

Context Window Compression and Hallucination Rates in Domain-Specific RAG Systems for Religious Texts

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: What is the correlation between context window compression ratios and hallucination rates in domain-specific RAG systems evaluated on religious text corpora. 19 claims were extracted from source literature; 6 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: RAGalyst: Automated Human-Aligned Agentic Evaluation for Domain-Specific RAG. Research question: What is the correlation between context window compression ratios and hallucination rates in domain-specific RAG systems evaluated on religious text corpora?.

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 5.8/10.

3 Results

16 papers retrieved. 19 claims extracted; 6 independently verified. Quality review score: 5.8/10.

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
RAGalyst is an automated, human-aligned agentic framework designed for the evaluation of domain-specific RAG systems.	✓	0.40
RAGalyst features an agentic pipeline that generates synthetic question-answering (QA) datasets from source documents.	✓	0.25
RAGalyst incorporates an agentic filtering step to ensure data fidelity in generated datasets.	✓	0.17
RAGAS is adopted as the primary baseline for evaluating the performance of RAGalyst because it is the only actively main	×	0.05
Existing end-to-end RAG evaluation frameworks rely on some degree of manual validation for QA dataset generation.	×	0.09
Rule-based metrics often fail to capture subtle semantic nuances in RAG evaluation.	×	0.13
LLM-based metrics in existing works are rarely benchmarked for alignment with human judgment.	×	0.13
AutoCalibrate auto-tunes evaluation prompts to improve alignment with human judgment.	×	0.08
DSPy is a declarative framework that enables the programmatic creation and refinement of prompts for LLMs.	×	0.03
The PoLL evaluation framework advocates using a panel of various LLM evaluators to reduce bias and variance in generatio	×	0.06
High-quality QA datasets are often unavailable or insufficient in specialized domains.	×	0.08
Alberti et al. pioneered a round-trip consistency approach combining answer extraction, question generation, and answer	×	0.05
Shakeri et al. proposed an end-to-end transformer-based generator that outputs both the question and the answer from a g	×	0.08
RAGalyst uses prompt optimization to achieve a strong correlation with human annotations for Correctness and Answerabili	✓	0.22
The RAGalyst framework was applied to evaluate RAG components across three domains: military operations, cybersecurity,	✓	0.19
No single embedding model, LLM, or hyperparameter configuration proved univ ersally ally optimal across the tested domains.	✓	0.18
RAGalyst is available on Github.	×	0.09
Modern Large Language Models suffer from hallucinations, defined as generating content that appears plausible but is fac	×	0.06
Mitigating hallucinations is especially important in safety-critical applications such as military	×	0.13

References

- <http://arxiv.org/abs/2506.09886v2>
- <http://arxiv.org/abs/2511.04502v1>
- <http://arxiv.org/abs/2507.22931v3>