

Retrieval Algorithm Impact on Hallucination Rates in RAG Systems for Religious Texts

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 7 peer-reviewed papers addressing the following research question: How does the choice of retrieval algorithm (BM25 vs. dense vs. hybrid) impact the hallucination rate in RAG systems evaluated on religious text corpora when using the same compression ratio. 9 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Hybrid Retrieval for Hallucination Mitigation in Large Language Models: A Comparative Analysis. Research question: How does the choice of retrieval algorithm (BM25 vs. dense vs. hybrid) impact the hallucination rate in RAG systems evaluated on religious text corpora when using the same compression ratio?.

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 3.8/10.

3 Results

7 papers retrieved. 9 claims extracted; 0 independently verified. Quality review score: 3.8/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
The hybrid retrieval system assigns a weight of 0.7 to Sparse and 0.3 to Dense for specific queries.	×	0.07
The hybrid retrieval system assigns a weight of 0.7 to Dense and 0.3 to Sparse for generic queries.	×	0.06
The hybrid retrieval pipeline includes query expansion using WordNet.	×	0.09
The hybrid retrieval pipeline uses Reciprocal Rank Fusion for combining results.	×	0.10
The retrieval phase of a LLM-driven RAG system contains two main components: the indexed database and the retriever.	×	0.06
The indexed database DB is an external knowledge-base which is a structured collection of documents $d_i \in D$, for $i = \{1, \dots\}$.	×	0.05
The retriever encodes q and all D in a vector space and applies a chosen similarity function f_{sim} to compute a similarity	×	0.03
The hybrid retrieval system uses dynamic weighting for combining results from sparse and dense retrievers.	×	0.11
The hybrid retrieval system uses top-k results from sparse and dense retrievers for generic queries.	×	0.12

References

- <http://arxiv.org/abs/2404.07220v2>
- <http://arxiv.org/abs/2504.05324v1>
- <http://arxiv.org/abs/2604.01733v1>