

DeepSeek R1 and Tulu 3 Pass@1 Accuracy on HumanEval Under Single-GPU Constraints

Assignee Research

June 4, 2026

Abstract

This report synthesises findings from 7 peer-reviewed papers addressing the following research question: What is the difference in pass@1 accuracy between Deepseek R1 and Tulu 3 on the HumanEval benchmark when restricted to single consumer-grade GPU inference. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Quantitative Analysis of Performance Drop in DeepSeek Model Quantization. Research question: What is the difference in pass@1 accuracy between Deepseek R1 and Tulu 3 on the HumanEval benchmark when restricted to single consumer-grade GPU inference?.

2 Methodology

Systematic literature search across multiple databases yielded 7 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.8/10.

3 Results

7 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 5.8/10.

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.

References

- <http://arxiv.org/abs/2409.15080v1>
- <http://arxiv.org/abs/2505.02390v2>
- <http://arxiv.org/abs/2503.10486v2>