

From Card Catalogs to Semantic Search

Building a Human-Centered Discovery Platform Powered by AI Technologies

Carolyn Caizzi and Amy Deschenes

ABSTRACT

The first phase of the Reimagining Discovery project at Harvard Library sought to address the challenge of fragmented search experiences of special collections materials using artificial intelligence (AI) technologies, such as embedding models and large language models (LLMs). The resulting platform, Collections Explorer, simplifies and enhances the search experience for more effective special collections discovery. The project team took a user-centered and trustworthy approach to implementing AI, grounding the choices of the platform in user empowerment and librarian expertise. The development process included extensive user research, including interviews, usability testing, and prototype evaluations, to understand and address user needs.

Collections Explorer was developed using a multi-component architecture that integrates multiple types of AI. The team evaluated more than 12 models to select ones that were the best fit for the need, as well as being ethical and sustainable. Detailed system prompts were developed to guide LLM outputs and ensure the reliability of information. The methodical and iterative approach helped to create a flexible and scalable platform that could evolve to support other material types in the future. Initial research showed that potential users are enthused at the prospect of AI-powered features to enhance discovery, especially the item-level summaries and related search suggestions. The project demonstrated the potential of integrating AI technologies into library discovery systems while maintaining a commitment to trustworthiness and user-centered design.

INTRODUCTION

Reimagining Discovery is a multi-year strategic project to transform how researchers, students, and the global community access and explore Harvard Library's extensive collections. The central goal of the project is to make the experience of finding library materials easier for all users. This article focuses on the creation of an experimental AI-enabled discovery platform for special and archival collections, called Collections Explorer. Building the Collections Explorer platform is only one part of the Reimagining Discovery project. There are additional workstreams to launch an improved finding aid discovery platform using Arclight and rebuild the library's metadata hub, which ingests metadata from disparate sources and delivers normalized metadata to other applications. The work involved with Reimagining Discovery will simplify and enhance Harvard Library's large and complex online search experience.

The Harvard Library discovery environment includes a variety of discovery platforms for different types of materials. The scope of the materials is vast and spans many formats, including scholarly

About the Authors

Carolyn Caizzi (ccaizzi@berkeley.edu; corresponding author) is Associate University Librarian for Digital Initiatives and Information Technology, University of California, Berkeley Library. **Amy Deschenes** (amy_deschenes@harvard.edu) is Interim Director of User Experience & Discovery, Harvard University Library. © 2026.

Submitted: 25 July 2025. Accepted for Publication: 22 October 2025. Published: 16 March 2026.

articles, books, images, maps, datasets, and archival materials. In total, the library provides access to billions of items, ranging from print books and journals to digitized photographs, geospatial data, and archival finding aids that describe unique historical collections. These materials are made available through a wide range of discovery platforms, each optimized for different formats and research needs. The environment is perceived by users as complex and siloed.¹ As of 2025, the library uses Ex Libris Primo for the main catalog, article discovery, and image discovery; Spotlight for digital collections; GeoBlacklight for maps and geospatial data; Dataverse for data; DSpace for the institutional repository; and Arclight for finding aids.

Based on user research focused on special collections and archives discovery, the project team identified a significant opportunity to make finding these materials easier, especially for people who may be unfamiliar with the techniques of archival discovery. Around the same time this opportunity to simplify and enhance discovery was identified, technological advances with artificial intelligence (AI), both with embedding models and large language models (LLMs), presented themselves.² Through Collections Explorer, the team sought to capitalize on this significant opportunity to improve the user experience for online discovery of special collections by utilizing cutting-edge AI technology solutions.

The development of Collections Explorer is grounded in a commitment to human-centered AI, prioritizing user empowerment and professional librarian expertise. While leveraging advanced AI technologies, the project maintains a critical human review process where librarians actively review and guide technological implementations. This approach ensures that the platform not only utilizes AI capabilities but also reflects the core values of libraries: providing reliable, contextual, and trustworthy information access. By keeping librarians integral to the technological development process, the team sought to create an AI-enabled discovery tool that enhances rather than replaces human research expertise, maintaining the library's role as a thoughtful and responsible knowledge partner.

Achieving the goal of building a trustworthy platform to simplify the search experience for users required more than simply improving search functionality; it required a deep understanding of users' needs. This understanding is the foundation of user-centered design, which is defined as "a philosophy based on the needs and interests of the user, with an emphasis on making systems usable and understandable."³ To achieve this, the team included users at every step of the process, from exploratory research and design reviews to usability testing.⁴ It was important that the team defined users as specifically as possible to build the right solution.⁵ The team intentionally chose one main audience for the first phase of work: people who are interested in special collections and archival materials but have not had training or background in how to find them.

In addition to clearly defining our audience, a significant amount of time was spent at the beginning of the project to determine the scope, strategy, and technical approach. When that information was finalized, the team built an alpha version of the platform that was tested with users. The feedback from that testing was incorporated into the next iteration of the platform, and the technical approach was further refined. The strategy, scope, and technical direction of the Collections Explorer platform were shaped not only by local user research but also by prior scholarship on discovery systems, the impact of AI on search, and best practices in user-centered design.

LITERATURE REVIEW

Library discovery has undergone significant transformations, evolving from physical card catalogs to online discovery systems that attempt to mirror the intuitive search experiences of the broader web. Early library discovery can be traced back to printed book catalogs, which gave way to card catalogs in the late 19th century as libraries required more flexible, updatable systems.⁶ The introduction of machine-readable cataloging (MARC) in the 1960s, pioneered by Henriette Avram, set the stage for computerized catalogs, culminating in the emergence of Online Public Access Catalogs (OPACs) in the 1980s.⁷

By the late 1990s, OPACs had transitioned to web-based interfaces, and by the mid-2000s, libraries moved toward discovery layers—interfaces designed to provide unified search across various library resources.⁸ This shift was a response to criticisms of Boolean-based search models and efforts to introduce more user-friendly, relevance-ranked searching similar to commercial search engines like Google.⁹ Karen Calhoun, in her influential report for the Library of Congress, argued that traditional catalogs had become unsustainable and urged libraries to integrate catalogs with web-based discovery tools and user workflows.¹⁰ This evolution also marked a broader convergence between traditional librarianship and information science, as libraries increasingly adopted indexing methodologies and search algorithms akin to those used in web search environments.¹¹

Despite their intent to simplify access, discovery layers have faced ongoing scrutiny regarding their effectiveness in supporting academic research. Some argue that discovery tools often prioritize user autonomy and search aesthetics over research precision, leading to mixed results.¹² This reflects broader concerns regarding user reliance on keyword-based searches without a clear understanding of system mechanics.¹³ Modern keyword search still suffers from vocabulary mismatch issues, as it depends on exact word matches rather than conceptual understanding.¹⁴ As library discovery continues to evolve, there is a need for systems that integrate with how users find information more broadly on the web rather than expect users to adapt to rigid search paradigms.¹⁵

As part of their work, the project team conducts regular user research to understand our users and address their needs. To build useful discovery tools, the team needs to have a solid understanding of users through observation and interviews.¹⁶ Even before naming the project “Reimagining Discovery,” research was conducted to inform the project proposal. Generative research helps define a problem and can reveal ideas for solutions.¹⁷ The generative research conducted for the project was accomplished through interviews with special collections users and researchers using generative AI for their academic work. The team conducted interviews to gain insights, discover problems, and inform their objectives for this library discovery work.¹⁸

This generative research shaped the goals and objectives for the project. However, to keep the project user-centered throughout its entirety, the team conducted extensive evaluative research at regular intervals. To assess how well the tool works for our intended audience, they conducted usability testing and prototype evaluations and solicited design feedback from users.¹⁹

