ARTICLE

Transparent AI Tools for Libraries

Creating the OneSearch Assistant

Lucy Campbell and Keven Jeffery

ABSTRACT

The rapid integration of generative artificial intelligence (GenAI) into academic library discovery systems offers new opportunities but also introduces challenges, as vendor tools streamline natural language searching while obscuring search logic. This article presents a case study of a librarian-designed AI discovery tool that is focused on transparency and pedagogy. The assistant converts natural language queries into structured, editable Boolean search strings; surfaces search logic; and retrieves results across library holdings. Developed through iterative prompt engineering, this tool demonstrates how GenAI can support research practices without sacrificing rigor. Findings from qualitative feedback and quantitative assessment indicate that the tool improves the consistency and quality of Boolean query construction while fostering user engagement with search as a reflective process. This case study argues that libraries can play a critical role in shaping AI discovery tools that align with information literacy goals and institutional scholarly standards. By foregrounding transparency and user agency, librarian-led AI design positions GenAI as a form of instructional scaffolding rather than a replacement for expertise.

INTRODUCTION

Since the release of ChatGPT in November 2022, generative artificial intelligence (GenAI) has been rapidly adopted in academic libraries. Both services and workflows have been altered, but perhaps the most prolific area changed has been discovery systems. The spread of AI tools in library discovery has been swift. Vendors have rapidly released tools designed for specific databases and the data they hold. The predominance of AI tools has resulted in challenges for library and information professionals, particularly in academic research libraries. For example, these tools frequently hide the underlying search logic from the user, creating an opacity that misses opportunities for learning about advanced search. Results are siloed to specific platforms instead of surfacing all content owned or subscribed to by a library. These factors create missed opportunities for the user in terms of information literacy and reproducibility across databases.

The vast array of AI discovery tools creates a problem for research libraries. While acknowledging the benefits of AI, the authors question how to create tools that support discovery without undermining critical information literacy skills or restricting content. This case study outlines a project at San Diego State University (SDSU) Library to build an AI Discovery Assistant that can surface search logic, retrieve results from across holdings, and provide learning opportunities for users. The OneSearch Assistant is a librarian-designed, GenAI-enhanced discovery aid that foregrounds transparency and pedagogy, while working with existing metadata indexes. It uses multiple AI prompts (agentic AI) to redirect known items to the discovery system and converts queries into structured, editable Boolean strings.

About the Authors

Lucy Campbell (lgcampbell@sdsu.edu; corresponding author) is Electronic and Continuing Resources Librarian, San Diego State University. **Keven Jeffery** (kjeffery@sdsu.edu) is Digital Initiatives Librarian, San Diego State University.

Submitted: 21 April 2026. Accepted for Publication: 29 May 2026. Published: 21 September 2026.

https://doi.org/10.5860/ital.v45i3.17697

LITERATURE REVIEW

Academic library interfaces are increasingly adopting machine learning and GenAI to support search and discovery. Library vendors routinely integrate AI-powered research assistants directly into database platforms, positioning them as tools that speed up the search process and lower barriers to entry for novice researchers. These systems promise increased usability and a reduced cognitive load, particularly for undergraduate students who are accustomed to web searching.[1]

At the same time, vendor-designed AI discovery tools are largely opaque. How queries are ranked or filtered is rarely visible to users. Instead, these tools are presented as conversational aids that pop up and prompt users to ask for help. For example, the Web of Science SmartSearch offers suggestions along the way and assures users, "There's no need to master complex Boolean syntax, just type your query or keywords naturally and SmartSearch constructs the Boolean query behind the scenes."[2] These tools are also often restricted to vendor-owned content as opposed to a library's full holdings. The ProQuest EbookCentral Research Assistant displays an AI panel that is restricted to individual book chapters that must be owned by the subscribing library, while JSTOR's AI Research Tool recommends related materials only from within JSTOR's journal articles and books.[3]

These features create tension between accessibility and agency. On one hand, AI-enhanced discovery lowers barriers for users lacking familiarity with search logic and discipline-specific vocabularies. On the other hand, black-box AI undermines user agency by discouraging refinement and replication of search strategies. This loss of transparency is particularly consequential for advanced research, such as systematic reviews, where repeatability is essential. As Wang et al. note, Boolean search strategies are still foundational because they can be documented and reused.[4]

Tay has suggested Boolean alone is of little benefit in search environments built on AI probabilistic ranking and relevance.[5] However, Boolean logic has long been central to librarian-mediated search and instruction. The three operators "AND," "OR," and "NOT," serve as mechanisms for translating research questions into structured queries that can be executed for maximum retrieval. Boolean remains a fundamental logic that underpins all search, even as linking, ranking, and machine learning have improved modern interfaces.[6] In academic databases, advanced search options typically begin with the building of a Boolean search string and then allow for additional filtering (e.g., by resource type, publication date range, or peer review). Boolean logic is still taught by librarians and instructors looking to ensure students understand how to structure a search string. While this may change in the future, in the near term the vast number of controlled databases continues to support Boolean search.

