

Faithfulness Constraints Reduce Throughput and Accuracy in Phi-3-Mini and Mistral-7B on Low-Resource Tasks

Assignee Research

June 6, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: What is the comparative performance degradation in throughput and answer accuracy for Phi-3-mini and Mistral-7B-v0.1 when enforcing strict faithfulness constraints during retrieval-augmented. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: SLM-Bench: A Comprehensive Benchmark of Small Language Models on Environmental Impacts—Extended Version. Research question: What is the comparative performance degradation in throughput and answer accuracy for Phi-3-mini and Mistral-7B-v0.1 when enforcing strict faithfulness constraints during retrieval-augmented generation on low-resource tasks?.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.3/10.

3 Results

12 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.3/10.

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/2603.21389v1>
- <http://arxiv.org/abs/2508.15478v2>
- <http://arxiv.org/abs/2510.25621v1>