Recent advancements in AI-driven discovery tools suggest a new trajectory for library search systems, opening opportunities to rethink how users interact with information. Companies like EBSCO and Clarivate are integrating features such as natural language search and AI-generated article insights to enhance the discovery experience.²⁰ These tools aim to address longstanding search challenges by enabling more intuitive, question-based queries and providing

summarization features that help reduce information overload. This shift aligns with broader trends in digital knowledge management, where linked data and machine learning are being leveraged to improve retrieval relevance and overall user experience.²¹

Building on these developments, the project also aimed to explore the integration of generative AI features into the discovery interface, incorporate natural language processing, and develop a semantic search tool.²² Semantic search is a type of search engine technology that understands the intent and context behind words and phrases. Unlike traditional keyword searches, semantic search retrieves results based on their underlying meaning instead of merely matching the exact terms used in the query. These innovations represent a fundamentally new approach to library discovery—one that requires thoughtful implementation. The team is committed to a human-centered AI (HCAI) approach, which emphasizes empowering users rather than replacing them. HCAI shifts the focus from purely technical goals to human-centered ones, ensuring that systems uphold core human values.²³ The team worked to review outputs from the embedding model and LLM, ensuring there was always a human in the loop for guiding and reviewing the technology choices. By participating in the development of these technologies, librarians can help shape the roadmap of emerging discovery tools, ensuring that the library continues to serve as a critical research partner and that professional values remain central to the adoption of AI in libraries.

Providing access to trustworthy information has always been a fundamental value in librarianship, as has creating trust between libraries and their patrons.²⁴ Even beyond the library environment, trust in systems and tools is critical for their successful adoption and continued use.²⁵ With this in mind, a critical component of the work is ensuring that the adoption of semantic search and AI-powered features results in a trustworthy and ethical library discovery tool. To that end, the team followed best practices for developing a transparent and explainable system to ensure its trustworthiness.²⁶

PRELIMINARY EXPERIMENTATION

The groundwork for Collections Explorer was laid in fall 2023, when a group of library staff received internal grant funding to build a prototype of a natural language search tool that queried our library catalog, HOLLIS, which is built on Primo from Ex Libris. Enabling support for natural language allows users to use full questions or conversational search terms like “What is the history of germ theory?” instead of constructing keyword searches like “germ theory AND history.” The objective of this grant proof-of-concept project was to build a working proof-of-concept that enabled natural language searching and test it with real users to inform future AI work for library discovery.

The goal of this proof-of-concept, dubbed “Talk with HOLLIS,” was to learn more about LLMs and the opportunities they present for enhancing library discovery. The team increased their knowledge and understanding of LLM technology during the pilot project. The prototype gave the team the ability to get direct feedback from potential users and understand what existing pain points a natural language search of library collections could address.

During the pilot project, the team was able to approximate the level of effort involved in developing a more robust discovery system built to power natural language searching and determine where best to focus the initial phase of work for Reimagining Discovery. More details about the testing are discussed in the User-Centered Design and Research section.

PROJECT STRATEGY

Product Approach

Building on the success and momentum of the “Talk with HOLLIS” proof-of-concept, the team collaborated with Mozilla.ai to determine our strategic approach and define the first phase of work on this new discovery platform. The team chose to work with Mozilla.ai because of the company’s commitment to building ethical, trustworthy, and open AI tools. Mozilla.ai provided support in determining the technical approach, provided recommendations on specific embedding models and LLMs, and worked with local developers to deepen the team’s knowledge of developing AI technologies.

Based on what was learned during the time working with Mozilla.ai, the team decided to build a search tool for special collections using an embedding model. In the early stages of the work, taking a chatbot-style approach was considered. However, the team realized that the technology resources needed to support it and the risk of misinformation was too great to make a chatbot model sustainable for the intended users. Instead, the work focused on developing a platform that would enable semantic search and support natural language searching, with a few generative AI features.

In addition to determining the technical approach for the project, the team worked to decide on the product strategy for the platform. The team defined the product strategy for the first phase of the project by reviewing existing user research and determining that the largest opportunity was to make our special collections and unique materials easier to find for a non-researcher audience. Harvard Library has a significant number of digitized items available to anyone online, but user research indicates that these items can be difficult to find. Identifying this pain point ensured that the technology would directly address a real problem, rather than being an innovation for its own sake.

The team adopted a product strategy framework from a content strategy framework used on other digital projects.²⁷ A meeting was held very early in the process with project stakeholders and the technical team to determine this strategy. By bringing together this group, the team was able to discuss all the ideas people had for the platform and determine the focus for the first phase of work. This helped align the team around a common goal and provide a clear scope for the platform. It was decided shortly after this meeting to call the platform “Collections Explorer.”

The team went through a meeting to determine answers to the following questions:

- Who is the main audience for this product?
- What are the goals of the product?
- What is the value proposition?
- What content or features are needed to fulfill the goals?
- How will the success of the product be measured?

By answering these questions, the following product strategy was determined:

- The main audience is a person who wants to take their first step to easily explore special collections.
- The goals are to enable natural language searching of special collections, offer AI-generated related search queries and item-level summaries, and provide easy access to digital special collections.

- The value proposition is that the tool helps users take the first step in finding the library's unique materials.
- The team will validate that the goals are addressed by providing prompt examples, an explanation of how the system works, and AI-powered features (embedding model and LLM-powered).
- The success of the system will be measured using web analytics, a feedback survey, and usability testing.

Team Approach

With the product strategy defined, the core technical team for Collections Explorer was assembled. At the time the alpha version was released, the core technical team included two full-stack developers, one front-end developer, a user experience specialist, and the project manager. Since that time, the team has added another developer and technical lead. The project manager, user experience specialist, and two developer positions are term positions and focus solely on the project. The remaining permanent technical staff on the project also perform maintenance work on other systems as needed. It is important for the success of the project that the user experience specialist is embedded with the technical team as one of the project goals was to build a user-centered platform. The developers regularly observed user-testing sessions, and the user experience specialist worked especially closely with the team's front-end developer on user interface design.

This team adopted the Scrum framework to move quickly and to incorporate ongoing user research as it iterated on Collections Explorer. The team deviates from Scrum in that they function without a dedicated product owner; instead, the project manager functions as product owner, in collaboration with the team's engineering manager and the director of user experience. The project manager also functions as the Scrum master. The team organizes their work into two-week sprints, holds daily stand-ups, and conducts biweekly planning and retrospectives. The team faithfully records sprint demos that it shares with the library, as this helps promote transparency in the work and enables interested stakeholders to see the incremental progress toward larger deliverables and releases. By having a clear product strategy and using this approach to manage technical work, the project was able to release a testable alpha version of Collections Explorer in seven sprints (14 weeks). This version of Collections Explorer underwent usability and accessibility testing to inform the next phase of work.

In addition to creating a product strategy and assigning people to a core technical team to build Collections Explorer, the project planning process defined a communication strategy and brought together a library stakeholder group. The library stakeholder group consists of decision-makers in the library in the area of archives management, fine arts, digital collections, and special projects, with the full-time project manager as facilitator. This group meets biweekly to discuss project processes and provide input on prioritizing larger swaths of work to meet project goals. For example, this group provided input on what committees and what other colleagues the project manager should engage with to facilitate a successful launch of the project as a whole. The group also provided input on the metadata sources and created a definition of special and archival resources for the first release of Collections Explorer.

The other key participants in the project were two graduate and one undergraduate student workers, known as project insights interns. These students served as pilot testing participants for user studies, provided design feedback on interface prototypes, and served as a mini-focus group for smaller questions the team had about the student experience using AI for academic research

and finding digitized collection items online. They also helped create test queries the team used for quality assurance and evaluate which LLM would be best for the project's needs. They worked a few hours a week, usually asynchronously, and occasionally attended team meetings to give feedback. The project insights interns were one way that the team involved potential users in testing and feedback work, which was an essential component of the entire project.

