

# From Retrieval to Reasoning: Knowledge Graph Augmented Generative AI for Cultural Archives

RAJESH KUMAR GNANASEKARAN

RICHARD MARCIANO

University of Maryland, College Park

**Abstract:** This paper presents an applied study on the design and evaluation of ChatLoS v3, a Knowledge Graph (KG) augmented Generative AI (Gen-AI) system that enables conversational access to cultural heritage archives. Building on prior versions, ChatLoS v1, which originally implemented a Gen-AI based conversational interface to query against one of the Maryland State Archives (MSA)'s Legacy of Slavery (LoS) cultural archives dataset collection using retrieval-augmented generation (RAG) mechanism, and ChatLoS v2, which introduced an agentic approach to analytical reasoning on the same dataset, this study develops a framework that integrates a Neo4j-based KG with Gen-AI as a natural language based conversational workflow. The result is a system capable of providing natural language access to archival datasets while maintaining transparency, provenance, and explainability. The study's methodology combines system design and expert evaluation grounded in Activity Theory (AT) to explore how this approach reshapes the interaction between users and digital archives. The evaluation, conducted with professional archivists, assessed accuracy, usability, explainability, and trust across system versions. Findings indicate that KG-enhanced Gen-AI not only increases usability but also reinforces trust by grounding responses in verifiable archival evidence. The study further engages ethical concerns surrounding the use of AI in records documenting enslaved individuals, situating the system within conversations on reparative descriptive practice and emphasizing transparency, provenance preservation, and responsible narrative synthesis over speculative generation. Future work includes modularizing the pipeline to support broader institutional deployment and comparative benchmarking across archival domains.

## Introduction

Cultural heritage archives preserve vital historical narratives, yet their digital forms often remain underused. Many collections depend on conventional keyword search systems that do not support natural language or contextual exploration. This limits how scholars and the public can engage with archival materials. At the same time, Large Language Models (LLMs)<sup>1</sup> and Generative AI (Gen-AI) tools offer a new paradigm for access<sup>2</sup>, allowing users to interact conversationally with data. While these systems improve accessibility, their tendency to hallucinate or misrepresent sensitive material makes them risky for archives. In response, this project develops ChatLoS v3, a Gen-AI system enhanced with Knowledge Graph (KG)<sup>3</sup> reasoning to combine the interpretive potential of Gen-AI with the rigor of archival theory. This work continues a line of inquiry first presented at the Society of American Archivists (SAA) in 2020, in collaboration with Philip Nicholas and colleagues, which introduced Computational Archival Science (CAS) as an emerging research agenda (Nicholas et al 2020). That earlier study demonstrated how interdisciplinary collaborations among archivists, historians, and technologists could reveal new insights into the Maryland State Archives (MSA)'s Legacy of Slavery (LoS)<sup>4</sup> collections through the use

<sup>1</sup> <https://aws.amazon.com/what-is/large-language-model/>

<sup>2</sup> <https://appian.com/learn/topics/enterprise-ai/what-is-generative-ai>

<sup>3</sup> <https://neo4j.com/use-cases/knowledge-graph/>

<sup>4</sup> <https://slavery2.msa.maryland.gov/pages/Search.aspx>

of open-source tools such as OpenRefine<sup>5</sup> and Tableau<sup>6</sup>. The current paper extends that research trajectory by transforming those exploratory workflows into a fully integrated, reasoning-driven system leveraging Gen-AI that links data across collections through a Neo4j KG database. By doing so, it operationalizes the earlier vision of computational collaboration into an explainable, generative system that maintains archival provenance. Earlier versions of ChatLoS explored Retrieval-Augmented Generation (RAG)<sup>7</sup> (v1) and Agentic Reasoning (v2) explored and published in other forums (Gnanasekaran and Marciano, 2024) (Gnanasekaran and Marciano, 2025) to simplify data exploration. These systems enabled conversational querying but could not perform complex reasoning across multiple datasets or maintain transparency in how answers were generated. The current study introduces ChatLoS v3, a Neo4j-based KG augmented framework that explicitly models archival entities and relationships. This approach aims to produce transparent, explainable, and contextually accurate responses to user queries.

## Problem Statement

Despite large-scale digitization, navigating across related archival datasets remains challenging (Marciano, 2021). Names, places, and document types appear in separate systems that do not easily connect. Earlier prototypes of ChatLoS used RAG (v1) and agent-based approaches (v2) to improve search and counting tasks. Those versions, however, could not reliably resolve multi-dataset questions, trace an individual across collections, or expose the computational steps for verification. The core problem is to provide a conversational interface that can reason across linked records while preserving archival principles of provenance and context. To address this, ChatLoS v3 integrates a KG and an explainability layer that reveals the Cypher queries and records used for each answer.

