POSTER

Athena: A Conversational Book Discovery Interface Combining LLM-Powered Retrieval-Augmented Generation and Interactive Graph Visualization

MATT MURTAGH-WHITE, Trinity College Dublin, Dublin, Leinster, Ireland

YUNKAI XU, Pennsylvania State University, University Park, PA, United States

NICOLE LEÓN, Pennsylvania State University, University Park, PA, United States

FRANK E RITTER, Pennsylvania State University, University Park, PA, United States

Open Access Support provided by:

Pennsylvania State University

Trinity College Dublin

Published: 28 September 2025

Citation in BibTeX format

UIST '25: The 38th Annual ACM Symposium on User Interface Software and Technology
September 28 - October 1, 2025
Busan, Republic of Korea

Conference Sponsors:
SIGCHI
SIGGRAPH

Athena: A Conversational Book Discovery Interface Combining LLM-Powered Retrieval-Augmented Generation and Interactive Graph Visualization

Matt Murtagh White*
ADAPT Centre
Trinity College Dublin
Dublin, Ireland
mmurtagh@tcd.ie

Yunkai Xu*
College of Information Sciences and Technology
Pennsylvania State University
State College, Pennsylvania, USA
yqx5322@psu.edu

Nicole León*
College of Information Sciences and Technology
Pennsylvania State University
State College, Pennsylvania, USA
nicoleleon@psu.edu

Frank E. Ritter
College of Information Sciences and Technology
Pennsylvania State University
State College, Pennsylvania, USA
frank.ritter@psu.edu

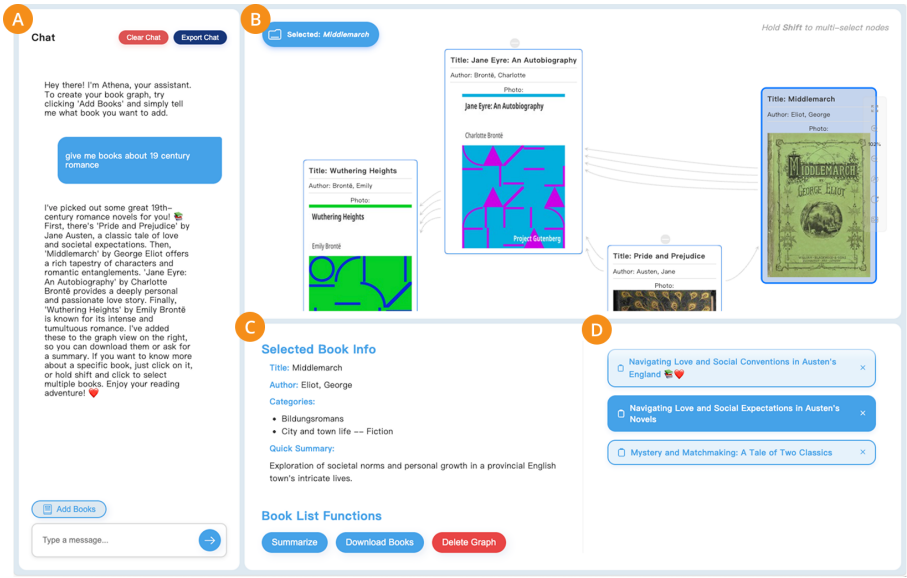

Figure 1: The Athena user interface, combining a chatbot with an interactive graph to support book exploration. (A) The Chatbot Panel enables natural language queries and returns rich, context-aware responses. (B) The Relationship Graph visualizes books and their interconnections, allowing users to explore the relationships of chatbot's response. (C) The Book Information Panel displays metadata and provides actionable functions such as summarization or download. (D) The Summary Panel stores GPT-generated thematic interpretations, supporting reflection and comparison.

Abstract

Despite advances in digital libraries, keyword-based search remains rigid, offering limited support for exploratory and sense-making tasks. We introduce Athena, a book discovery system that integrates LLM-powered Retrieval-Augmented Generation (RAG) with interactive graph visualization. This hybrid system allows users to engage in natural language dialogue, navigate relational graphs of retrieved books, and generate cross-book summaries, offering an alternative to static keyword search. A preliminary user study found that Athena reduced cognitive load, improved usability, and encouraged exploratory behavior, although user trust in AI-generated

*All authors contributed equally to this research.