EARLY EXPLORATORY RESEARCH AND DESIGN PROCESS

Early exploratory research that informed the Reimagining Discovery project charter began in 2022. This research was conducted for the discovery phase to inform the habits and pain points of users of special collections and digital discovery platforms. The team learned from surveys and interviews with special collections users that the experience of finding those materials was confusing and fragmented. After the launch of ChatGPT (OpenAI) in November 2022, the team began to hear anecdotally how ChatGPT was changing academic research habits. In fall 2023, the team conducted a survey of undergraduate and graduate students to understand how generative AI adoption was affecting their search habits and expectations.²⁸

In early 2024, the team conducted feedback sessions on the Talk with HOLLIS proof-of-concept and interviewed users about their expectations for library discovery in the age of AI. Participants included five undergraduate students and six graduate students. All participants were familiar with HOLLIS. Some participants were expert users of ChatGPT, while others had never used it. They were asked questions about ChatGPT use and then performed task-based usability testing of Talk with HOLLIS. The features they most desired from an AI-powered library discovery tool were summaries of each result and how it relates to their prompt, recommendations for related search terms, and transparency in which LLMs and data sources are used. The team also learned that users hold libraries to an even higher standard for trust and accuracy than they do commercial companies. Because of this, the team wanted to make sure that how AI is used for this search product is transparent and understandable by users, so the team created a "How This System Works" page (see Figure 1). This research was used to inform the product strategy and planning sessions with Mozilla.ai. See Appendix A for the full report.

As the team began to envision what the user interface for Collections Explorer would be, they used a user-centered design process and a design feedback framework. Prototypes for the interface were created using Figma (see prototype screenshots in Figures 2 and 3). This allowed the team to experiment with different designs and preview them with stakeholders before any coding happened. The user experience specialist on the project ran sessions and gathered feedback to inform the prototypes that were created. The team intentionally kept the initial versions very simple because they knew that they would do user testing to inform more detailed features like filters. The user experience specialist gathered feedback on the design prototypes from the technical team and stakeholders using a feedback approach called "Likes, Wishes, and What Ifs."²⁹ The designs were then given to the front-end developer to build the interface as part of the technical team's work.

Figure 1. Collections Explorer alpha “How This System Works” page.

How This System Works

Collections Explorer uses AI to search Harvard Library’s collections. Currently, it only searches special collections.

Question

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

Question

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

Question

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

Question

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

Question

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

Question

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

Figure 2. Collections Explorer alpha homepage.

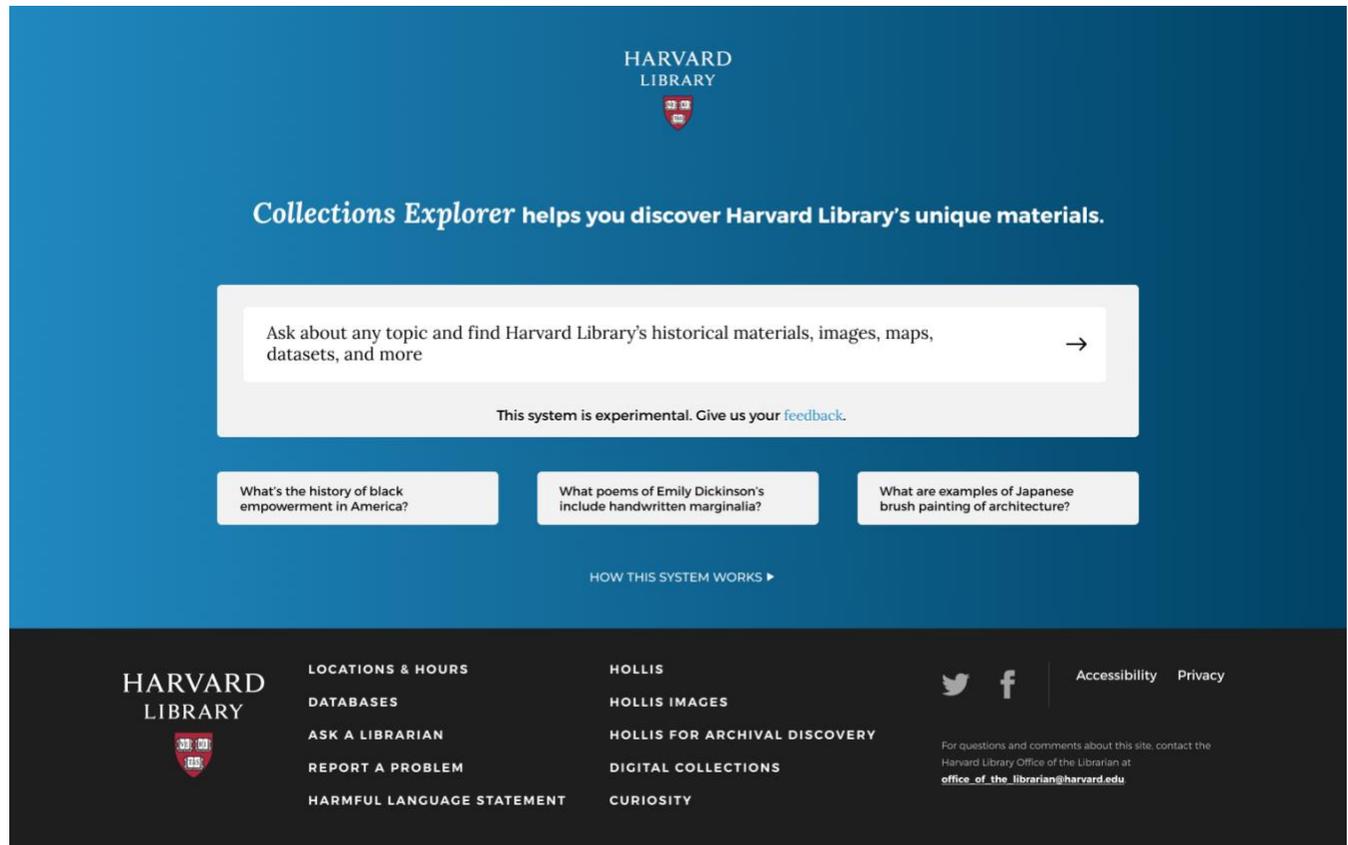


Figure 3. Collections Explorer alpha results page.

HARVARD LIBRARY

Collections Explorer

→

This system is experimental. We recommend you double-check the responses.

Include individual items in my results

DIGITAL COLLECTION

Contagion

May Include Digitized Materials

Why Did I Get This Result? ▼

DIGITAL COLLECTION ITEM

Hancock family papers

Contagion

May Include Digitized Materials

Why Did I Get This Result? ▼

FINDING AID

COVID-19 Response ephemera

May Include Digitized Materials

Why Did I Get This Result? ▼

FINDING AID ITEM

Hancock family papers

Letter by Secretary of International Medical Congress, drawing Attention to Address by Professor Pasteur, relative to Germ Theory in connection with Inoculation of Animals as Means of protecting them against Virulent Diseases

May Include Digitized Materials

Why Did I Get This Result? ▼

FINDING AID ITEM

Hancock family papers

The inhibitory action of Listerine. St. Louis: Lambert Pharmacal Co., 1908.

May Include Digitized Materials

Why Did I Get This Result? ▼

COLLECTION WEBPAGE

Contagion

Why Did I Get This Result? ▼

FINDING AID

COVID-19 Response ephemera

May Include Digitized Materials

Why Did I Get This Result? ▼

FINDING AID ITEM

Hancock family papers

Letter by Secretary of International Medical Congress, drawing Attention to Address by Professor Pasteur, relative to Germ Theory in connection with Inoculation of Animals as Means of protecting them against Virulent Diseases

May Include Digitized Materials

Why Did I Get This Result? ▼

More Results

YOU MIGHT ALSO TRY

How did public education campaigns of the 19th century intersect with germ theory?

