

Impact of Retrieval Step Count on Vendi-RAG Performance in Multi-Hop QA Benchmarks

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: What is the impact of varying the number of retrieval steps on the downstream task performance of Vendi-RAG when evaluated on the HEQA benchmark versus HotPotQA. 15 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: Vendi-RAG: Adaptively Trading-Off Diversity And Quality Significantly Improves Retrieval Augmented Generation With LLMs. Research question: What is the impact of varying the number of retrieval steps on the downstream task performance of Vendi-RAG when evaluated on the HEQA benchmark versus HotPotQA?.

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 4.4/10.

3 Results

10 papers retrieved. 15 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
Experiments were conducted on three multi-hop QA benchmark datasets: MuSiQue, HotpotQA, and 2WikiMultiHopQA.	✓	0.15
The sensitivity analysis of the VSR process used 100 randomly sampled queries from the dataset.	×	0.03
Setting the parameter $s = 0.0$ serves as a baseline representing a pure similarity search scenario relying exclusively on	×	0.02
In the sensitivity analysis, increasing the parameter s from 0.0 to 1.0 causes both Kendall's τ and Spearman's ρ to decrease	×	0.01
At parameter $s = 0.2$, the Kendall's τ value is 0.797 and the Spearman's ρ value is 0.828.	×	0.00
At parameter $s = 1.0$, the Kendall's τ value is 0.074 and the Spearman's ρ value is 0.078.	×	0.00
On the 2WikiMultiHopQA dataset, Vendi-RAG-4o achieved an F1-score of 69.9%.	×	0.10
On the 2WikiMultiHopQA dataset, Adaptive-RAG-4o achieved an F1-score of 63.4%.	×	0.07
On the HotpotQA dataset, Vendi-RAG-4o achieved an Exact Match score of 56.5%.	×	0.09
On the HotpotQA dataset, Adaptive-RAG-4o achieved an Exact Match score of 52.1%.	×	0.07
On the MuSiQue dataset, Vendi-RAG-4o achieved an Accuracy of 63.4%.	×	0.12
On the MuSiQue dataset, Adaptive-RAG-4o achieved an Accuracy of 62.8%.	×	0.08
The Vendi Score (VS) explicitly quantifies semantic diversity in a set of documents.	×	0.13
The Vendi Score attains its maximum value n when all documents in the set are orthogonal.	×	0.06
Similarity search (SS) often results in redundant documents with high similarity.	×	0.03

References

- <http://arxiv.org/abs/2502.11228v2>

- <http://arxiv.org/abs/2404.14464v1>
- <http://arxiv.org/abs/2407.20114v3>