Research Question 1 (RQ1): How does the integration of a Knowledge Graph augmented Generative AI framework enhance reasoning, explainability, and provenance in archival data access?

## Methodology

The study follows a design science approach using iterative prototyping and expert evaluation. The methodology includes three stages: (1) a KG construction; (2) system integration through the KG-RAG pipeline; and (3) user evaluation grounded in Activity Theory (AT). Each stage was designed to test the system's ability to address RQ1 by measuring improvements in reasoning capability, explainability, and adherence to archival provenance. In the first stage, a KG was developed using Neo4j to unify three major archival datasets from the LoS collection, the Certificates of Freedom, Domestic Traffic Advertisements, and Manumissions<sup>8</sup>. Entities such as enslaved individuals, slave owners, documents, and locations were modeled as nodes, while relationships (e.g., 'listed in,' 'owned by,' 'freed through') were represented as edges. This structure enables the system to link individuals across datasets and trace their life paths and are shown in Figure. 1, Figure. 2, and Figure. 3. The current prototype graph contains five nodes and seven relationships derived from these three datasets. Entity resolution was conducted using a manual validation to reduce duplication and ambiguity. This technical progression builds directly on earlier interdisciplinary data preparation work documented by (Nicholas et al. 2020), which identified the challenges of cross-dataset linking using manual tools. The new architecture automates these relationships through Cypher queries<sup>9</sup> (Cypher is Neo4j's declarative graph query language used to retrieve structured graph relationships) and graph embeddings by extending its scope into automated reasoning. The second

<sup>5</sup> <https://openrefine.org/>

<sup>6</sup> <https://www.tableau.com/>

<sup>7</sup> <https://aws.amazon.com/what-is/retrieval-augmented-generation/>

<sup>8</sup> <https://slavery2.msa.maryland.gov/pages/Search.aspx>

<sup>9</sup> <https://neo4j.com/docs/getting-started/cypher/>

stage integrated this graph into a Graph-RAG pipeline that translates natural language queries into Cypher commands executed on Neo4j KG database. Graph-RAG refers to a hybrid architecture that combines structured knowledge graph retrieval with LLM-based natural language summarization. Unlike traditional RAG, retrieval is constrained to explicit graph relationships, reducing hallucination risk. The results are summarized by the LLM, producing human-readable answers with embedded provenance information (links to original documents from LoS online database). The final stage involved expert evaluation conducted with two expert archivists at the MSA using Activity Theory. The evaluation draws on Activity Theory (AT) (Engeström, 1987), examining the archival activity system through three components: the subject (archivist/researcher), the object (archival inquiry), and the mediating artifact (ChatLoS v3). The system was analyzed as a tool reshaping knowledge production and interpretive practice within the archival environment. Participants engaged in structured testing sessions comparing ChatLoS v1, v2, and v3. A standardized rubric assessed accuracy, explainability, usability, and trust. This evaluation was to interpret how the AI system mediated archival practices, evolving from collaboration among humans to collaboration between humans and AI systems.

## Results

The integration of a KG-RAG enabled ChatLoS v3 to perform reasoning tasks from a conversational interface that were not possible in earlier versions. For example, when asked to 'trace the life of an individual across the datasets linked via the KG', the system identified an individual named "John Howard's" records from three datasets: a sale advertisement (1829), a manumission record (1830), and a Certificate of Freedom (1830) only with a natural language query. It then generated a narrative summary connecting these events, illustrating how the individual's life journey could be reconstructed across disparate archival sources as shown in Figure. 4 and Figure. 5. The natural language responses also included references to the original document hosted by MSA's online database, which is shown in Figure. 6. This demonstrates the system's capacity for multi-hop reasoning and contextual synthesis. ChatLoS v3 produced correct aggregate counts in all test cases, including a query asking for the number of advertisements in "Dorchester County" between 1825 and 1830, where it correctly returned seventy-eight results. Unlike prior versions, which occasionally produced inconsistent data, ChatLoS v3's reliance on Cypher queries ensured factual precision. Participants consistently reported higher confidence in the responses, citing the system's transparent presentation of its reasoning steps as a major factor. The architecture is reproducible in other archival contexts provided structured datasets are available. Core components include (1) entity normalization, (2) graph schema modeling, (3) Cypher query translation, and (4) LLM summarization with provenance embedding.

