

SOVEREIGN: What is the inference latency overhead of EVOR’s synchronous query-knowledge base evolution compared to tradit

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Large Language Models (LLMs) showcase impressive capabilities but encounter challenges like hallucination, outdated knowledge, and non-transparent, untraceable reasoning processes.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 inference latency overhead of EVOR’s synchronous query-knowledge base evolution compared to traditional RAG approaches on code generation tasks?.

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

10 papers retrieved. 10 claims extracted, 10 verified. Tribunal: 7.5/10 → 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
Large Language Models (LLMs) encounter challenges like hallucination, outdated knowledge, and non-transparent, untraceable	✓	0.34
Retrieval-Augmented Generation (RAG) has emerged as a promising solution by incorporating knowledge from external databases	✓	0.39
RAG enhances the accuracy and credibility of the generation, particularly for knowledge-intensive tasks.	✓	0.27
RAG allows for continuous knowledge updates and integration of domain-specific information.	✓	0.27
RAG synergistically merges LLMs' intrinsic knowledge with the vast, dynamic repositories of external databases.	✓	0.34
This comprehensive review paper offers a detailed examination of the progression of RAG paradigms, encompassing the Naive	✓	0.36
The paper meticulously scrutinizes the tripartite foundation of RAG frameworks, which includes the retrieval, the genera	✓	0.32
The paper highlights the state-of-the-art technologies embedded in each of these critical components.	✓	0.24
This paper introduces up-to-date evaluation framework and benchmark.	✓	0.21
This article delineates the challenges currently faced and points out prospective avenues for research and development.	✓	0.27

References

- <https://doi.org/10.1007/s11704-026-60308-3>
- <https://doi.org/10.1109/access.2014.2332453>

- <https://doi.org/10.48550/arxiv.2312.10997>