Nevertheless, studies of undergraduate searching show students struggling to apply Boolean logic effectively. Lowe et al.'s comparative study found that natural language queries (NLQs) often performed as well as Boolean searches for introductory research, leading the authors to question the value of teaching Boolean as an early research skill.[7] Rather than abandon Boolean, they argue for reframing it as an advanced practice for rigorous research. These practices are particularly critical for systematic reviews, where precision and recall are paramount. Studies of published systematic reviews reveal high rates of error in reported search strategies, including logical inconsistencies and inappropriate truncation.[8] These findings underscore the importance of tools that make search logic visible and allow for refinement. They also align closely with the emphasis on search as an interactive, reflective practice as described by the Association of College and Research Libraries' (ACRL) *Framework for Information Literacy*.[9]

Indeed, many structured frameworks seek to position GenAI not as a replacement for human expertise but as instructional scaffolding that supports learning. The American Library Association's draft *Guidance on the Use of Artificial Intelligence in Libraries* (April 2026) emphasizes AI as a complement, rather than a replacement, for human intelligence and preserves the "right of patrons to human assistance."[10] The EDUCAUSE framework for *AI Literacy in Teaching and Learning* positions AI as a collaborator that assists with research tasks rather than an authority.[11] Leo S. Lo's AI literacy guide positions AI as a supportive partner in human-AI collaboration intended to enhance human abilities.[12] Finally, the ACRL's *AI Competencies for Academic Library Workers* emphasizes human-centered engagement with AI technologies.[13] Libraries too have begun articulating frameworks for AI engagement at the institutional level. Grounded in professional values such as transparency, accessibility, and sustainability, they also champion human oversight. The SDSU Library has developed its own guiding principles and evaluative rubric to both assess existing AI tools and carefully design in-house tools such as the OneSearch Assistant.[14]

INSTITUTIONAL CONTEXT

At SDSU, a 2025 survey found that 62.5% of students used AI applications in their studies. Since then, the institution has moved quickly to address AI in academia.[15] At the institutional level, students, faculty, and staff have access to ChatGPT Edu, Copilot, and Gemini. The university has focused on data privacy and equity by ensuring that tools are available to the entire campus community. AI literacy is an increasing focus, as students and faculty actively experiment with GenAI in learning and research. As an R1 research institution that is part of the nation's "first and largest AI empowered university system," SDSU is very supportive of efforts by the library to develop an AI tool.[16]

The SDSU Library uses a unified discovery environment powered by Ex Libris (Clarivate) Alma/Primo and shared across 22 campuses. It provides access to owned, licensed, and open access (OA) resources through one interface, relying on metadata indexing and relevance ranking to support retrieval. As with many discovery systems, much functionality is hidden from users, who interact primarily through a simplified search box. The release of the Primo Research Assistant in September 2024 inspired the development of the OneSearch Assistant. Primo's assistant processes NLQs to show the five most relevant results and provides a summary of their content. The authors noted that this vendor-designed assistant surfaces and summarizes five items from the Clarivate Central Discovery Index (CDI) and largely obscured the process from the user.[17] The Boolean search string it produced was located outside of the Research Assistant and accessed by clicking "View more results from your library search." Even more problematic from an information literacy standpoint was the way the Primo assistant constructed Boolean strings. Entire sentences were nested within parentheses, and keywords were almost nonexistent (Figure 1).

While the potential of an AI research assistant was immediately apparent to the authors, these poorly constructed Boolean strings did not meet the instructional and research needs at SDSU. While a poorly constructed Boolean string may increase recall, it reduces precision and does not reflect the structure or core concepts of the initial research question. Librarians were motivated to develop an assistant that foregrounds Boolean logic and makes AI search strategies visible and editable. Importantly, it would also allow for all content owned and subscribed by the library to be surfaced.

Figure 1. Sample input comparing Boolean outputs.

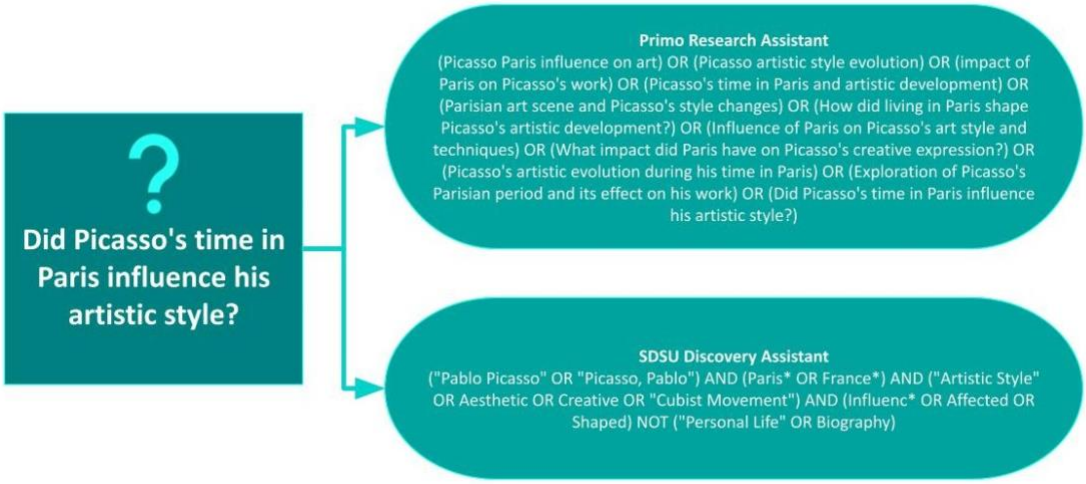

SYSTEM DESIGN/DEVELOPMENT

Initially, the goal was to develop an AI-powered research assistant similar to the Primo Research Assistant, with some important distinctions. It would turn natural language into Boolean searches using keywords instead of sentences. It would search metadata to surface all library content, rather than content in Clarivate's CDI.[18] It would summarize the first five articles as key takeaways as opposed to summaries, and it would expose the search while providing opportunities to filter results and edit the query. These goals were established to support transparency and explainability while encouraging self-directed learning. Librarians wanted to create an AI discovery assistant that aligned with information literacy outcomes and could be used as an instructional device in classrooms.

