scene collections rather than two representations of the same physical	×	0.05
For gaussian avatars, Habitat-GS exports canonical gaussians from six trained AnimatableGaussians identities, with three	×	0.11
Each avatar is driven by GAMMA-generated motion trajectories with pre-computed joint matrices and proxy capsules.	×	0.02
Evaluation metrics include Success Rate (SR), Success weighted by Path Length (SPL), Distance to Goal (DTG), Collision R	×	0.05
Gemini 3.0 Pro is used as an automated evaluator to score images rendered from an embodied AI simulator on dimensions su	×	0.09
Training on 3DGS scenes improves agent visual robustness and cross-domain generalization on PointNav.	✓	0.15
Training with dynamic gaussian avatars equips agents with human-aware navigation capabilities on dynamic point-goal navi	✓	0.26
The system remains efficient under varying scene complexity and avatar counts.	×	0.15
The simulator’s ResourceManager is the main entry point and controller, while the Asset Importing Module manages asset i	×	0.05
The Navigation Module includes a PathFinder and NavMesh, while the Gaussian Data Structures include GaussianSplattingDat	×	0.05
The Rendering Module includes GaussianSplattingDrawable and GaussianAvatarDrawable, with backend engine support in C++.	×	0.03

## References

- <http://arxiv.org/abs/2507.02910v1>
- <http://arxiv.org/abs/2602.12375v1>
- <http://arxiv.org/abs/2604.12626v1>