

Adaptive Diversity-Quality Trade-offs in Vendi-RAG for Robust Multi-Hop Question Answering

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: What is the effect of adaptive diversity-quality trade-offs in Vendi-RAG on answer robustness when evaluated against adversarial distractors in the HotpotQA benchmark. 8 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.2/10. 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 effect of adaptive diversity-quality trade-offs in Vendi-RAG on answer robustness when evaluated against adversarial distractors in the HotpotQA benchmark?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.2/10.

3 Results

13 papers retrieved. 8 claims extracted; 2 independently verified. Quality review score: 5.2/10.

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
Vendi-RAG was evaluated on three multi-hop QA benchmark datasets: MuSiQue, HotpotQA, and 2WikiMultiHopQA.	✓	0.23
The sensitivity analysis of the VSR process was conducted using 100 randomly sampled queries from the dataset.	×	0.03
The sensitivity analysis evaluated the retrieval pipeline across multiple s values ranging from 0.0 to 1.0.	×	0.03
Setting $s = 0.0$ serves as a baseline representing a pure similarity search scenario.	×	0.02
Kendall’s τ and Spearman’s ρ were used to quantify deviations from the baseline in the sensitivity analysis.	×	0.03
As s increases from 0.0 to 1.0, both Kendall’s τ and Spearman’s ρ decrease progressively.	×	0.04
Vendi-RAG uses a retrieval approach based on the Vendi Score (VS), which explicitly quantifies semantic diversity in a s	✓	0.19
The Vendi Score (VSk(D)) reflects the effective number of unique documents in D , attaining its maximum value n when all	×	0.06

References

- <http://arxiv.org/abs/2509.12382v1>
- <http://arxiv.org/abs/2502.11228v2>
- <http://arxiv.org/abs/2404.14464v1>