What are techniques for antiseptic surgery?

How does animalcular theory relate to germ theory?

How did germ theory change healthcare?

What is the role of mass media in promoting germ theory about tuberculosis?

NEED ARTICLES OR BOOKS

View in HOLLIS

Give us your feedback.

HARVARD LIBRARY

LOCATIONS & HOURS

DATABASES

ASK A LIBRARIAN

REPORT A PROBLEM

HARMFUL LANGUAGE STATEMENT

HOLLIS

HOLLIS IMAGES

HOLLIS FOR ARCHIVAL DISCOVERY

DIGITAL COLLECTIONS

CURIOSITY

Accessibility Privacy

For questions and comments about this site, contact the Harvard Library Office of the Librarian at office_of_the_librarian@harvard.edu

TECHNICAL APPROACH AND AI FEATURES

While the Collections Explorer front end was being developed using React, the technical team concurrently focused on designing and implementing the back end for an alpha version. The back end supports core product goals outlined in the project strategy: enabling natural language searching, providing AI-generated related search queries, and offering item-level summaries. These capabilities are powered by two complementary AI approaches: semantic search and retrieval using an embedding model, and the generation of related queries and summaries using LLMs.

The retrieval system relies centrally on a component called a vector store, which manages the collection database and performs natural language query matching. Given the substantial size of Harvard Library's special collections metadata records—tens of gigabytes in aggregate—the team selected FAISS, a highly performant in-memory vector store optimized for efficient retrieval at scale. FAISS was integrated via LangChain's vector store interface, an increasingly popular open-source standard.

Critical to the functionality of the vector store is the embedding model. Embedding models convert textual information into numerical vectors, enabling semantic comparisons between texts based on their positions in a shared vector space. In this representation, semantically similar texts naturally cluster closer together. The technical team evaluated models using the Massive Text Embedding Benchmark (MTEB) scores—an aggregate performance measure across 15 retrieval benchmarks—and selected BAAI/bge-base-en-v1.5. This model was chosen for its strong performance and efficient resource usage, alongside its optimization for paragraph-length content. The team decided it would more thoroughly review top contender models in later phases for the production system.

Also critical to retrieval, a semantic application programming interface (API) was developed to handle user queries, converting natural language queries into embeddings using the selected model and retrieving the most relevant matches from the vector store. This semantic approach enables users to discover relevant materials without relying solely on exact keyword matches.

For the alpha release of Collections Explorer, the team processed several hundred finding aids, thousands of metadata records from digitized special collections items, and various collection descriptions sourced from the Harvard Library website. These metadata records and descriptions were embedded and used to create a static index within the vector store. Given the alpha's limited scope, the team deferred creating a dynamic metadata processing pipeline until the beta production phase.

To address additional features, the team also built an LLM API, enabling communication with various LLMs hosted on AWS Bedrock. An experimental evaluation was conducted with 12 different models, assessing their outputs based on accuracy, clarity, and appropriateness of tone. Student interns played a critical role in this evaluation process. Ultimately, Mistral Small was selected for generating explanatory item summaries (“Why Did I Get This Result?”), helping users contextualize matches by assessing the relevance of the resource to their query. Claude Haiku was chosen for additional features including suggested alternative searches (“You Might Also Try”) and facilitating secondary resource searches (“Need Books or Articles”) by converting natural language prompts into keyword queries executed against Harvard Library's Primo instance.³⁰

To ensure reliability and prevent inaccuracies, all LLM outputs were guided by carefully optimized system prompts. This methodical approach enabled the rapid development of an alpha version of

Collections Explorer in just 14 weeks, allowing early user testing and feedback integration ahead of the subsequent production-ready phase.

EVALUATIVE RESEARCH

Once the alpha was ready, the team conducted initial usability testing on Collections Explorer with casual users (our primary audience) and advanced researchers (our secondary audience) in fall 2024. The team received positive and constructive feedback from participants during the first round of user testing. The participants were encouraged by the potential of an AI-powered system for finding collection items at Harvard Library. Most notably, they were excited about the system's AI-generated features, like the “Why Did I Get This Result?” summary and the “You Might Also Try” search suggestions. The example prompts on the landing page (see Figure 4) were seen as approachable and friendly, making the system easier to navigate. There were three specific areas of focus for the test: search results presentation, filters, and generative AI features.

Participants demonstrated a strong preference for individual item-level results over collection-level results, finding items to be more useful. Collection records were perceived as time-intensive since they would require more manual review. Results relevancy was also a significant concern, with users expressing a desire for clear indication of result completeness. Participants encountered difficulty understanding special collections terminology, finding terms such as *finding aid* and *digital collection* unhelpful. They identified the most important metadata to include on the results page as access to the digitized version, year, collection name, and location (i.e., repository).

Participants wanted to be able to filter on digitization status, year, location, resource type, and language—preferences that aligned with existing web analytics. Users demonstrated a clear preference for immediate access to digitized materials, with advanced researchers valuing the option to expand results to include non-digitized items. Users wanted to be able to access help from librarians directly within the search interface.

The “Why Did I Get This Result?” feature was received very positively by users, who wanted to get immediate access to the information rather than have it behind a click. Advanced researchers had comments on the item descriptions, desiring a summary that allowed them to make a personal relevance assessment, instead of being told whether the item was relevant or not. The “You Might Also Try” recommendation section was also well-received, but participants wanted suggested prompts that were only relevant to Harvard Library’s collections. Participants thought it was important for the interface to clearly denote where AI was used to generate content.

The usability testing of the prototype provided critical insights that are shaping the next phase of development. See Appendix C for a complete list of changes made based on usability testing. Participants expressed a strong preference for individual item-level results over collection-level descriptions, citing efficiency and relevance. Users also valued transparency about AI functionality, requesting clearer labeling and visible explanations without needing to click. Advanced researchers preferred AI-generated relevance statements that focused on objective item descriptions rather than inferred user intent. Users appreciated the suggested searches but wanted them more closely tailored to the Harvard Library context. Feedback also highlighted the importance of clarity around digitization. Users wanted to know whether an item is digitized and accessible without navigating through multiple pages. For the beta release, the team worked to improve metadata visibility and interface design to prioritize access status and enable filtering by digitization, year, repository, and language.

See Figures 4 and 5 for screenshots of the improved beta design, and see Appendix B for a full report of the usability testing.

Figure 4. Collections Explorer beta homepage.