The SDSU instance of Primo is branded as OneSearch; the authors therefore decided to brand the homegrown research assistant as the OneSearch Assistant. The OneSearch Assistant is built outside of library discovery. Users must opt in to the tool by clicking a banner. They can then enter a query, and the tool uses prompts to either get the record for a known item or generate a Boolean query to search the discovery system. To begin, the researcher enters their query. If the initial prompt identifies the input as a known item, it formats a search for the title and author. This search is sent to the Primo Search API, and a list of known items is returned. The user can select one of the items, view all results in discovery, or continue to the Boolean conversion prompt (Figure 2). If the initial prompt does not identify a known item, it will send the query to a second prompt, which performs the Boolean conversion, searches the Primo Search API with the Boolean string, and returns the number of results found (Figure 3). On the results page, a third prompt presents key takeaways and suggested keywords based on a sample of the first five journal article abstracts (Figure 4). A researcher can change their query and rerun the process or update the Boolean string with the aid of a search operator key (Figure 5). Once satisfied with the results retrieved, the researcher can then manually apply facets, like "articles," "books," or "peer-reviewed," and send their search to the discovery system (Figure 6). These options are presented to the user on the OneSearch landing page (Figure 7).

Figure 2. OneSearch Assistant workflow.

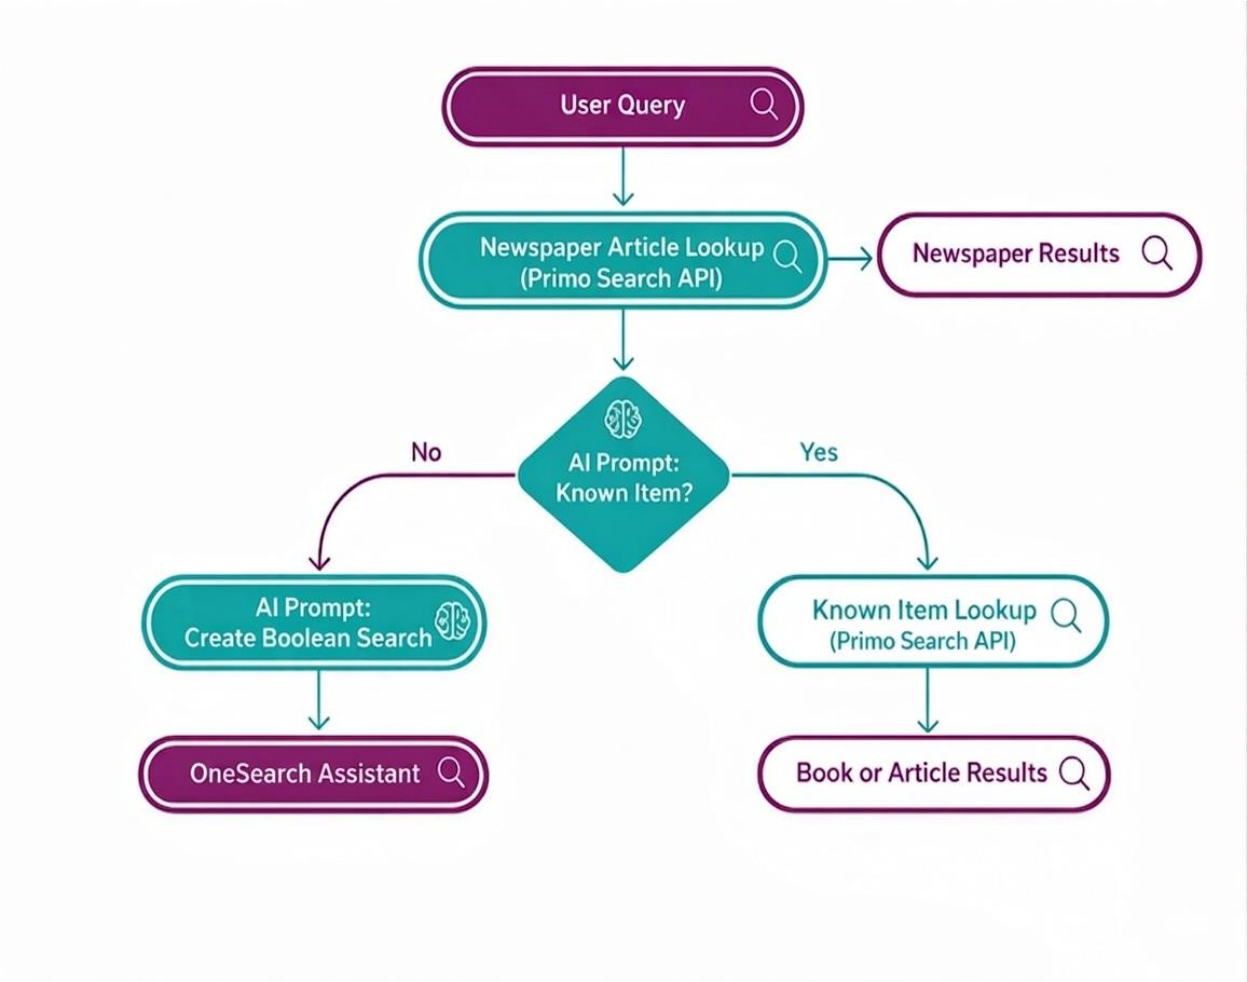

Figure 3. Boolean conversion prompt.

