Diverse Retriever Portfolios Reduce Hallucinations in Open-Domain RAG Systems

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: Can diverse retriever portfolios reduce hallucination rates in RAG systems on open-domain QA benchmarks by optimizing for query-type coverage rather than aggregate retrieval accuracy. 13 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Retriever Portfolios: A Principled Approach to Adaptive RAG. Research question: Can diverse retriever portfolios reduce hallucination rates in RAG systems on open-domain QA benchmarks by optimizing for query-type coverage rather than aggregate retrieval accuracy?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

4 papers retrieved. 13 claims extracted; 1 independently verified. Quality review score: 4.5/10.

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 evaluates retriever portfolios on four QA benchmarks: HotpotQA, 2WikiMultiHopQA, TriviaQA, and MusiQue.	×	0.09
The evaluation uses two answer models: Gemma-3-27B-It and Llama-3.1-70B-Instruct.	×	0.01
A size-k portfolio is evaluated by its best-of-k retrieval score, defined as the maximum support-document score achieved	×	0.05
The full heterogeneous candidate pool consists of 360 candidates, including DS and Vendi retrievers with MPNet and E5 ba	×	0.05
The portfolio is trained once on pooled training queries from all four benchmarks and evaluated on their corresponding t	×	0.04
At portfolio size k=5, the top-k average baseline achieves 0.492 support recall and 0.432 support F1.	×	0.02
At portfolio size k=5, the learned portfolio achieves 0.594 support recall and 0.500 support F1.	×	0.03
The learned portfolio adds lower-average but complementary Vendi and GraphDense variants to cover queries missed by earl	×	0.05
The top-k average baseline list is dominated by closely related GraphDense/E5 configurations.	×	0.00
The proposed method yields better retrieval recall and answer accuracy compared to single-retriever baselines.	✓	0.16
The proposed method yields better retrieval recall and answer accuracy compared to inference-time tuning methods like Ve	×	0.12
The proposed method significantly reduces latency and token usage compared to baselines.	×	0.07
Retrieval-augmented generation (RAG) conditions generation on both the query and the retrieved context to improve factua	×	0.10

References

- http://arxiv.org/abs/2605.31176v1

- http://arxiv.org/abs/2210.02627v1
- http://arxiv.org/abs/2510.22344v1