## Ethical Considerations: AI, Archival Harm, and Reparative Discussion

The application of generative AI to records documenting enslaved individuals requires careful ethical consideration. These records were created within oppressive systems and reflect structural inequities and archival silences. While KG grounding reduces hallucination risk, AI-generated narrative synthesis may inadvertently reframe traumatic histories through computational abstraction. This study situates ChatLoS v3 within ongoing archival conversations around reparative descriptive practice, which seeks to confront harmful legacy description, center marginalized individuals, and contextualize records created through systems of racial violence. Reparative description does not simply enhance access; it critically interrogates how archival systems frame individuals, particularly enslaved people, primarily through transactional, legal, or property-based language. Generative AI introduces both opportunity and risk within this context. On one hand, conversational synthesis can help surface fragmented life traces across datasets, potentially humanizing individuals whose records are otherwise dispersed and decontextualized. On the other hand, automated narrative construction risks re-inscribing archival bias if the system merely reproduces the structural logic embedded in historical documentation. ChatLoS v3 addresses this risk by prioritizing transparency, exposing Cypher queries, and embedding direct links to original records rather

than generating speculative interpretation. By foregrounding provenance and computational traceability, the system resists presenting AI-generated summaries as interpretive truth. Instead, it positions AI as a mediating tool that supports but does not replace critical archival interpretation. In this sense, ChatLoS v3 aligns with reparative principles by enhancing access while maintaining contextual integrity, acknowledging archival bias, and avoiding speculative completion of historical gaps. The system is intentionally designed to support reflective engagement with the archive rather than automated narrative authority.

## Findings

The findings directly address and answer the RQ1 by demonstrating that a KG augmented framework strengthens reasoning and transparency in archival Gen-AI systems. By grounding responses in explicit graph relationships, ChatLoS v3 eliminates the 'black box'<sup>10</sup> effect common to generative models. Each answer is traceable to its source, preserving the archival principles of provenance and context. This progression fulfills the computational vision articulated by (Nicholas et al. 2020), evolving from human-facilitated interdisciplinary exploration to a semi-automated reasoning environment. The inclusion of AT in the evaluation provided insight into how the system mediated human interaction. Users transitioned from passively retrieving records to actively exploring narratives. This transformation suggests that Gen-AI can serve as a partner in archival research rather than a replacement for human expertise. Participants described ChatLoS as a tool that 'thinks with' them, encouraging inquiry while maintaining respect for data integrity. These observations underline the potential of explainable Gen-AI to enhance both technical performance and ethical alignment.

## Limitations

While promising, this study has several limitations. First, the KG currently represents a subset of the full LoS corpus. Expansion to the complete dataset may introduce entity resolution complexity and scalability challenges. Second, system development required expertise in graph modeling and query design. Although end users interact through natural language, institutional adoption may require technical training and infrastructure investment. Third, the expert evaluation involved a limited sample size of two archivists, which constrains generalizability. Broader user studies involving historians, educators, and community researchers are needed. Finally, archival datasets themselves contain historical bias and incomplete records. The KG does not eliminate structural archival silences but instead provides structured pathways for navigating them.

## Conclusion

This study demonstrates that integrating KG reasoning within Gen-AI systems significantly improves the quality, transparency, and interpretability of archival access. ChatLoS v3 fulfills its research objective by enabling reasoning grounded in provenance, ensuring factual accuracy, and supporting scholarly interpretation through explainable responses. Future research will focus on large-scale user evaluations, quantitative accuracy benchmarking, bias mitigation strategies in entity linking, and integration of community-informed metadata practices. Future work includes modularizing the pipeline to support broader institutional deployment and comparative benchmarking across archival domains. Additionally, participatory feedback mechanisms will be developed to allow users to flag potential inaccuracies or ethical concerns in generated narratives. By embedding trust and explainability into the design, ChatLoS establishes a foundation for responsible AI use in cultural heritage archives.

---

<sup>10</sup> <https://www.ibm.com/think/topics/black-box-ai>

Relationship Type	Source Node	Target Node
OWNED_ESLAVED_PERSON	SlaveOwner	EnslavedPerson
PLACED_DOMESTIC_TRAFFIC_AD	SlaveOwner	DomesticTrafficAd
LISTED_FOR_DOMESTIC_TRAFFIC_SALE_IN	EnslavedPerson	DomesticTrafficAd
SIGNED_MANUMISSION	SlaveOwner	ManumissionRecord
MANUMITTED_THROUGH	EnslavedPerson	ManumissionRecord
WAS_GRANTED_FREEDOM_THROUGH	EnslavedPerson	CertificateOfFreedom
WHO_ESLAVED_WAS_FREEDOM_FROM	CertificateOfFreedom	SlaveOwner

