SOVEREIGN: What is the impact of domain-specific data augmentation on cross-modal alignment scores in multimodal LLMs eva

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Large Language Models (LLMs) showcase impressive capabilities but encounter challenges like hallucination, outdated knowledge, and non-transparent, untraceable reasoning processes.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 domain-specific data augmentation on cross-modal alignment scores in multimodal LLMs evaluated on standard vision-language understanding datasets?.

2 Methodology

Multi-query arXiv search (1 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

3 papers retrieved. 8 claims extracted, 8 verified. Tribunal: 8.5/10 $\rightarrow$ 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
Large Language Models (LLMs) encounter challenges such as hallucination, outdated knowledge, and non-transparent, untrac	✓	0.31
Retrieval-Augmented Generation (RAG) incorporates knowledge from external databases to address LLM challenges.	✓	0.24
RAG enhances the accuracy and credibility of generation, particularly for knowledge-intensive tasks.	✓	0.24
RAG allows for continuous knowledge updates and the integration of domain-specific information.	✓	0.24
RAG paradigms are categorized into Naive RAG, Advanced RAG, and Modular RAG.	✓	0.19
The tripartite foundation of RAG frameworks consists of retrieval, generation, and augmentation techniques.	✓	0.22
The paper introduces an up-to-date evaluation framework and benchmark for RAG systems.	✓	0.23
The paper delineates current challenges and points out prospective avenues for research and development in RAG.	✓	0.18

References

- https://doi.org/10.1145/3649506
- https://doi.org/10.48550/arxiv.2303.10130
- https://doi.org/10.48550/arxiv.2312.10997