

Designing Knowledge Graph Enhanced Retrieval Systems for Institutional Data

Ryan Halim, Jessica Lim

International Journal of Research and Innovation

Abstract

This demo article explores the design of a retrieval system that combines semantic search with an institutional knowledge graph. The graph represents fictional policies, organizational units, services, documents, and their temporal relationships, while vector retrieval identifies passages relevant to a natural-language query. A routing layer decides when graph traversal, text retrieval, or a hybrid strategy should be applied. Evaluation uses synthetic documents and invented questions that test entity resolution, multi-hop relationships, source freshness, and citation completeness. The simulated results suggest that graph constraints can improve precision for relationship-heavy questions and help expose conflicting versions of institutional records. The approach also introduces maintenance costs, including entity governance, schema evolution, and synchronization with source systems. The proposed design therefore emphasizes provenance, access control, and explicit uncertainty. Every document, entity, metric, and result described in this paper is fictional. It is original demo content prepared solely to test the OJS user interface, search index, article metadata, and PDF galley workflow.

Keywords: knowledge graph, retrieval augmented generation, institutional data

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 "Designing Knowledge Graph Enhanced Retrieval Systems for Institutional Data" 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.