

Fine-Tuned Phi-3-Mini and Mistral-7B Hallucination Rates in Long-Context RAG Benchmarks

Assignee Research

June 6, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: How do fine-tuned versions of Phi-3-mini and Mistral-7B on religious domain data compare in hallucination rates versus their base versions on long-context RAG benchmarks. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: MUST-RAG: MUSical Text Question Answering with Retrieval Augmented Generation. Research question: How do fine-tuned versions of Phi-3-mini and Mistral-7B on religious domain data compare in hallucination rates versus their base versions on long-context RAG benchmarks?.

2 Methodology

Systematic literature search across multiple databases yielded 16 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

16 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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.

References

- <http://arxiv.org/abs/2506.22486v1>
- <http://arxiv.org/abs/2507.23334v2>
- <http://arxiv.org/abs/2511.07328v2>