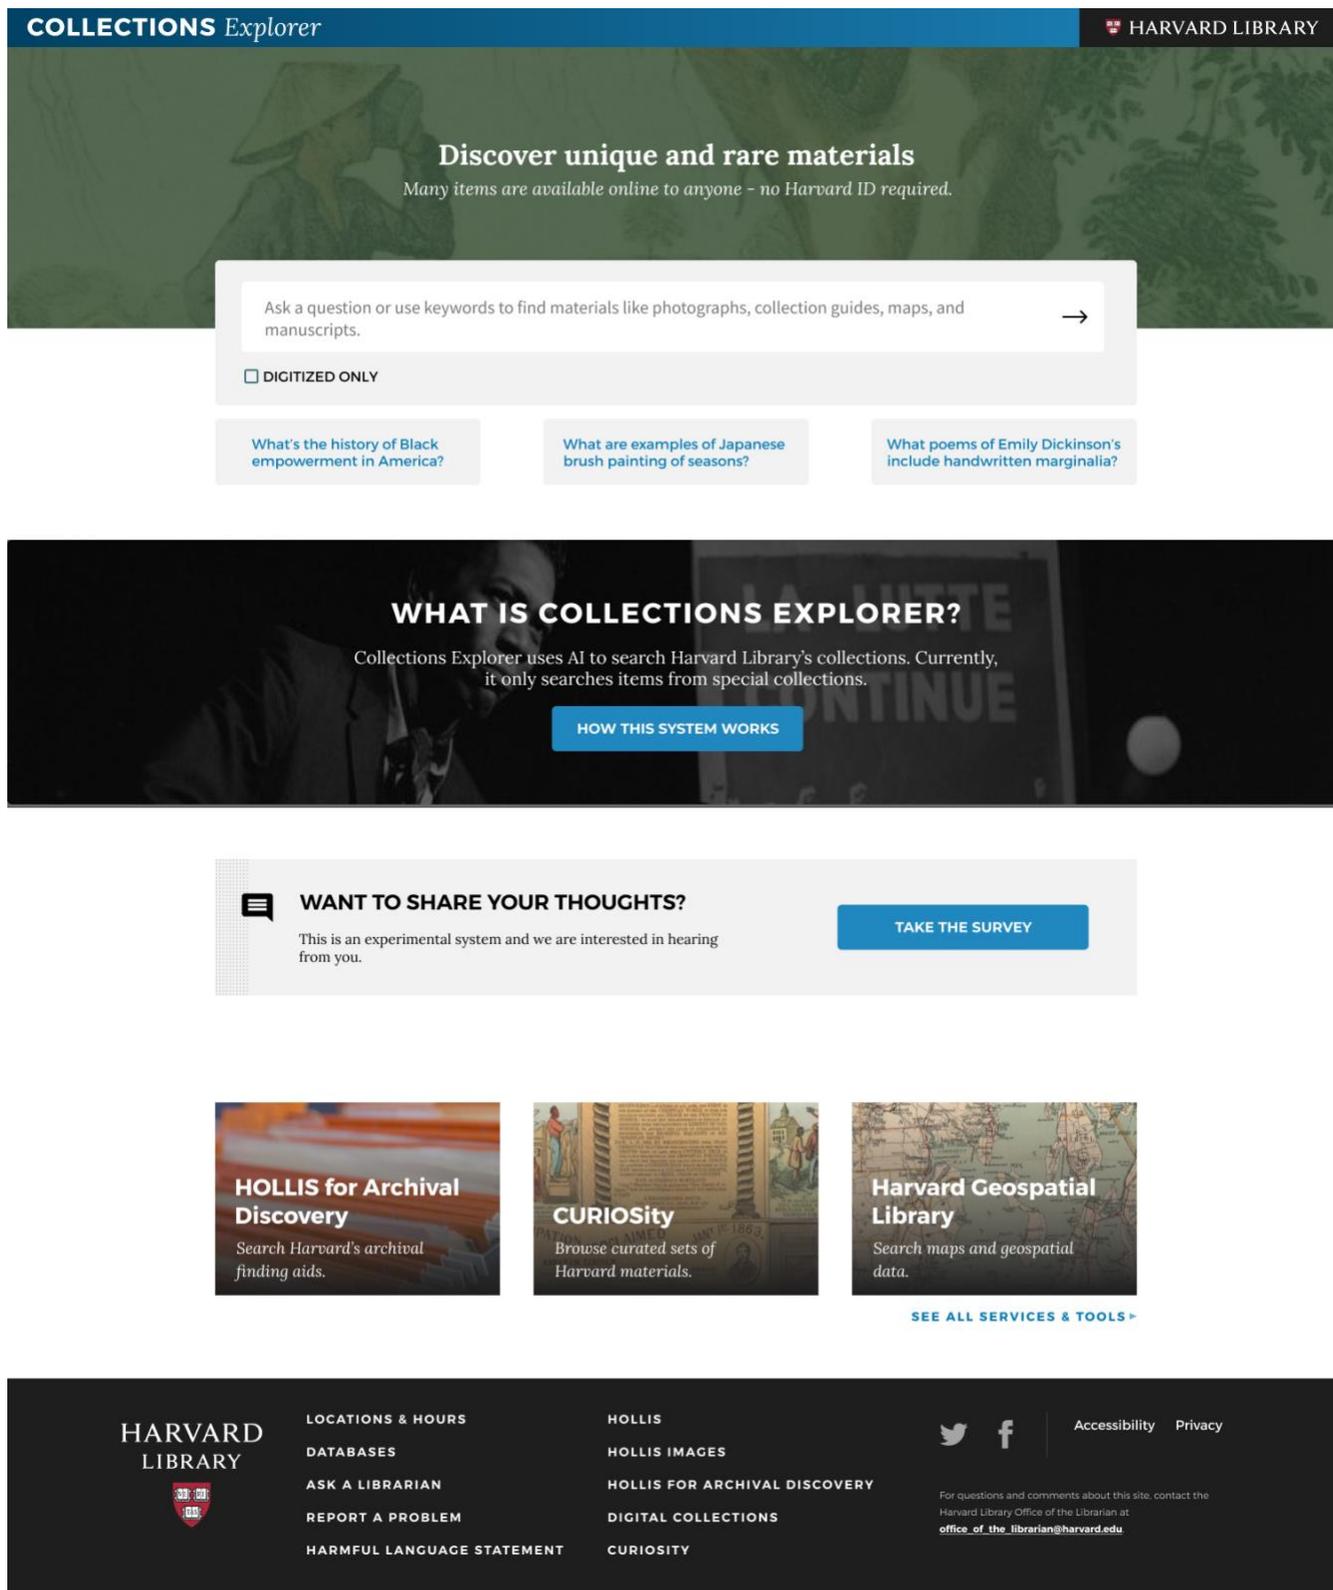


Figure 5. Collections Explorer beta results page.

The screenshot shows the Collections Explorer interface. At the top, there is a search bar with the query "What does Harvard have on the history of germ theory development?". Below the search bar, there are filter options: "DIGITIZED ONLY", "DATE (1920-1980)", "LOCATION", "RESOURCE TYPE (2)", and "LANGUAGE". A "COLLAPSE FILTERS" button is also present. The results section shows "10 OF 200 RESULTS" and a "HIDE SUMMARIES" toggle. The results are displayed as a list of items, each with a "Name" field, a "Date" field, and a brief description. On the right side, there is a "YOU MIGHT ALSO TRY" section with several suggested queries, and a "NEED ARTICLES OR BOOKS" section with a "TRY HOLLIS CATALOG" button. At the bottom, there is a footer with the Harvard Library logo, navigation links, and social media icons.

GETTING TO LAUNCH AND FUTURE PLANS

Prior to launch, the technical team built a robust ingest pipeline to manage the flow of metadata records from disparate sources. This pipeline allows for continual updating and adding to the search index, as Harvard Library creates new and updates metadata records in databases of record, like ArchivesSpace or Alma. The team also evaluated multiple vector databases based on requirements, some of which came out of usability testing, like the ability to support both semantic search and metadata filtering. Ultimately, the team implemented Elasticsearch, which is capable of supporting tens of millions of vectors.

Additionally, the technical team partnered with librarians to evaluate which embedding model to use for production. The librarians evaluated the relevancy of results that three different models returned for multiple queries, including foreign language ones. Ultimately, the team chose Cohere v3 multi-lingual, which performed the best in terms of precision and focused on ensuring that embedded texts are split into the right size for the embedding model. The team will repeat the evaluation of LLMs available for generating the summary of a result. This feature replaced the “Why Did I Get This Result?” and became AI-generated summaries due to the feedback gathered during usability testing. Lastly, as a result of the usability testing, the entire user interface was redesigned, from the homepage to the layout of the results to the removal of the collection-level toggle and the addition of metadata filters.

Collections Explorer is being designed with flexibility in mind: AI models can be swapped as the landscape evolves, and hybrid search (combining semantic and keyword approaches) will ensure relevance and precision. In order to launch Collections Explorer, the team also fine-tuned relevancy in the search engine and performed a thorough security assessment. The interface will continue to evolve through more evaluative research conducted during the project. Post-launch, the next priority for Collections Explorer will be to expand the scope of content and build features to enable the discovery of digital images, full-text documents, and born-digital materials.

