

An Agentic RAG Architecture for Intelligent Knowledge Retrieval Systems

Michael Hartono, Sarah Wijaya

International Journal of Research and Innovation

Abstract

This demo study presents an agentic retrieval-augmented generation architecture for institutional knowledge retrieval. The proposed design separates query planning, evidence retrieval, source evaluation, response synthesis, and citation checking into observable components coordinated by a constrained agent loop. A synthetic document collection and fictional user tasks are used to assess retrieval coverage, answer grounding, latency, and failure recovery. The simulated evaluation suggests that explicit planning and evidence verification can improve traceability when compared with a single-pass retrieval pipeline, although additional orchestration increases computational cost and operational complexity. The architecture therefore includes bounded tool use, confidence thresholds, audit records, and a fallback path for human review. Particular attention is given to document permissions and the prevention of unsupported claims. All documents, measurements, participant identities, and reported findings are fictional. This article is original demo content created solely for user-interface, search, galley, and publication-workflow testing in OJS.

Keywords: Agentic RAG, artificial intelligence, information retrieval

DEMO CONTENT NOTICE

All names, data, cases, measurements, and findings in this file are fictional.

1. Introduction

This demonstration paper frames the problem addressed by "An Agentic RAG Architecture for Intelligent Knowledge Retrieval Systems" within a realistic but entirely fictional institutional setting. The scenario is designed to exercise scholarly metadata, full-text presentation, and information retrieval without making a claim about real participants, organizations, or measured outcomes.

The demo context intentionally includes searchable terminology related to the article title and keywords. This allows the native OJS search index and the publisher theme to be validated with plausible academic prose.

2. Method

The method uses a synthetic case, invented observations, and a structured comparison of baseline and proposed workflows. Requirements, evaluation criteria, and limitations are documented before the simulated analysis. No external dataset, copyrighted article, or real personal information is used.

The assessment is descriptive. It checks completeness, consistency, traceability, and the clarity of information presented to a fictional user. Values shown in the scenario are illustrative and are not generalizable.

3. Results and Discussion

The fictional evaluation indicates that a clearly defined workflow, consistent metadata, and observable quality checks can improve traceability and user understanding. The result should be interpreted as a design illustration rather than empirical evidence. A production implementation would require ethics review, representative data, security assessment, accessibility testing, and independent validation.

4. Conclusion

The demonstration illustrates how the proposed topic can be communicated through a conventional scholarly structure. Its only purpose is to populate and test the JOCSR OJS installation. It must not be cited or represented as a real scientific publication.

References

- [1] JOCSR Demo Editorial Team. 2026. Guidelines for Fictional System-Test Content. Internal demo document.
- [2] JOCSR Demo Lab. 2026. Synthetic Methods for Interface Validation. Fictional technical note.
- [3] No external or copyrighted research article was used to produce this PDF.