Figure 1. Relationships for the ChatLoS Knowledge Graph in Neo4j

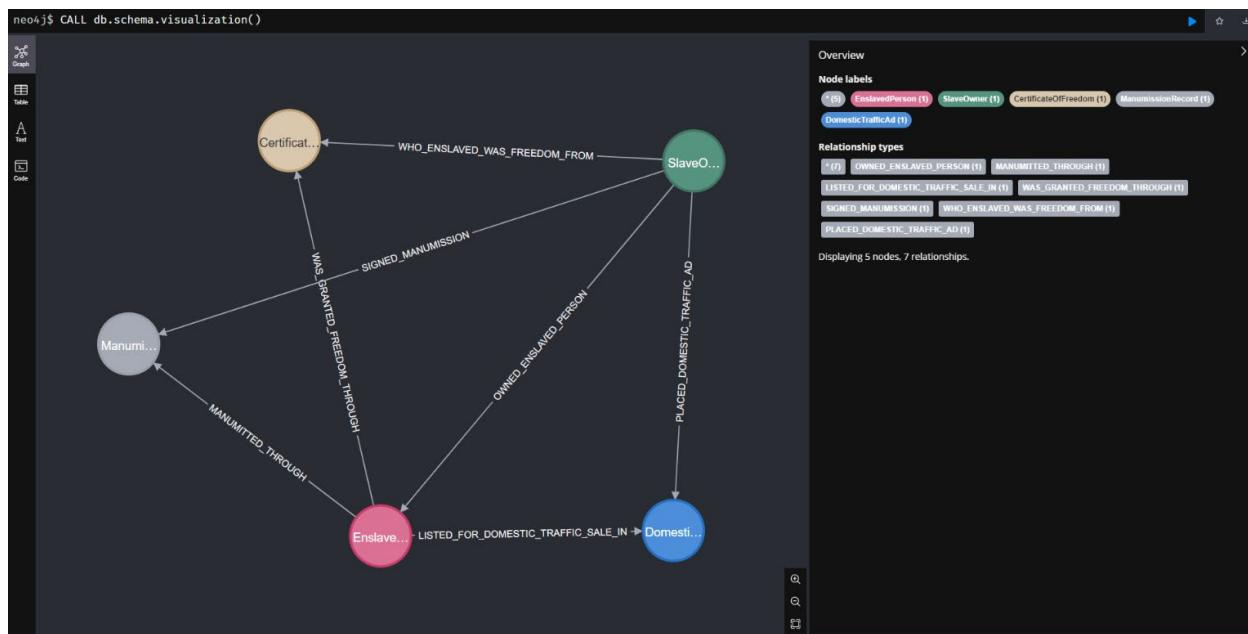


Figure 2. A subgraph showing the entity–relationship of the ChatLoS Knowledge Graph in Neo4j

<b>Node Label</b>
EnslavedPerson
SlaveOwner
ManumissionRecord
CertificateOfFreedom
DomesticTrafficAd

**Figure 3. Node Labels of the ChatLoS Knowledge Graph in Neo4j**

This progression strongly suggests that John Howard transitioned from being enslaved, being listed for sale, gaining his freedom, and then receiving official recognition of that freedom. This kind of detailed tracing can further our understanding of the process of liberation during this period.

Generative AI explanation behind the response

What the AI did:

The AI interpreted your question: 'Can you trace the journey of individuals from being listed in Domestic Traffic Ads to receiving a Certificate of Freedom, including any manumission records?', generated a Cypher query to explore graph relationships in the Knowledge Graph, executed it against Neo4j, and composed an answer from the query results.

Internal Logs (Advanced):

```
> Entering new GraphCypherQACChain chain...
Generated Cypher:
cypher
MATCH (e:EnslavedPerson)-[:LISTED_FOR_DOMESTIC_TRAFFIC_SALE_IN]->(ad:DomesticTrafficAd)
  (e)-[:MANUMITTED_THROUGH]->(m:ManumissionRecord),
  (e)-[:WAS_GRANTED_FREEDOM_THROUGH]->(f:CertificateOfFreedom)
RETURN e.first_name AS FirstName,
       e.last_name AS LastName,
       ad.Ad_Image_Metadata AS AdImageLink,
       ad.Sale_Date AS SaleDate,
       m.DateManumitted AS ManumissionDate,
       f.date_freed AS FreedomDate
ORDER BY FirstName, LastName, SaleDate
```

```
Full Context:
[{'FirstName': 'John', 'LastName': 'Howard', 'AdImageLink': 'http://msa...
```

```
> Finished chain.
DEBUG: Chain Result: {'query': 'Can you trace the journey of individual...
```

 Clear Chat

**Figure 4. Example ChatLoS v3 Output Tracing an Individual's Life Path**

# ChatLoS enhanced with KG-RAG using Neo4j and Generative AI—Cross-dataset querying with DTA, Manumissions, and Certificates of Freedom



Can you trace the journey of individuals from being listed in Domestic Traffic Ads to receiving a Certificate of Freedom, including any manumission records?