CONCLUSION

In September 2025, Collections Explorer (beta MVP) launched for the public.³¹ Collections Explorer marks a significant advancement in library discovery, showcasing how AI technologies can be thoughtfully integrated to meet evolving user needs. By prioritizing human-centered design and careful technological evaluation, the project reflects librarians’ commitment to empowering users and providing trustworthy information.

The development process demonstrated the value of iterative design, user research, and technical adaptability. From initial concept to beta release, the project balanced user requirements with organizational goals, offering insights into how researchers interact with special collections in an AI-enhanced environment. Moving forward, Collections Explorer will continue to explore ways AI can make special collections more accessible and engaging, ultimately transforming how users discover and connect with Harvard Library’s extensive resources.

STATEMENT OF AI USE

Claude Haiku 3.5 and ChatGPT 4o were used by the writers to help with rough drafts of paragraphs (based on the authors’ notes or project reports), brainstorm ideas for how to find sources, suggest transition sentences to connect sections from one to the other, and suggest alternative wording, especially to avoid the use of “we” and “us.” Below are five example prompts with outputs that were used by the authors during the writing process.

Prompt 1: Rough Draft of Introduction Paragraph

User: I am going to share the draft text of an article that I am co-writing. Please read it and let me know when you are ready to help with an adjustment. In the literature review and in the technical approach & AI features section, we mention the importance of having a human-in-the-loop for reviewing LLM outputs which adds to the trustworthiness and library-centric nature of the collections explorer platform. Can you draft a paragraph that I can add to the introduction that sets up this topic. Please follow the writing style in the introduction as it is currently and do not use em-dashes.

Claude: Here's a draft paragraph for the introduction that highlights the human-in-the-loop approach:

The development of Collections Explorer is grounded in a commitment to human-centered AI that prioritizes user empowerment and professional librarian expertise. While leveraging advanced AI technologies, the project maintains a critical human oversight process where librarians actively review and guide technological implementations. This approach ensures that the platform not only utilizes cutting-edge AI capabilities but also reflects the core values of library science: providing reliable, contextual, and trustworthy information access. By keeping librarians integral to the technological development process, the team seeks to create an AI-enabled discovery tool that enhances rather than replaces human research expertise, maintaining the library's role as a thoughtful and responsible knowledge partner.

Prompt 2: Rough Draft of Usability Testing Paragraph

User: I am writing an article about usability testing that I conducted last year. Please provide a paragraph, suitable for an academic article, that describes these findings. Please avoid adverbs and overly flowery language.

Participants prefer seeing individual items in search results rather than collection-level descriptions because they want direct access to individual items. They avoid the collection level records because they know it will be time consuming to review everything included.

Participants noted inconsistencies in the relevancy of the results and wanted better ranking within the search results. Participants want to quickly understand when HARVARD doesn't have items relevant to their query.

Participants didn't understand the definitions or differences between collection container terms (e.g., finding aid, digital collection, collection webpage). These terms are not helpful.

Participants desired information about digitization status because they wanted to understand the accessibility of items in the collection without navigating multiple pages or interfaces.

Participants want to quickly understand if an item is relevant before viewing it because they want to evaluate its usefulness before committing more time. The most important metadata fields are year and container collection.

Claude: Here's a draft paragraph for an academic article:

The usability testing revealed key user preferences in collection search interfaces. Participants demonstrated a strong preference for individual item-level results over collection-level descriptions, seeking direct access to specific materials and avoiding time-intensive record reviews. Search result relevancy and ranking emerged as significant concerns, with users

expressing a desire for clear indication of result completeness. Participants encountered difficulty comprehending collection container terminology, finding terms such as “finding aid” and “digital collection” opaque and unhelpful. Metadata transparency was crucial, with users valuing immediate access to digitization status and core descriptive information like year and collection container. These findings underscore the importance of user-centric design in archival search interfaces, emphasizing the need for clear, direct access to research materials.

Prompt 3: Transition Sentences

User: I’m going to paste the introduction to the peer-reviewed article I’m writing. Can you suggest a few sentences to add to the end of it and transition into the Literature Review? [Pasted introduction text and beginning of Lit Review section]

ChatGPT: Here are a few transition options you can use to smoothly lead into the Literature Review section. These build on your introduction and set up the rationale for engaging with existing research:

Option 1: Emphasizing the need for a conceptual foundation

To situate our work within the broader context of library discovery, user-centered design, and AI applications in cultural heritage, we turn now to the existing literature that informs the conceptual and technical foundation of this project.

Option 2: Highlighting the importance of precedent and related work

Before delving into the design and development process in more detail, it is important to review prior work in the fields of library discovery, archival access, and AI-driven user interfaces. This literature provides valuable insight into both the challenges we faced and the innovations we sought to build upon.

Option 3: Framing the literature as informing strategy and methodology

The strategy, scope, and technical direction of the Collections Explorer platform were shaped not only by user research, but also by prior scholarship on discovery systems, special collections access, and AI in library contexts. The following literature review outlines key contributions that influenced our approach.

Option 4: Signaling a contribution to ongoing scholarly conversations

Our work contributes to an ongoing conversation about how libraries can adapt discovery systems to meet the needs of diverse users in an AI-enabled landscape. The following literature review situates this project within that evolving dialogue.

Prompt 4: Remove “We” and “Us”

User: Can you remove our and we from this sentence and make it clearer?

With this in mind, a critical component of this work is to ensure that in our adoption of semantic search and AI-powered features, we build a trustworthy and ethical library discovery tool.

ChatGPT: Certainly. Here’s a clearer, more objective version of the sentence without using “we.” With this in mind, a critical component of the work is ensuring that the adoption of semantic search and AI-powered features results in a trustworthy and ethical library discovery tool.

Or, for slightly smoother flow:

A critical component of this work is ensuring that the integration of semantic search and AI-powered features supports the development of a trustworthy and ethical library discovery tool.

Prompt 4: Describe the History of Card Catalogues

User: Is there a history online of the evolution of card catalogues? If so, please reference it and summarize it.

ChatGPT: Yes, the history of the card catalog in U.S. libraries is well-documented, and its development played a crucial role in library organization. Here's a brief overview:

1. **Early Library Cataloging (Pre-19th Century)**
Before the adoption of card catalogs, libraries primarily used printed book catalogs or handwritten manuscript lists to organize their collections. These methods were prevalent until the late 19th century.
2. **The Birth of the Card Catalog (Mid-to-Late 19th Century)**
The shift toward card catalogs began in the mid-19th century, influenced by the need for more flexible and updateable systems. The Library of Congress adopted the use of cards in 1860 to manage its growing collection. However, it wasn't until the 1876 American Library Association (ALA) conference that the card catalog was formally endorsed as a practical cataloging system.
3. **Melvil Dewey's Influence (1876 Onward)**
Melvil Dewey, a major force in library standardization, promoted the card catalog system in the late 19th century. His Library Bureau, founded in 1876, began mass-producing standardized catalog cards and card cabinets, which helped libraries across the country transition to this system.
4. **Widespread Adoption (1880s–1920s)**
By the 1880s and early 1900s, most U.S. libraries, including the New York Public Library and Harvard University, had embraced card catalogs. The Library of Congress started selling preprinted catalog cards in 1901, allowing libraries nationwide to adopt a uniform cataloging system.
5. **The Dominance of Card Catalogs (20th Century)**
Card catalogs remained the primary cataloging system for U.S. libraries throughout the 20th century. Libraries expanded their catalogs by using typewriters and later computer-printed cards.
6. **The Decline and Shift to Digital (1980s–Present)**
The introduction of computerized library catalogs (OPACs) in the 1980s and 1990s led to the gradual decline of card catalogs. By the early 2000s, most major libraries had retired their physical card catalogs in favor of digital databases.

