

Retrieval-Augmented Generation Effects on Alignment and Faithfulness in 7B vs 70B Models for Domain-Specific QA

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: How does the integration of retrieval-augmented generation affect the alignment and faithfulness of 7B models compared to 70B models in domain-specific QA tasks (e.g., legal, religious) when measured. 9 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Investigating Retrieval-Augmented Generation in Quranic Studies: A Study of 13 Open-Source Large Language Models. Research question: How does the integration of retrieval-augmented generation affect the alignment and faithfulness of 7B models compared to 70B models in domain-specific QA tasks (e.g., legal, religious) when measured by metrics like FAITH and CONTRADICTION?.

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 5.2/10.

3 Results

10 papers retrieved. 9 claims extracted; 2 independently verified. Quality review score: 5.2/10.

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
The study investigates 13 open-source Large Language Models in the context of Quranic studies.	✓	0.24
The system employs a Retrieval-Augmented Generation (RAG) architecture combining retrieval-based and generative methods.	×	0.09
Semantic search is performed over a vectorized dataset obtained from Qur'anic surah descriptions.	×	0.02
Generated responses include references to original dataset entries, such as surah descriptions or specific virtues.	×	0.06
Human evaluators assessed response quality based on three dimensions: Context Relevance, Answer Faithfulness, and Answer	✓	0.16
Context Relevance is calculated using the precision@k metric.	×	0.07
The dataset selection criteria included Authenticity, Descriptive Richness, Clarity and Accessibility, and Relevance.	×	0.03
The dataset source underwent a thorough review to confirm compliance with recognized Islamic scholarship and the absence	×	0.02
Evaluation data was logged and stored on a platform after responses were evaluated, scored, and commented on.	×	0.03

References

- <http://arxiv.org/abs/2509.12382v1>
- <http://arxiv.org/abs/2503.16581v1>

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