ISSN 2161-3974
Full-Length Paper

Journal of eScience Librarianship 13 (1): e800
DOI: https://doi.org/10.7191/jeslib.800

Journal of eScience Librarianship
putting the pieces together: theory and practice

The Implementation of Keenious at Carnegie Mellon University

Joelen Pastva, Carnegie Mellon University, Pittsburgh, PA, USA, jpastva@andrew.cmu.edu
Dom Jebbia, Carnegie Mellon University, Pittsburgh, PA, USA
Maranda Reilly, Carnegie Mellon University, Pittsburgh, PA, USA
Ashley Werlinich, Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

In the fall of 2022, the Carnegie Mellon University (CMU) Libraries began investigating Keenious—an artificial intelligence (AI)-based article recommender tool—for a possible trial implementation to improve pathways to resource discovery and assist researchers in more effectively searching for relevant research. This process led to numerous discussions within the library regarding the unique nature of AI-based tools when compared with traditional library resources, including ethical questions surrounding data privacy, algorithmic transparency, and the impact on the research process. This case study explores these topics and how they were negotiated up to and immediately following CMU's implementation of Keenious in January, 2023, and highlights the need for more frameworks for evaluating AI-based tools in academic settings.

Received: October 2, 2023 **Accepted**: February 5, 2024 **Published**: March 5, 2024

Keywords: Keenious, artificial intelligence, AI, libraries, recommender, ethics

Citation: Pastva, Joelen, Dom Jebbia, Maranda Reilly, and Ashley Werlinich. 2024. "The Implementation of Keenious at Carnegie Mellon University." *Journal of eScience Librarianship* 13 (1): e800. https://doi.org/10.7191/jeslib.800.

Data Availability: Assessment plan survey questions are available under the article Supplementary Files.

Background

As an American research university focused on AI, engineering, and robotics, CMU is uniquely positioned to explore emerging AI tools in the library as well as the laboratory. CMU Libraries, like the university itself, prioritizes innovation; as stated in the CMU Libraries strategic plan, the central goal of our work is to "create a 21st century library that serves as a cornerstone of world-class research and scholarship" (Carnegie Mellon University n.d.). In order to enrich the scholarly information ecosystem, our librarians seek out new tools and resources that could potentially change our information seeking networks for the better. As such we have long been interested in improving the resource discovery process, and strive to find better ways to point users to relevant research made available by the library while also reinforcing information literacy best practices.

A common problem space for libraries is that database search engines require some baseline knowledge of a topic to find relevant content, which can be challenging for inexperienced researchers who may feel overwhelmed when searches return millions of results. Even when content appears to be useful, it often takes a significant investment of time to read portions of articles in order to determine relevance, which can be a daunting process. Libraries have also increasingly acknowledged that research begins outside of the library, and any tools that can improve the research process while also pointing back to library resources are highly desirable (Frederick and Wolff-Eisenberg 2020). Hoping to address these issues, we were intrigued by Keenious, a recommender tool that utilizes search algorithms and AI to analyze input text to suggest relevant academic articles.

As an article recommender tool that utilizes text and documents of interest as its starting point, Keenious is well positioned to help take the guesswork out of how to initialize searches. We also view Keenious as an opportunity to create new pathways to library-subscribed resources because of its integration with the library's link resolver, improving our ability to link to full text content. Additionally, the use of topics in Keenious to encourage exploration of related content had the potential to train users on the benefits of controlled vocabularies for improved and reproducible search results. Finally, the Keenious plugins for Microsoft Word and Google Docs offer multiple ways to meet researchers where they are and better integrate with the research and writing process rather than pushing users to an outside tool or website.

