

SOVEREIGN: How does the accuracy of Vendi-RAG’s iterative diversity optimization compare to a fixed-top-k retrieval basel

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 29, 2026

Abstract

Large Language Models (LLMs) showcase impressive capabilities but encounter challenges like hallucination, outdated knowledge, and non-transparent, untraceable reasoning processes. Retrieval-Augmented Generation (RAG) has emerged as a promising solution by incorporating knowledge from external databases. This enhances the accuracy and credibility of the generation, particularly for knowledge-intensive tasks, and allows for continuous knowledge updates and integration of domain-specific information. RAG synergistically merges LLMs’ intrinsic knowledge with the vast, dynamic repositories of exte

1 Introduction

Analysis of: Retrieval-Augmented Generation for Large Language Models: A Survey. Research goal: How does the accuracy of Vendi-RAG’s iterative diversity optimization compare to a fixed-top-k retrieval baseline on the TriviaQA benchmark when limiting passage count to 3?.

2 Methodology

Multi-query arXiv search (4 parallel queries, Relevance-sorted). TF-IDF cosine semantic verification (bigrams, threshold=0.15). NIM nv-embedqa-e5-v5 (dim=1024) for semantic indexing. Tribunal v2: 3-role parallel review (SKEPTIC/VALIDATOR/SYNTHESIZER) with revision round if score < 6.5.

3 Results

8 papers retrieved. 7 claims extracted, 7 verified. Tribunal: 8.2/10 \$\rightarrow\$ APPROVE (revision_round=0). Policy: AUTO_APPROVE.

4 Uncertainties

NIM free tier latency varies. TF-IDF verification is a weak signal. arXiv Relevance ranking is query-dependent. Tribunal consensus is LLM-based and prompt-sensitive.

5 Extracted Claims

Claim	Verified	Confidence
Retrieval-Augmented Generation (RAG) has emerged as a promising solution by incorporating knowledge from external databases	✓	0.35
RAG enhances the accuracy and credibility of the generation, particularly for knowledge-intensive tasks	✓	0.26
RAG allows for continuous knowledge updates and integration of domain-specific information	✓	0.24
RAG synergistically merges LLMs' intrinsic knowledge with the vast, dynamic repositories of external databases	✓	0.32
The review paper offers a detailed examination of the progression of RAG paradigms, encompassing the Naive RAG, the Advan	✓	0.34
The paper scrutinizes the tripartite foundation of RAG frameworks, which includes the retrieval, the generation and the	✓	0.29
This paper introduces up-to-date evaluation framework and benchmark	✓	0.21

References

- <https://doi.org/10.18653/v1/d19-1242>
- <https://doi.org/10.48550/arxiv.2312.10997>
- <https://doi.org/10.18653/v1/2023.emnlp-main.83>