

SOVEREIGN: What is the impact of query-knowledge base co-evolution on retrieval latency and generation quality as measure

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Large Language Models (LLMs) showcase impressive capabilities but encounter challenges like hallucination, outdated knowledge, and non-transparent, untraceable reasoning processes.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: What is the impact of query-knowledge base co-evolution on retrieval latency and generation quality as measured by pass@1 and pass@10 metrics?.

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

5 papers retrieved. 8 claims extracted, 7 verified. Tribunal: 8.2/10 → 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) enhances the accuracy and credibility of generation for knowledge-intensive tasks	✓	0.30
LLMs encounter challenges like hallucination, outdated knowledge, and non-transparent, untraceable reasoning processes	✓	0.29
RAG has emerged as a promising solution by incorporating knowledge from external databases	✓	0.27
RAG allows for continuous knowledge updates and integration of domain-specific information	✓	0.24
RAG synergistically merges LLMs' intrinsic knowledge with external databases	✓	0.28
The paper offers a detailed examination of Naive RAG, Advanced RAG, and Modular RAG paradigms	✓	0.27
The paper scrutinizes the tripartite foundation of RAG frameworks including retrieval, generation, and augmentation tech	✓	0.24
This paper introduces up-to-date evaluation frameworks and benchmarks	×	0.14

References

- <https://doi.org/10.1109/access.2023.3295776>
- <https://doi.org/10.65563/jeaai.v1i7.65>
- <https://doi.org/10.48550/arxiv.2312.10997>