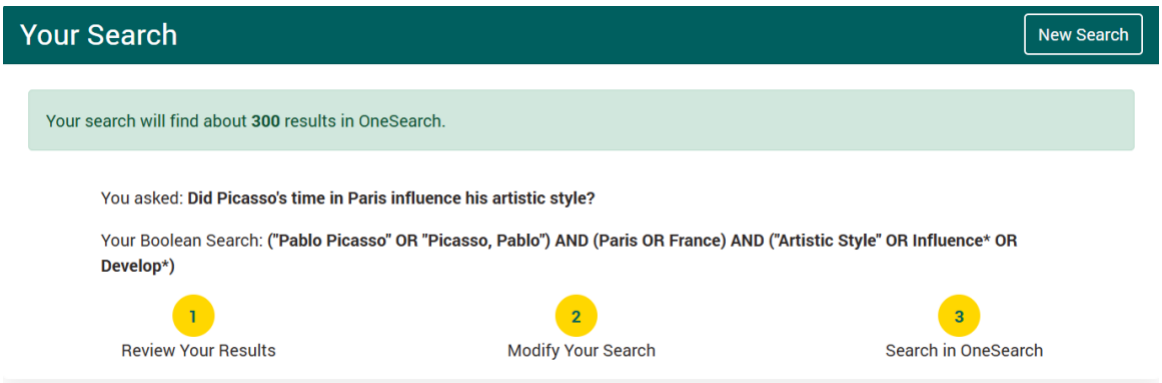

Figure 4. Key takeaways prompt.

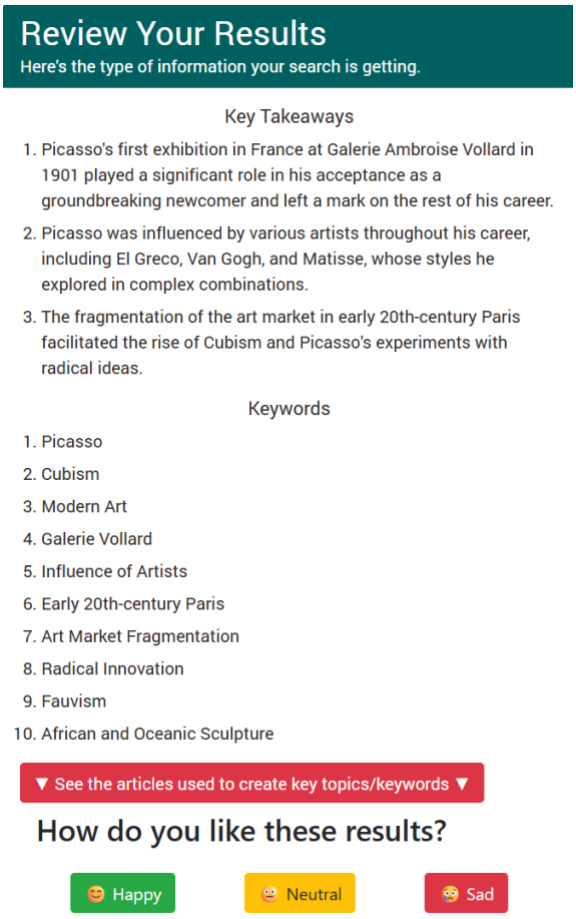

Figure 5. Manual search revision.

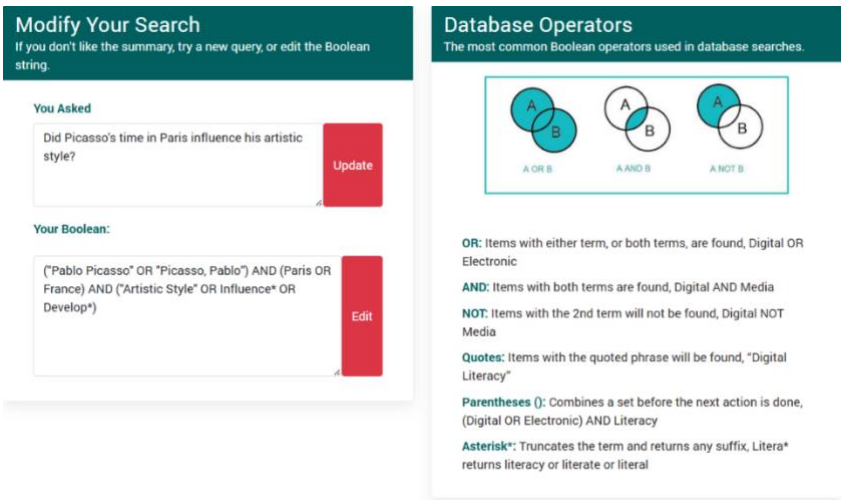

Figure 6. Form search of OneSearch.

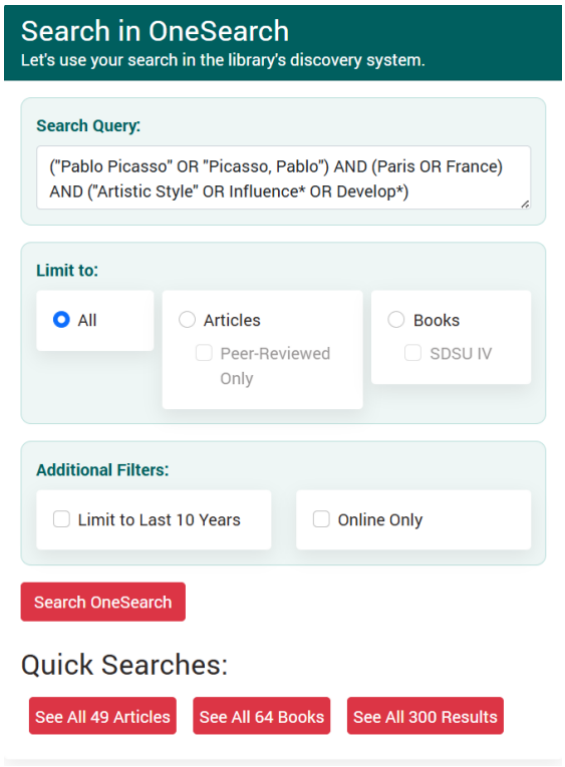

Figure 7. OneSearch landing page.

