Retrieval-Augmented Generation Efficiency and Factual Accuracy Across Religious and Scientific Texts

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: What is the correlation between context window utilization efficiency and factual accuracy in retrieval-augmented generation systems when processing long-form religious texts compared to structured. 18 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.4/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: On the Consistency of Multilingual Context Utilization in Retrieval-Augmented Generation. Research question: What is the correlation between context window utilization efficiency and factual accuracy in retrieval-augmented generation systems when processing long-form religious texts compared to structured scientific abstracts?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.4/10.

3 Results

14 papers retrieved. 18 claims extracted; 1 independently verified. Quality review score: 4.4/10.

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
All codes and data for the study are released at https://github.com/Betswish/mRAG-Context-Consistency.	×	0.03
In multilingual RAG (mRAG), retrieval can be performed either monolingually or cross-lingually.	×	0.08
Monolingual retrieval is performed only over passages in the same language as the query.	×	0.11
Monolingual retrieval can be successful for high-resource languages.	×	0.05
Monolingual retrieval is marginally useful or even harmful when the question is posed in a low-resource language.	×	0.05
Relevant information for low-resource language queries is likely to be available only in different languages.	×	0.13
Cross-lingual retrieval attempts to extract useful information simultaneously from multiple languages.	×	0.05
Cross-lingual retrieval leads to visible gains in low-resource languages.	×	0.06
Recent mRAG evaluations by Chirkova et al. (2024) and Park and Lee (2025) test the LLM pipeline (retrieval + generation)	×	0.06
The study evaluates how LLMs leverage retrieved passages in different languages in various multilingual QA tasks.	✓	0.21
The study reveals that LLMs have remarkably robust input understanding but much more brittle generation abilities in mul	×	0.06
The study applies a RAG answer attribution method based on feature attribution analysis.	×	0.09
The study considers both single-passage and multi-passage mRAG setups.	×	0.05
Chirkova et al. (2024) investigated mRAG systems across 13 languages.	×	0.07
Chirkova et al. (2024) highlighted the limited gains of in-language retrieval in their setup.	×	0.07
Nie et al. (2023) proposed the Prompts Augmented by Retrieval Crosslingually (PARC) pipeline.	×	0.08
The PARC pipeline augments contexts with semantically similar sentences retrieved from high-resource languages to enhanc	×	0.07
Gao et al. (2022) introduced a retrieval-augmented method for multilingual keyphrase generation leveraging English annot	×	0.06

References

- http://arxiv.org/abs/2504.00597v4
- http://arxiv.org/abs/2204.07288v1
- http://arxiv.org/abs/2504.19754v1