This history reflects a gradual evolution driven by the need for flexibility, efficiency, and standardization in library cataloging.

Sources [these all pointed to the same article online]: Coyle, Karen. 2016. "The Evolving Catalog." *American Libraries Magazine*, January 4.

<https://americanlibrariesmagazine.org/2016/01/04/cataloging-evolves/>

APPENDICES

Appendix A: Report on Talk with HOLLIS and AI Discovery Interviews

In January 2024, interviews and usability testing were conducted with 11 Harvard students to understand their usage of generative AI for academic work, as well as to solicit feedback on the usability of Talk with HOLLIS. Talk with HOLLIS is an experimental generative AI chatbot that retrieves results from HOLLIS via natural language prompting.

Methodology and Participants

Through in-depth interviews, the team sought to learn the following:

- How do students use ChatGPT for academic work?
- Why do students use Google versus HOLLIS versus ChatGPT?
- How is the overall usability of Talk with HOLLIS (TWH)?
- What feature ideas do students have for future library chatbots?

Participants included five undergraduate students and six graduate students. All participants were familiar with HOLLIS. Some participants were expert users of ChatGPT, while others had never used it. They were asked questions about ChatGPT use and then tested Talk with HOLLIS.

Interview Insights on Student Use of HOLLIS and ChatGPT

- HOLLIS is a secondary step in the research process for students. They go to HOLLIS with known item searches or intentional search goals. Preliminary topic exploration, when they're at the idea stage, is done elsewhere, like Google or ChatGPT.
- Participants expect different functionality from a chatbot than they would from a search engine. They want to engage in conversation and think of ChatGPT as a partner rather than a source of information.
- ChatGPT is used for quick answers when the source or context isn't as important to students.
- Students use ChatGPT as a learning partner, writing editor, translation tool, and to find "lost citations," but they don't trust it for everything. They say that they have to "fact-check" the output.
- Most students know that ChatGPT won't connect them to Harvard resources.

Usability Insights on Talk with HOLLIS

- Students are excited that the libraries are experimenting with generative AI; they have high expectations and specific ideas of what this tool could do.
- The chatbot is less overwhelming than HOLLIS.
- Participants think the Talk with HOLLIS chatbot is too similar to HOLLIS. They want a library chatbot to be a supplementary tool, not a replacement for HOLLIS. The features that they most want are the following:
 - The ability to ask follow-up questions and refine the output through follow-up prompts.
 - An explanation of why these specific results are in the output.
 - Recommendations on related materials.
 - A distinct identity that is separate from what HOLLIS already does.
 - The ability to seamlessly connect with other library services (i.e., text a call number, request pick up, make a research appointment, make an archives appointment, etc.).
- Students expect tools like this to be on brand with other library sites.

- Students want to understand how the chatbot was trained and what data sources are used. They want to know the limitations of the chatbot.
- Harvard tools are held to a higher standard regarding reliability and trust.

Appendix B: Report on Collections Explorer Alpha Usability Testing

In October 2024, the User Research Center conducted user testing on Collections Explorer, an experimental tool built by Harvard Library to enable easier searching of Harvard's unique special collections. The system uses semantic search technology to understand the meaning of your question and not just the individual words. Feedback from this phase will inform future iterations of Collections Explorer to enhance the user experience.

Methodology and Participants

We hosted think-aloud usability testing with 11 participants.

Recommendations and Insights

The participants were encouraged by the potential of an AI-powered system for finding collection items at Harvard. Most notably, they were excited about the system's AI-generated features, like the "Why Did I Get This Result" summary and the "You Might Also Try" search suggestions. The example queries were seen as approachable and friendly, making the system easier to navigate. Overall, participants found the tool easy to use and saw potential for broader promotion to diverse audiences, including students and educators.

Homepage Insights

Content on the homepage needs to be expanded so the audience for the search is clearer (for casual users) and is clearer what is included here (for advanced users).

1. Participants believe that public members would not have easy access to collection items.
2. Advanced users think this search is too simplistic for their needs.
3. Participants are delighted by the variety of example questions relating to diverse perspectives.

Search Results Insights

4. Participants prefer seeing individual items in search results rather than collection-level descriptions because they want direct access to individual items. They avoid the collection-level records because they know it will be time consuming to review everything included.
5. Participants noted inconsistencies in the relevancy of the results and wanted better ranking within the search results. Participants want to quickly understand when Harvard doesn't have items relevant to their query.
6. Participants didn't understand the definitions or differences between collection container terms (e.g., finding aid, digital collection, collection webpage). These terms are not helpful.
7. Participants desired information about digitization status because they wanted to understand the accessibility of items in the collection without navigating multiple pages or interfaces.
8. Participants want to quickly understand if an item is relevant before viewing it because they want to evaluate its usefulness before committing more time. The most important metadata fields are year and container collection.

Filter Insights

9. Participants expect more specific filtering and toggle functionality to refine searches quickly using digitized items and relevant metadata.
10. The most desired filters are digitized/non-digitized, year, repository, resource type, and language (aligns with analytics from HOLLIS).
11. Participants did not understand the language associated with the individual results toggle and often didn't notice it.
12. All participants would prefer to see items they can immediately access in the results. Advanced researchers who want a more complete view can use a filter to include non-digitized items.
13. Participants expect access to Chat with a Librarian from Collections Explorer.

AI-Related Insights

14. Non-relevant search results were often related to additional semantic terms relating to Harvard. Update the system prompt to ignore these kinds of phrases (i.e., "in Yenching Library")
15. Participants expected the information under "Why Did I Get This Result?" to be immediately available without a click.
16. Advanced researcher participants expect the language under "Why Did I Get This Result?" to describe the item without a relevancy statement. They want to decide if something is relevant for themselves and bristle at the term's use.
17. Participants like the "You Might Also Try" section. However, they wish the language of the suggested prompts referred more to Harvard Library collections than general search terminology.
18. Participants often speculated on what features might be AI-powered. They would like to know this within the interface. This also supports the project goal of transparency in AI use.

Appendix C: Changes Made Based on Usability Testing

- Prioritizing item-level results in the search results.
- Displaying item-level and collection-level records intermingled in the results list. Originally, the platform showed collections by default and users had the option to opt into including item-level records using a toggle. This proved confusing for users.
- Displaying the record summary by default instead of behind a click; changed the wording of the summary so as not to make relevancy statements; changed the title of the content area to simply "About This Item."
- Indicating what content on the page is created with generative AI.
- Adding the "digitized only" checkbox to the homepage; added this and additional filters (year, location, resource type, and language) to the results page.
- Updating the homepage to provide more context for the platform and how it is distinct from other library discovery platforms.