UIST Adjunct '25, Busan, Republic of Korea

ACM ISBN 979-8-4007-2036-9/25/09
https://doi.org/10.1145/3746058.3758400

content varied. We outline future directions focused on scaling to larger, more diverse user studies and systematically analyzing how conversational and visual features influence trust, satisfaction, and external validation behaviors.

CCS Concepts

• **Information systems** → *Search interfaces.*

Keywords

Human-centered Design, Searching System, Interaction Design

ACM Reference Format:
Matt Murtagh White, Yunkai Xu, Nicole León, and Frank E. Ritter. 2025. Athena: A Conversational Book Discovery Interface Combining LLM-Powered Retrieval-Augmented Generation and Interactive Graph Visualization. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST Adjunct '25), September 28–October 01, 2025, Busan, Republic of Korea.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3746058.3758400

1 Introduction

Recent advances in large language models (LLMs) have enabled a shift from keyword-based retrieval to dialogue-driven search [10, 17]. This shift is especially promising for exploratory domains where users generally have vague goals, abstract concepts, or unfamiliar domain knowledge [7, 12]. However, in large-scale digital collections such as Project Gutenberg [13], access remains largely limited to keyword search and hierarchical browsing. These modes, while efficient for targeted retrieval, offer limited support for curiosity-driven or exploratory engagement [3, 11].

Exploratory search, which is often vague, iterative, and evolving [8], requires systems that support sense-making and discovery, not just retrieval [1, 16]. Despite significant advances in visual interfaces and metadata design [2, 9, 18], mainstream digital libraries still remain largely plain, static, and text-driven. They offer limited scaffolding for sense making, reflection, or iterative exploration [8, 14, 15]. Prior studies have shown that visual interfaces can aid users in navigating complex information spaces [5, 7]. Building on this, we argue that combining LLM-powered conversational interfaces with interactive visualizations presents a promising direction to better support open-ended, curiosity-driven exploration in digital archives.

Therefore, we present Athena, a hybrid interface that combines large language model–driven dialogue with visual and interactive exploration. Users can chat with an AI chatbot powered by retrieval-augmented generation (RAG) (Figure 1 A), explore book relationships through an interactive graph(Figure 1 B), view detailed metadata(Figure 1 C), and generate thematic comparisons and summaries across selected books(Figure 1 D). Athena is designed to scaffold curiosity-driven discovery and support more meaningful engagement with large literary collections.

2 System Design

Athena supports both structured and open-ended book discovery through a tightly integrated conversational and visual interface. User interaction begins with natural language queries to a chatbot, which interprets intent and retrieves relevant books from the

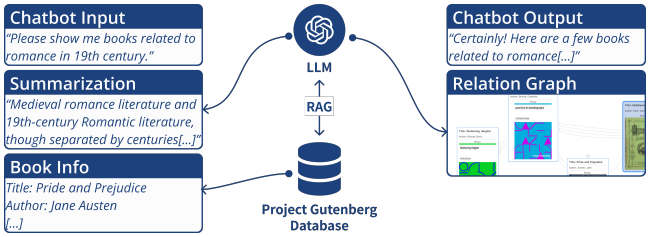

Figure 2: Overview of our system design, which combines an LLM(GPT-4o) with the Project Gutenberg database.

Project Gutenberg corpus via the Gutendex API [6], enriching results with metadata including subjects, authors, languages, cover images, and download links (Figure 2).

A core feature is the dynamic relation graph (Figure 2), where books are visualized as nodes and their semantic relationships—such as shared authorship, genre, or thematic similarity—are inferred in real time by an LLM based on book metadata. The LLM generates edges, justifications for these links, and concise 15-word summaries for each book, allowing users to quickly parse the rationale for inclusion.