In preliminary discussions surrounding Keenious, our implementation team sought to identify any similar or comparable tools to better frame a needs assessment. Word processor plugin features which integrate with research and writing (such as citation management tools RefWorks and Zotero) have been well received, and we felt this Keenious feature would prove similarly useful. For article recommendations, the Libraries have access via the Primo discovery tool to Ex Libris's bX, which recommends articles based on link resolver usage data collected across discovery systems, databases, and publisher platforms and claims to be platform- and content-neutral (Ex Libris n.d.). Similar to other features offered by content platforms such as Scopus, bX 1) requires users to have run a successful search, 2) is a passive feature built into a larger discovery interface,

and 3) offers no ability to refine recommendations or transparency regarding the recommendation process. Usage data from CMU's Primo Analytics indicated that bX was not often utilized for content discovery in the Primo environment. It was felt that Keenious, as a standalone, portable utility with filtering, searching, and citation generation would be more readily adopted by CMU users.

Although AI technology is increasingly likely to be integrated with modern research tools, some tools more prominently highlight AI as a performance indicator or selling point. CMU Libraries subscribes to Third Iron's AI-powered LibKey suite, which simplifies direct linking to library-subscribed content through searches originating in Primo (LibKey Discovery), library databases (LibKey Link) and external searches such as Google via a browser extension (LibKey Nomad). Based on usage data, LibKey has contributed to a noticeable improvement in utilization of library electronic resources. The functionality of LibKey as a linking solution is viewed as complementary to resource recommender tools such as Keenious, with each addressing a separate but related problem in the research process. A separate team at CMU Libraries also intends to trial SciteAI, a platform that utilizes deep learning to classify and contextualize article citations.

Project details

We initially learned about Keenious in the fall of 2022 through direct vendor contact with the library's Head of Resource and Discovery Services and Director of Library Services. As potentially the first American university library to adopt Keenious, we did not have the usual local network of peer institutions to consult with questions or concerns. As the rollout of the General Data Protection Regulation (GDPR) in the EU taught us, the data privacy landscape in the US differs enough from the EU to limit our ability to lean on the experiences of European library adopters of Keenious for direct comparison. Instead we had to largely develop our own metrics and practices to understand if Keenious would be useful to our users, identify the ethical questions associated with implementing the tool, and determine how to assess the tool as it evolves in the future. We reached out to various stakeholders across library functional areas for feedback, including the library's Collection Advisory Council and Discovery Access Working Group before deciding to move forward with a year-long trial starting in January of 2023. We assembled our implementation team to include representation across library departments. Our team includes Joelen Pastva (Director of Library Services), Dom Jebbia (Library Associate), Maranda Reilly (Electronic Resources Manager), and Ashley Werlinich (Liaison to English and Drama), and our diverse perspectives will let us more effectively and conscientiously assess the implementation of Keenious.

A key component to our implementation of Keenious was creating awareness about the tool, as well as presenting the tool transparently so users could understand tool functionality and navigation with an informed gaze. Before implementation, we developed a solid plan to raise awareness about the tool to our librarians and to the wider university community. This plan included the creation of a Keenious LibGuide, several emails to our library instructor listserv, and an informational session for our library instructors on how to use Keenious (Pastva 2023). By creating several pathways library instructors could use to seek

information, we made sure liaison librarians would have a clear understanding of the tool and be able to represent it effectively and comfortably to our library users. In addition, the Keenious LibGuide helped us to represent the tool to faculty, student, and staff users when they did not have an instruction session that incorporated Keenious into research demonstrations to guide their understanding.

The technical implementation of Keenious was managed by the Electronic Resources Manager, and required standard information about CMU-specific email domains, IP ranges, and our link resolver to connect with library content. The Director of Library Services promoted the tool via the library's social media accounts and website. Our Liaison to English and Drama represented library instructors within our team's discussions around implementation, and developed an assessment plan alongside our Library Associate.

Who is affected by this project?