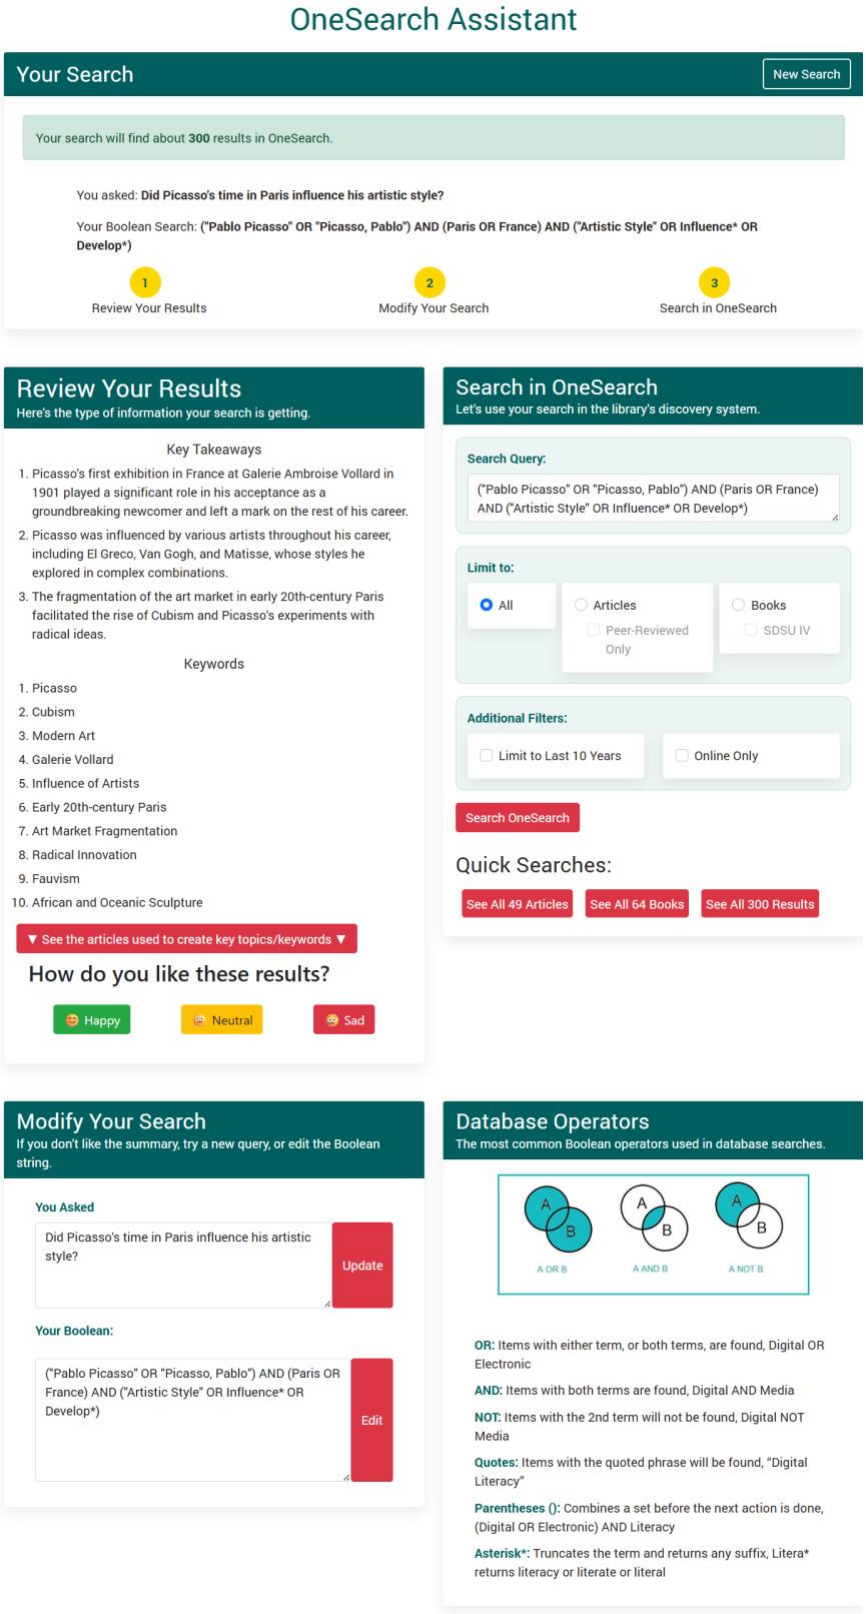

The process of building the discovery tool encompassed many iterative steps of testing and refinement that are still ongoing. Initially, the assistant was built using a locally hosted open AI model (Llama 3.1:70B) on a campus sandbox. The application was then moved to Microsoft Azure Foundry (Redmond, Washington), a production platform for building and deploying AI applications. It offers a range of large language models (LLMs), allowing for consideration of sustainability and speed.

Prompt Engineering Process

The key to the success of the OneSearch Assistant is the underlying prompt that defines how the AI constructs Boolean strings and generates key takeaways. As shown by Wang et al., LLMs can generate complex Boolean strings, but their outputs vary widely depending on prompt design.[19] Prompt development was a central part of the design process. The authors tested prompts on various AI platforms, including ChatGPT, Google Gemini, and Microsoft CoPilot. There are certain behaviors programmed into these platforms that must explicitly be removed. For example, models explain what they are doing, but an explanation or follow-up question muddies the clean Boolean logic. Part of the prompt engineering process was to instruct the AI not to explain itself. The creativity (temperature) and probability (P-value) of the AI model were also experimented with, reducing the possibility of hallucinations.