Users can select books directly from the graph for comparison, summarization, or to provide focused context for conversational queries. When one or more books are chosen, the LLM synthesizes key themes and differences into a comparative summary, or enables follow-up questions with those books as explicit context, using metadata and summaries from Gutendex. This end-to-end system of retrieval, visualization, and summarization via conversational input is designed to reduce barriers to exploratory search, enhance transparency, and scaffold reflection within large digital collections.

3 Preliminary User Study

Method. We conducted a within-subjects, counterbalanced study with eight graduate participants, comparing Athena, a RAG-powered conversational agent with GPT-4o and interactive book-graph visualization, to the standard Project Gutenberg website. Each participant completed eight information-seeking tasks per system (four lookup, four exploratory), in a randomized order. Behavioral data, including query counts and types, time-on-task, summary use, external verification, and navigation patterns, were logged via screen recordings as participants verbalized their thought processes (think-aloud protocol). Exploration and trust were inferred from these behavioral logs and post-task surveys, which included the NASA Task Load Index (TLX, 0–10 scale) [4] and Likert items on trust, usability, and exploration support. All procedures were approved by IRB.

Findings. Participants reported substantially lower cognitive and emotional workload with Athena than with Project Gutenberg, as reflected in NASA TLX ratings for mental demand (M = 2.25 vs. 5.13), effort (M = 2.75 vs. 4.63), and frustration (M = 2.13 vs. 5.00), and higher perceived performance (M = 7.88 vs. 5.00). Participants also completed both lookup and exploratory tasks more quickly on Athena, suggesting greater efficiency for both focused and open-ended search. Qualitative analysis of behavioral logs and think-aloud protocols revealed that Athena users engaged in more

frequent exploratory actions—such as multi-turn queries and book comparisons and switched less often to external resources (e.g., web search) for verification. This reduction in verification may be due to richer contextual responses provided by the chatbot, or differing levels of trust and perceived sufficiency of information between the two interfaces. Overall, these results suggest that Athena's multimodal design supports more sustained exploration and potentially higher satisfaction during open-ended search, although further research is needed to disentangle causal factors and assess long-term trust and engagement.

4 Future Work

Athena's conversational and visual interface fostered more exploratory search behaviors than the baseline, with participants issuing more multi-turn queries and conducting book comparisons. Subjective workload was lower and tasks were completed more efficiently with Athena. Participants also engaged less in external validation (e.g., web search), suggesting that richer contextual responses and transparent relational explanations may enhance user confidence, though this raises open questions about trust calibration, as lower validation could reflect either improved sufficiency or misplaced confidence. Overall, these results highlight the promise of hybrid conversational–visual search for large corpora. While our study focused on book recommendation and search tasks, the underlying interaction model is broadly applicable to many other domains where users must explore complex information spaces. We plan to expand this work by recruiting a broader participant pool and systematically investigating how specific interface features influence users' trust, satisfaction, and validation strategies over time.

Acknowledgments

We would like to thank Frank Ritter for his contributions to this paper. This work was conducted with the financial support of Taighde Éireann Centre for Research Training in Artificial Intelligence under Grant No. 18/CRT/6223 and the European Union's Horizon Europe Research and Innovation Programme (grant agreement #101070125).

References

[1] Jürgen Bernard, Debora Daberkow, Dieter Fellner, Katrin Fischer, Oliver Koepler, Jörn Kohlhammer, Mila Runnwerth, Tobias Ruppert, Tobias Schreck, and Irina Sens. 2015. Visinfo: a digital library system for time series research data based on exploratory search—a user-centered design approach. *International Journal on Digital Libraries* 16 (2015), 37–59.

[2] Miriam Boon, Orland Hoeber, Larena Hoeber, Dale Storie, and Veronica Ramshaw. 2024. Exploratory Search and Beyond: A Study of Complex Search Scenarios within a Public Digital Library. *Proceedings of the Association for Information Science and Technology* 61, 1 (Oct. 2024), 44–55. doi:10.1002/pra2.1007

[3] Zeljko Carevic, Maria Lusky, Wilko van Hoek, and Philipp Mayr. 2018. Investigating exploratory search activities based on the stratagem level in digital libraries. *International Journal on Digital Libraries* 19 (2018), 231–251.

