Hybrid Retrieval Integration in Vendi-RAG: ROUGE-L Performance on ELI5

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: What is the impact of combining sparse and dense retrieval methods (hybrid retrieval) on the ROUGE-L performance of Vendi-RAG on the ELI5 dataset compared to using each method individually. Large Language Models (LLMs) excel in language comprehension and generation but are prone to hallucinations, producing factually incorrect or unsupported outputs. Retrieval Augmented Generation (RAG) systems address this issue by grounding LLM responses with external knowledge. 12 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: Hybrid Retrieval for Hallucination Mitigation in Large Language Models: A Comparative Analysis. Research question: What is the impact of combining sparse and dense retrieval methods (hybrid retrieval) on the ROUGE-L performance of Vendi-RAG on the ELI5 dataset compared to using each method individually?.

2 Methodology

Systematic literature search across multiple databases yielded 15 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

15 papers retrieved. 12 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 hybrid retrieval system assigns a weight of 0.7 to Sparse and 0.3 to Dense for specific queries.	×	0.09
The hybrid retrieval system assigns a weight of 0.7 to Dense and 0.3 to Sparse for generic queries.	×	0.09
The hybrid retrieval system uses WordNet for query expansion.	×	0.09
The hybrid retrieval system uses Reciprocal Rank Fusion for combining results.	×	0.13
The hybrid retrieval system uses a dynamic weighting mechanism for combining results.	×	0.05
The hybrid retrieval system uses a Sparse Retriever and a Dense Retriever.	×	0.12
The hybrid retrieval system combines the top-k results of the Sparse Retriever and the Dense Retriever.	✓	0.17
The hybrid retrieval system is designed to address lexical chasm, i.e., the gap or mismatch between the vocabulary used	×	0.05
The retrieval phase of a LLM-driven RAG system contains two main components: the indexed database and the retriever.	×	0.05
The indexed database is an external knowledge-base which is a structured collection of documents.	×	0.04
The retriever encodes the query and all documents in a vector space.	×	0.05
The retriever applies a chosen similarity function to compute a similarity score between the vector representations of t	×	0.04

References

- http://arxiv.org/abs/2502.11228v2
- http://arxiv.org/abs/2504.05324v1
- http://arxiv.org/abs/2109.10739v1