Additional challenges included handling the specifics of search. The authors studied how Primo and the Clarivate CDI rank search results based on metadata, controlled vocabularies, normalization processes, and boosting algorithms.[20] The prompt had to speak to the specific ways search is performed in both places, for example, by searching for names with both first name, last name order used in Alma, and last name, first name used by the CDI. Another idiosyncrasy of Primo search is that the use of wildcards within a phrase search is not supported.[21] That means that a truncating asterisk nested inside parentheses is ignored. Root words (e.g., *psycholog* or *aspir*) are treated as whole words and searched as such. In order to improve relevancy and ensure accurate results, the prompt had to clarify that nesting and truncation could not be combined. Primo only treats Booleans as operators when they are in all caps, so this was another instruction that had to be clear in the prompt.[22]

Throughout the prompt engineering process, additional changes were implemented through a multistage assessment process that included qualitative and quantitative methods. These included classroom testing, student and faculty feedback, analysis of usage logs, and evaluation of Boolean strings using a rubric. Each stage revealed limitations in the prompt that directly informed revisions.

Prompt Feedback and Testing

Initial testing took place in classroom settings during an honors class of approximately a dozen students led by a librarian. The class, HONOR113-01: Fighting Misinformation, is focused on information literacy instruction and includes a session on advanced search strategies. It was identified as an appropriate opportunity for testing. During the class, students used the assistant to build searches based on their own research questions. These sessions allowed the authors to observe students interacting with the tool and revealed several recurring issues, including inconsistent handling of truncation and a tendency to include material type requests (e.g., book, article) as search concepts. Student feedback was also collected during class time. Some student groups were impressed with the tool, one noting that it “got better results in the last five minutes than we got in the last hour.” However, others were quick to identify issues; for example, “It

included unnecessary words and omitted words that were necessary to our search and added words that are not necessary to our search."[23]

In response, the prompt was revised to impose stricter output constraints and clearer sequencing of instructions. Explicit directives were added to remove references to material types and require consistent use of quotation marks with phrases. The order of operations in the prompt was also adjusted so that phrase identification occurred before truncation, reducing errors caused by truncating multiword terms. These changes improved the consistency of generated searches while reinforcing the instructional goal of modeling expert search construction.

The tool was also shared with Librarian Faculty and Instructional Technology Services Faculty Fellows on campus. Faculty feedback echoed praise from students; for example, one professor summarized, "Very cool! Displaying the Boolean search query information is great because it helps students think more about searching for literature (and that it can be an iterative process)."[24] However, the faculty also identified shortcomings, with one observing, "Some of the articles are fine, though none are the most obvious sources I'd expect students to find."[25] The faculty also highlighted the nuances of search between academic disciplines and recommended the addition of scientific names. For example, a search for chimpanzees would also add *Pan troglodyte*. As a result, the prompt was further revised to improve precision by reducing creativity and adding scientific names. A faculty member also suggested adding context to the prompt to direct its response. Noting that the AI responses often used archaic language that could be upsetting to users, researchers set context by initiating the prompt with a directive, "you are a progressive librarian." *Librarian* was used to set the context and *progressive* to encourage more modern and inclusive keywords.

Second Phase of Testing

In February 2025, a beta version of the OneSearch Assistant was ready for wider deployment and went live on the library interface. Users had to opt in to the tool by clicking a banner encouraging them to try it out. At this stage, quantitative data were gathered and evaluated to underpin observational and qualitative research. Anonymous log files captured real-world queries and subsequent Boolean search strings produced, in addition to the number of queries that resulted in successful searches of library discovery. Immediate issues were identified by reviewing tracked files, for example, treating date ranges as a list of multiple years in the Boolean (e.g., an NLQ "1980 to 2022" became "(1980* OR 1990* OR 2000* OR 2010* OR 2020* OR 2022)") and an inability to recognize newspaper headlines as known items, instead treating them as research questions. Researchers noted that 10% of searches performed were for known items. Consequently, instructions were added to the prompt to remove references to publication dates, and an AI prompt was added to the initial step to redirect known items prior to building the Boolean search.

Researchers also developed a simple rubric to score Boolean strings on a 10-point scale, with one point awarded for each of the following:

1. Handling the entire query
2. Fixing spelling
3. Identification of appropriate synonyms
4. Appropriate use of scientific names
5. Appropriate use of quotes
6. Appropriate use of AND
7. Appropriate use of OR

8. Appropriate use of NOT
9. Appropriate use of wildcards (truncation)
10. Handling proper names

If any of the 10 targets were irrelevant to a search, the point was automatically awarded. Researchers scored the first 100 queries. The goal of the rubric evaluation was an overall score of 8 out of 10, or two failures per conversion. Initial evaluations met this benchmark, with queries averaging a score of 8.7. However, some areas scored poorly. Appropriate use of quotes and wildcards respectively scored 5.5 and 5.9, well below the goal. Analysis found that Primo's treatment of combined quotes and truncations continued to pose a problem. Further revisions were necessary to address this issue. Synonyms and prompts were also areas of concern, each scoring slightly below target.

The tool sometimes got creative with synonyms or discarded original search terms and replaced them with synonyms that weren't quite relevant or what a human researcher would have chosen. These are better described as related terms rather than synonyms; for example, "Mast Cell Activation Syndrome" mapped to "Mastocytosis," a related but different disease. Sometimes the assistant mishandled the prompt; for example, "Chinese American martial arts" looked for "China" AND "USA" AND "Martial arts," entirely removing the crucial phrase "Chinese American."

In summary, a number of adjustments were made to the prompt following evaluation and assessment (see appendix):

- Set Context "you are a progressive librarian"
- Set output expectations earlier
- Adjust the order of operations
- Identify and quote phrases before truncating
- Exclude limiters/facets for publication date ranges and material types
- Include scientific names
- Generalize ages and time periods
- Use clearer language
- Replace "Do not" with "Never" in prompt
- Add a prompt to identify known items
- Add an API to search for newspaper headlines

A subsequent evaluation of queries using the rubric scored an average of 9.4 out of 10, demonstrating these changes improved the efficacy of the AI discovery assistant and led to fewer errors, such as truncation nested within quotes.