[4] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9

[5] Orland Hoeber. 2025. Design Principles for Exploratory Search Interfaces. In *Proceedings of the 2025 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '25)*. Association for Computing Machinery, New York, NY, USA, 12–22. doi:10.1145/3698204.3716443

[6] Gareth B. Johnson. 2020. Gutendex: An open API for Project Gutenberg metadata. https://github.com/garethbjohnson/gutendex. Accessed March 2025.

[7] Nabin Khanal, Chun Meng Yu, Jui-Cheng Chiu, Anav Chaudhary, Ziyue Zhang, Kakani Katija, and Angus G. Forbes. 2024. FathomGPT: A natural language interface for interactively exploring ocean science data. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) *(UIST '24)*. Association for Computing Machinery, New York, NY, USA, Article 95, 15 pages. doi:10.1145/3654777.3676462

[8] Gary Marchionini. 2006. Exploratory search: from finding to understanding. *Commun. ACM* 49, 4 (April 2006), 41–46. doi:10.1145/1121949.1121979

[9] Lori McCay-Peet, Anabel Quan-Haase, and Dagmar Kern. 2015. Exploratory search in digital libraries: a preliminary examination of the use and role of interface features. In *Proceedings of the 78th ASIS&T Annual Meeting: Information Science with Impact: Research in and for the Community* (St. Louis, Missouri) *(ASIST '15)*. American Society for Information Science, USA, Article 70, 4 pages.

[10] Fengran Mo, Kelong Mao, Ziliang Zhao, Hongjin Qian, Haonan Chen, Yiruo Cheng, Xiaoxi Li, Yutao Zhu, Zhicheng Dou, and Jian-Yun Nie. 2024. A Survey of Conversational Search. https://arxiv.org/abs/2410.15576

[11] Andrea Papenmeier, Dagmar Kern, Daniel Hienert, Alfred Sliwa, Ahmet Aker, and Norbert Fuhr. 2021. Starting Conversations with Search Engines - Interfaces that Elicit Natural Language Queries. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval* (Canberra ACT, Australia) *(CHIIR '21)*. Association for Computing Machinery, New York, NY, USA, 261–265. doi:10.1145/3406522.3446035

[12] Seokhyeon Park, Yumin Song, Soohyun Lee, Jaeyoung Kim, and Jinwook Seo. 2025. Leveraging Multimodal LLM for Inspirational User Interface Search. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 579, 22 pages. doi:10.1145/3706598.3714213

[13] Project Gutenberg. 2024. Project Gutenberg. https://www.gutenberg.org/. Accessed: 2025-05-01.

[14] Frank E. Ritter, Graham D. Baxter, and Elizabeth F. Churchill. 2014. Cognition: Human–Computer Communication. In *Foundations for Designing User-Centered Systems*. Springer, London, 201–223. doi:10.1007/978-1-4471-5134-0_7

[15] Frank E. Ritter, Graham D. Baxter, and Elizabeth F. Churchill. 2014. Cognition: Memory, Attention, and Learning. In *Foundations for Designing User-Centered Systems*. Springer, London, 123–164. doi:10.1007/978-1-4471-5134-0_5

[16] Rao Shen, Naga Srinivas Vemuri, Weiguo Fan, Ricardo da S. Torres, and Edward A. Fox. 2006. Exploring digital libraries: integrating browsing, searching, and visualization. In *Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries* (Chapel Hill, NC, USA) *(JCDL '06)*. Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/1141753.1141755

[17] Ryen W. White. 2024. Advancing the Search Frontier with AI Agents. *Commun. ACM* 67, 9 (Aug. 2024), 54–65. doi:10.1145/3655615

[18] Chang Xiao. 2023. AutoSurveyGPT: GPT-Enhanced Automated Literature Discovery. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) *(UIST '23 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 1, 3 pages. doi:10.1145/3586182.3616648