Yes, based on the available data, we can trace the journey of an individual named John Howard from being listed in Domestic Traffic Ads to receiving his Certificate of Freedom.

1. **Domestic Traffic Ads:** John Howard was listed for sale in a Domestic Traffic Ad on **August 15, 1829**. You can view the **original ad** through the following **source: 1829 Domestic Traffic Ad for John Howard**.
2. **Manumission Records:** Approximately **a year later**, his Manumission Record shows that he was granted freedom on **September 28, 1830**.
3. **Certificate of Freedom:** Subsequently, he was officially recognized as a free individual through his **Certificate of Freedom** dated **October 25, 1830**.

This progression strongly suggests that John Howard transitioned from being enslaved, being listed for sale, gaining his freedom, and then receiving official recognition of that freedom. This kind of detailed tracing can further our understanding of the process of liberation during this period.

Generative AI explanation behind the response

What the AI did:

The AI interpreted your question: 'Can you trace the journey of individuals from being listed in Domestic Traffic Ads to receiving a Certificate of Freedom, including any manumission records?', generated a Cypher query to explore graph relationships in the Knowledge Graph, executed it against Neo4j, and composed an answer from the query results.

Figure 5. Example ChatLoS v3 Output Tracing an Individual's Life Path—Continued



**Figure 6. Example ChatLoS v3 Output Referencing the link to the Original Document on MSA's online db**

### Acknowledgments

We wish to acknowledge the support of four IMLS (Institute of Museum and Library Services) LB21 (Laura Bush 21st Century Librarian) grants: (1) National Forum grant – “Computational Thinking to Library and Archival Science Education & Research, CT-LASER” (RE-73-18-0105-18) (2018-2020: [link](#)), (2) NDII (National Digital Infrastructures and Initiative) Project grant – “Piloting an Online Collaborative Network for Integrating CT into Library and Archival Education and Practice, Piloting Network” (RE-246334-OLS-20) (2020-2024: [link](#)), (3) Implementation grant – “Training of Archival & Library Educators with iNnovative Technologies, TALENT Network” (RE252287-OLS-22) (2022:2025: [link](#)), and (4) Applied research grant – “GenAI-4-Arch, Harnessing Generative AI to Support Exploration and Discovery in Library and Archival Collections<sup>11</sup>” (LG-256565-OLS-24) (2024:2026: [link](#)).

These grants are advancing data science education and infrastructure across the library, archives, and museum ecosystems. This work benefited immeasurably from the contributions of Dr. Michael J. Kurtz, RIP Dec. 17, 2022, who was an inspiring mentor, collaborator, and dear friend of many in this CAS research community. Dr. Kurtz was grounded in the core principles of archives and their social value, as well as a visionary who saw the potential of computational technology in this cultural work.

<sup>11</sup> <https://heritage-ai.org/>

## Resources

Engeström, Yrjö. 1987. "Learning by Expanding: An Activity-Theoretical Approach to Developmental Research." Helsinki: Orienta-Konsultit.

Nicholas, Philip, Rajesh Kumar Gnanasekaran, Lori Perine, Alexis Hill, and Richard Marciano. 2020. "Establishing a Research Agenda for Computational Archival Science through Interdisciplinary Collaborations between Archivists and Technologists." Society of American Archivists.

Marciano, Richard. 2021. "Afterword: Toward a New Discipline of Computational Archival Science." In Re-Envisioning the MLS: Perspectives on the Future of Library and Information Science Education, 179–199. Bielefeld University Press.

Gnanasekaran, Rajesh Kumar., Marciano, Richard. 2024. "Leveraging LLMs and Learning-as-a-Service to Create a Culturally Rich Conversational Chatbot." International Symposium on Grids & Clouds.

Gnanasekaran, Rajesh Kumar., Marciano, Richard. 2025. "Enhancing Access to Cultural Archives through Data Science, Generative AI, and Knowledge Graphs." Doctoral dissertation, University of Maryland.