AI GUIDING PRINCIPLES

Coinciding with the development of the OneSearch Assistant, the SDSU Library simultaneously drafted guiding principles for AI use. Librarians recognized the rapid development of AI technologies and the need to evaluate vendor-integrated AI tools before introducing them to the academic community. At the SDSU Library, administration, faculty, and staff may choose to adopt AI where it will add value to workflows and/or services. This provision of guiding principles and a self-guided evaluative rubric were designed to support the ethical integration of AI tools and help navigate AI technologies with confidence, uphold institutional values, and foster AI literacy and sustainability. The guiding principles identified are

> Human Oversight: We champion human expertise—AI assists but does not replace library staff and faculty work. We encourage human review of AI outputs, support building expertise, and champion recognizing when to use it and when not to.
>
> Explainability: We select tools that are understandable and provide documentation or training to ensure informed use. Where possible we will avoid vendor reliance by prioritizing open or in-house tools and acknowledge these tools may rapidly change.
>
> Ethics: We respect user autonomy, privacy, and intellectual freedom and support the option to opt out of adopting these tools. We clearly communicate when and how AI is used and we comply with institutional and legal policies.
>
> Accessibility: We commit to audit tools for bias and exclusion and promote inclusive equitable access to AI literacy by supporting underrepresented user communities in AI use
>
> Critical Literacies: We promote AI literacy as an extension of information literacy and encourage users to develop personal research skills and maintain academic integrity.
>
> Sustainability: We evaluate the environmental and financial impact of AI tools and choose interoperable, cost-effective, and maintainable solutions. We seek to minimize intensive use of AI when it has negative social and environmental impacts.[26]

Each of these principles was applied to a rubric and was scored out of 3, with a score of 1 being "Not Good," 2 being "Meh," and 3 being "Good." This timely document allowed the authors to deliberately design the AI discovery assistant to meet these principles, for example, making the tool collaborative by ensuring iterative opportunities to edit output, clearly explaining what the tool does, including a privacy statement, keeping the application opt-in, designing an accessible interface, citing sources, and using a model requiring less energy. Scoring the OneSearch Assistant on the rubric resulted in a score of 16 out of 18, well above the established threshold of 12 points for adoption.[27]

DISCUSSION

This case study demonstrates how GenAI can function as an instructional tool for library discovery. By surfacing Boolean logic and allowing users to inspect and edit queries, the OneSearch Assistant contrasts with existing discovery implementations that prioritize conversational style while concealing the mechanics of search and retrieval. It offers valuable insights into the processes of testing and assessment required to develop a tool that is effective, academic, and ethical.

At the same time, the project highlights tensions between usability and rigor. Natural language interfaces lower barriers for users but risk oversimplification and loss of repeatability. Findings suggest that while AI-generated searches can assist with complex querying, they require human oversight to avoid misinterpretation. These tensions are especially pronounced in advanced research, where convenience must be balanced with transparency. The OneSearch Assistant demonstrates that GenAI can support advanced searching when it is designed to preserve user control and expert judgment. This positions librarians as designers and mediators who can shape AI systems in alignment with institutional standards and research integrity, as opposed to passive adopters of vendor-produced tools.

For libraries developing or evaluating AI discovery tools, this case study suggests several design recommendations. AI systems should be iteratively tested to identify and address persistent issues. They should prioritize explainability, expose the process, and provide editable results. Tools should be built in-house to encourage users to revise and reflect. This approach supports the ACRL *Framework*'s emphasis on inquiry, authority, and strategic exploration, and helps situate AI literacy as an extension of information literacy rather than a separate skill set.

Finally, sustainable and ethical AI deployment depends on governance structures that extend beyond individual tools. Institutional AI guiding principles, evaluative rubrics, and opt-in adoption models provide mechanisms for aligning AI use with values such as equity, accessibility, transparency, and sustainability. Librarian involvement in both design and governance is critical to avoiding vendor silos and ensuring that implementations remain adaptable.

CONCLUSION

The OneSearch Assistant offers a unique search solution. It searches metadata across all library collections and allows for the application of facets. Its educational potential is noteworthy, as initial research suggests it can help students understand the complexity of library search by exposing Boolean search strings. The technical processes of prompt engineering and platform comparison provide valuable insights for librarians and developers interested in building NLQ into discovery systems. By using real student search data, the assessment ensures relevance and effectiveness in practical library settings. Moreover, it addresses the limitations of existing tools, aiming to create a more comprehensive search experience.

As GenAI capabilities evolve, this project offers a foundation for libraries to consider how librarian-designed AI tools can strengthen information literacy while preserving scholarly standards. Future work will continue to refine both the prompts and the interface. This research also provides some best practices for continued careful integration of AI into library workflows and services beyond discovery. It offers an approach to AI grounded in academic principles and assessment that ensures libraries remain active designers and stewards of responsible, transparent, and pedagogically aligned research.