ENDNOTES

- ¹ Carolyn Caizzi, Amy Deschenes, and Stu Snyderman, "Reimagining Discovery: Transforming Access to Collections with Artificial Intelligence-Driven Exploration," presented at the Coalition for Networked Information Spring 2024 Membership Meeting, <https://www.cni.org/topics/special-collections/reimagining-discovery-transforming-access-to-collections-with-artificial-intelligence-driven-exploration>.
- ² Stuart Snyderman and Martha Whitehead, "From Building a First-Generation Digital Library Infrastructure to Reimagining Discovery," *International Journal of Digital Curation* 19, no. 1 (2025): 13, <https://doi.org/10.2218/ijdc.v19i1.1068>.
- ³ Donald Norman and Stephen Draper, *User Centered System Design: New Perspectives on Human-Computer Interaction* (Hillsdale, NJ: Lawrence Erlbaum Associates, 1986).
- ⁴ Jesse James Garrett, *The Elements of User Experience: User-Centered Design for the Web* (Westlake Village, CA: New Riders, 2002).
- ⁵ Laura Klein, *Build Better Products: A Modern Approach to Building Successful User-Centered Products* (New York, NY: Rosenfeld Media, 2016).
- ⁶ Michael Stuart Freeman, "Pen, Ink, Keys, and Cards: Some Reflections on Library Technology," *College & Research Libraries* 52, no. 4 (1991): 328–35, <https://doi.org/10.5860/crl.52.04.328>; Karen Coyle, "The Evolving Catalog," *American Libraries Magazine*, January 4, 2016, <https://americanlibrariesmagazine.org/2016/01/04/cataloging-evolves/>.
- ⁷ Patrick Lo, "The Transition from the Library OPAC to the Library Resources Portal: Difficulties and New Approaches in User Education for the Academic Librarians and Their Changing Roles," *International Journal of Technology, Knowledge, and Society* 4, no. 2 (2008): 39–50, <https://doi.org/10.18848/1832-3669/CGP/v04i02/55848>.
- ⁸ Lo, "The Transition from the Library OPAC to the Library Resources Portal"; Marshall Breeding, "Chapter 1: Discovery Product Functionality," *Library Technology Reports* 50, no. 1 (n.d.): 5–32.
- ⁹ Karen Markey, "The Online Library Catalog: Paradise Lost and Paradise Regained?," *D-Lib Magazine* 13, nos. 1–2 (2007), <https://doi.org/10.1045/january2007-markey>; Kristen Antelman et al., "Toward a Twenty-First Century Catalog," *Information Technology and Libraries* 25, no. 3 (2006): 128–39, <https://doi.org/10.6017/ital.v25i3.3342>.
- ¹⁰ Karen Calhoun, *The Changing Nature of the Catalog and Its Integration with Other Discovery Tools* (Washington, DC: Library of Congress, 2006), <https://hdl.handle.net/1813/2670>.
- ¹¹ Betsy Van Der Veer Martens, "Discovery, Retrieval," in *Keywords In and Out of Context*, 111–26 (Cham, Switzerland: Springer, 2023), https://doi.org/10.1007/978-3-031-32530-4_8.
- ¹² Lorcan Dempsey, "Library Discovery Directions," foreword to *Resource Discovery for the Twenty-First Century Library: Case Studies and Perspectives on the Role of IT in User Engagement and Empowerment*, edited by Simon McLeish (London, UK: Facet Publishing, 2020), <https://doi.org/10.25333/P8J4-1411>; Richard Wisneski, "I Can't Get No Satis-Searching: Reassessing Discovery Layers in Academic Libraries Journal of Web Librarianship," *Journal of Web Librarianship* 18, no. 1 (2024): 1–14, <https://doi.org/10.1080/19322909.2024.2326687>.
- ¹³ Charles Hildreth, "The Use and Understanding of Keyword Searching in a University Online Catalog," *Information Technology and Libraries* 16, no. 2 (1997): 52–62.
- ¹⁴ Aaron Tay, "Boolean vs Keyword/Lexical Search vs Semantic—Keeping Things Straight," *Medium* (blog), November 26, 2023, <https://aarontay.medium.com/boolean-vs-keyword-lexical-search-vs-semantic-keeping-things-straight-95eb503b48f5>.
- ¹⁵ Dempsey, "Library Discovery Directions."
- ¹⁶ Donald A. Norman, *The Design of Everyday Things*, revised and expanded edition (Cambridge, MA: MIT Press, 2013).

- ¹⁷ Erika Hall, *Just Enough Research* (New York, NY: A Book Apart, 2013).
- ¹⁸ Steven Portigal, *Interviewing Users: How to Uncover Compelling Insights* (New York, NY: Rosenfeld Media, 2013); Chauncey Wilson, *Interview Techniques for UX Practitioners: A User-Centered Design Method* (San Francisco, CA: Morgan Kaufmann, 2013).
- ¹⁹ William Albert and Thomas Tullis, *Measuring the User Experience* (San Francisco, CA: Morgan Kaufmann, 2022), <https://doi.org/10.1016/C2018-0-00693-3>; Hall, *Just Enough Research*; Mike Kuniavsky et al., *Observing the User Experience: A Practitioner's Guide to User Research* (San Francisco, CA: Morgan Kaufmann, 2012).
- ²⁰ Matt Enis, "AI on the Horizon: Several Major Companies—including Clarivate, EBSCO, OCLC, and OverDrive—Are Introducing New AI Tools and Features That Will Impact Researchers, Patrons, and Library Work Directly," *Library Journal News*, August 5, 2024, <https://www.libraryjournal.com/story/ai-on-the-horizon/>.
- ²¹ Dempsey, "Library Discovery Directions."
- ²² Trey Grainger et al., *AI-Powered Search* (New York, NY: Simon and Schuster, 2025); Daniel Jurafsky and James H. Martin, *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, third edition (Upper Saddle River, NJ: Prentice Hall, 2025), <https://web.stanford.edu/~jurafsky/slp3/ed3book.pdf>; Hannah Bast et al., "Semantic Search on Text and Knowledge Bases," *Foundations and Trends® in Information Retrieval* 10, nos. 2–3 (2016): 119–271, <https://doi.org/10.1561/1500000032>.
- ²³ Mohammad Amir Khusru Akhtar et al., "Challenges in Developing Ethical and Socially Responsible Explainable AI," in *Towards Ethical and Socially Responsible Explainable AI*, vol. 551 (Cham, Switzerland: Springer, 2024), https://doi.org/10.1007/978-3-031-66489-2_3; Ben Shneiderman, *Human-Centered AI* (Oxford, UK: Oxford University Press, 2022), <https://academic.oup.com/book/41126>.
- ²⁴ Michael Gorman, *Our Enduring Values Revisited: Librarianship in an Ever-Changing World* (Chicago, IL: ALA Editions, 2015).
- ²⁵ Tita Alissa Bach et al., "A Systematic Literature Review of User Trust in AI-Enabled Systems: An HCI Perspective," *International Journal of Human-Computer Interaction* 40, no. 5 (2024): 1251–66, <https://doi.org/10.1080/10447318.2022.2138826>; John D. Lee and Neville Moray, "Trust, Self-Confidence, and Operators' Adaptation to Automation," *International Journal of Human-Computer Studies* 40, no. 1 (1994): 153–84, <https://doi.org/10.1006/ijhc.1994.1007>; Susan Kathleen Lippert, "An Exploratory Study into the Relevance of Trust in the Context of Information Systems Technology," *ProQuest Dissertations and Theses* (2001).
- ²⁶ Akhtar et al., "Challenges in Developing Ethical and Socially Responsible Explainable AI."
- ²⁷ Amy Hannah Deschenes, Meg McMahon, and Claire DeMarco, "Designing Content for Impact: A Content Strategy Approach for Libraries & Beyond," *Weave: Journal of Library User Experience* 8, no. 2 (2025), <https://doi.org/10.3998/weaveux.6954>.
- ²⁸ Amy Deschenes and Meg McMahon, "A Survey on Student Use of Generative AI Chatbots for Academic Research," *Evidence Based Library and Information Practice* 19, no. 2 (2024): 2–22, <https://doi.org/10.18438/eblip30512>.
- ²⁹ Richard Lemarchand, *A Playful Production Process: For Game Designers (and Everyone)* (Cambridge, MA: MIT Press, 2021).
- ³⁰ J. J. Chen, "Reimagining Harvard Collections: AI-Enhanced Discovery," presentation at the Code4Lib conference, Princeton, New Jersey, March 12, 2025, <https://2025.code4lib.org/talks/reimagining-harvard-collections-ai-enhanced-discovery>.
- ³¹ Stu Snyderman, "Launch of Collections Explorer (MVP)," *The Circ Blog*, Harvard Library, September 19, 2025, <https://circblog.library.harvard.edu/2025/09/19/launch-of-collections-explorer-mvp/>.