APPENDIX: REFINED PROMPT

You are a progressive librarian helping researchers construct a search of an article database using Boolean operators (i.e. AND, OR, NOT), parentheses, truncation with asterisk, and quotes. The only output should be a Boolean search string. Never provide any commentary or explanation or note explaining actions. For the user provided research query do the following these steps: Correct spelling and grammar. Remove quotation marks and symbols. Remove references to material types like books, articles, reviews, etc. Remove filler and stop words. Remove references to publication dates. Generalize historical time periods. Generalize ages (e.g. use child instead of "age 5". Put multiple word phrases in quotes, for example "20th Century". Always enclose a phrase of two or more words in double quotes, for example "social network". Never quote a single word term, for example ("social"). Do not truncate possessive names, for example "Alzheimer's Disease". For a personal name, write the first and last name together in quotes, then follow it with the last name, first name in quotes, all within parentheses joined by OR (e.g., "Charles Darwin" OR "Darwin, Charles"). Use OR to connect synonyms, for example Cats OR Dogs. Do not exclude unwanted terms with NOT. Identify keywords and select the most common single word term synonyms for each keyword. Include the synonyms along with the original keyword in the search. Include related scientific terms or scientific names as a quoted phrase, for example "pan troglodytes". Group synonyms with parentheses, for example (Students OR Scholars). Never truncate words in quoted phrases. For single terms only, reduce words to their root by stemming and apply wildcard truncation with an asterisk(*) to capture plurals, variations, and related terms (e.g., crime would match crime, crimes, criminal, criminals, etc.). Never truncate words in quoted phrases*. Prior to outputting the response verify that every multiple word phrase or multiple word name or multiple word term are enclosed in double quotes. Never provide any commentary or a note or description or correction of your process. Do not ask any questions. END the output after the first Boolean search string has been created.

*It was necessary to repeat the instruction "never truncate words" in quoted phrases in order to get the prompt to comply consistently.

ENDNOTES

[1] Ex Libris, *The Impact of Generative AI on Libraries* (Jerusalem: Ex Libris, October 2023).

[2] Web of Science Help Center, "Smart Search," updated August 20, 2025, https://webofscience.zendesk.com/hc/en-us/articles/32587676013713-Smart-Search.

[3] JSTOR, "Research Tools," accessed January 12, 2026, https://about.jstor.org/products/jstor-platform/features-and-tools/research-tool/.

[4] Shuai Wang, Harrisen Scells, Bevan Koopman, and Guido Zuccon, "Can ChatGPT Write a Good Boolean Query for Systematic Review Literature Search?," in *Proceedings of SIGIR '23* (New York: Association for Computing Machinery, 2023), https://doi.org/10.1145/3539618.3591703.

[5] Aaron Tay, "The Horseless Carriage of AI Search: Why Using LLMs to Generate Boolean Alone Is Likely of Little Benefit," *Creative Spin* (blog), May 20, 2026, https://aarontay.substack.com/p/the-horseless-carriage-of-ai-search.

[6] M. Sara Lowe et al., "The Boolean Is Dead, Long Live the Boolean! Natural Language versus Boolean Searching in Introductory Undergraduate Instruction," *College & Research Libraries* 79, no. 4 (2018): 517, https://doi.org/10.5860/crl.79.4.517.

[7] Lowe et al., "The Boolean Is Dead."

[8] Jingfan Zang and Tony Russell-Rose, "A Prototype 'Debugger' for Search Strategies," in *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval (CHIIR '23)* (New York: Association for Computing Machinery, 2023), 417–21, https://doi.org/10.1145/3576840.3578321.

[9] Association of College and Research Libraries, *Framework for Information Literacy for Higher Education* (Chicago, IL: American Library Association, 2016), https://www.ala.org/acrl/standards/ilframework.

[10] American Library Association Artificial Intelligence Policy Working Group, *Guidance on the Use of Artificial Intelligence in Libraries*, draft for ALA Board Council and Members, April 7, 2026.

[11] Michelle Kassorla, Maya Georgieva, and Allison Papini, *AI Literacy in Teaching and Learning* (Louisville, CO: EDUCAUSE, October 17, 2024), https://www.educause.edu/content/2024/ai-literacy-in-teaching-and-learning/executive-summary.

[12] Leo S. Lo, "AI Literacy: A Guide for Academic Libraries in Chicago," *College & Research Libraries News* 86, no. 3, 120 (2025) https://doi.org/10.5860/crln.86.3.120.

[13] Association of College and Research Libraries, *AI Competencies for Academic Library Workers* (Chicago, IL: American Library Association, 2025), https://www.ala.org/acrl/standards/ai.

[14] San Diego State University Library, *SDSU Library AI Guiding Principles* (San Diego, CA: San Diego State University, 2025), https://libguides.sdsu.edu/ld.php?content_id=83074108.

[15] David Goldberg, "AI Survey Resources," San Diego State University, accessed April 9, 2026, https://aisurvey.sdsu.edu.

[16] "CSU Announces Landmark Initiative to Become Nation's First and Largest AI-Empowered University System," California State University (press release), February 4, 2025, https://www.calstate.edu/csu-system/news/Pages/CSU-AI-Powered-Initiative.aspx.

[17] Ex Libris, "Getting Started with Primo Research Assistant," Ex Libris Knowledge Center, last modified February 24, 2026, https://knowledge.exlibrisgroup.com/Primo/Product_Documentation/020Primo_VE/Primo_VE_(English)/015_Getting_Started_with_Primo_Research_Assistant.

[18] Ex Libris, "Getting Started with the Primo Research Assistant."

[19] Wang et al., "Can ChatGPT Write?"

[20] Ex Libris, "Search and Ranking in CDI," Ex Libris Knowledge Center, last modified February 17, 2026, https://knowledge.exlibrisgroup.com/Primo/Content_Corner/Central_Discovery_Index/Documentation_and_Training/Documentation_and_Training_(English)/CDI_-_The_Central_Discovery_Index/080Search_and_Ranking_in_CDI.

[21] Ex Libris, "Search and Ranking in CDI."

[22] Ex Libris, "Search and Ranking in CDI."

[23] HONOR-113-01 classroom discussion, San Diego State University, January 31, 2025.

[24] Jeffrey Osbourne, email message to author, February 11, 2025.

[25] Anonymous, email message to author, February 10, 2025.

[26] San Diego State University Library, SDSU Library AI Guiding Principles.

[27] San Diego State University Library, SDSU Library AI Guiding Principles.