The role of AI in academic libraries and the research data lifecycle has been a topic of discussion among scholars since the 1990s, as evidenced by works such as Getz (1991). As society moves further into the 4th industrial revolution, interest in utilizing AI in library services has only continued to grow exponentially. Information professionals implementing AI are keenly aware of the practical and ethical challenges that new technologies present (Berendt et al. 2023; Bubinger and Dinneen 2021; Cox, Pinfield, and Rutter 2019). Despite these concerns, libraries around the world are deciding that the benefit to users outweighs the potential risks of AI if properly implemented (Duncan 2022; Ali, Naeem, and Bhatti 2021; Andrews, Ward, and Yoon 2021; Asemi and Asemi 2018; Panda and Chakravarty 2022; R-Moreno et al. 2014). Although AI has repercussions for everyone in society and academia, the people implementing Keenious at CMU Libraries were able to identify three groups that would likely be the largest user base: librarians, researchers and faculty, and students.

Librarians

During the implementation of new AI tools in academic libraries, subject liaison librarians and information professionals are a crucial group to consider. Although they represent the smallest potential user group, they play a vital role in shaping how the library and its services are perceived by other members of the university community. As the primary point of contact for students, researchers, and academic departments, they form strong relationships that are critical to the success of the library.

Subject librarians are also responsible for providing instructional assistance and research guides that faculty members in other departments rely on for their own instruction. Their understanding of Keenious will inform their constituencies' understanding of the product. Thus, it is important to consider the ethical implications of AI on this group since any new features or changes could have a ripple effect throughout the university.

One ethical consideration for subject liaison librarians is the impact of AI on their job responsibilities. AI tools have the potential to automate certain tasks, which could change the nature of liaison work and require new skills. Additionally, AI may introduce biases into the research process or make it more difficult to find and evaluate relevant sources. This can occur in a number of different ways. For instance, the quality and diversity of training data has a major influence on neural network bias, but synthetic datasets and artificial diversity can also degrade the network's performance. Numerous cultural biases are introduced by the predominant usage of English text in the pre-training corpus of many AI research tools. As artificial intelligence continues to advance, and automation becomes more sophisticated, investigators are increasingly relinquishing control to machine agents when it comes to resource discovery and other regular tasks. Students often lack the prerequisite knowledge to understand nuances in a new subject, which makes it difficult to recognize biases and inaccuracies when using new tools. Therefore, it is important that library research guides and instructional materials actively engage with these new technologies.

Furthermore, AI could potentially affect the relationships that subject librarians have with their patrons. While AI can identify relevant resources and provide faster access to information, it lacks the human connection and ability to have substantive conversations that librarians bring to the research process. Subject librarians can identify creative thinking and perform scaffolded instruction that compliments AI tools, and enhances their services rather than replace them.

Overall, Keenious is a modest implementation of machine learning that creates a new kind of search engine. It aligns well with librarians' role in resource discovery. The Keenious implementation team at CMU researched the product and communicated extensively with the Keenious product development team to understand how the product was built, and how best to communicate its features to different audiences.

Faculty/Researchers

AI has significant implications for the numerous professional obligations of faculty and researchers, particularly with regards to student instruction and scholarly publication. Student instruction is critical to the mission of universities because it trains future scholars who will contribute to the advancement of knowledge and society. Likewise, scholarly publication is an important part of securing funding and reappointment for faculty members.

AI has the potential to improve the speed and quality of research, allowing researchers to analyze vast amounts of data and make new discoveries. However, the use of AI in research also creates challenges in reproducibility and systemic bias. Reproducibility is critical in ensuring that research findings can be verified and validated, which is a cornerstone of scientific inquiry. However, the use of AI can make reproducing results more difficult, as the algorithms used in AI may be complex and difficult to replicate.

Additionally, AI can introduce systemic bias into research, which can have significant implications for the validity and reliability of research findings. For example, if an AI algorithm is trained on data that is biased in some way, it may produce biased results that perpetuate existing inequalities or reinforce stereotypes.

When a library is implementing new AI services, it is important to have robust relationships with faculty power users who can provide input on how to integrate AI tools into their teaching and research. Faculty members are key drivers of innovation and change in academic institutions, and their support and expertise can be instrumental in ensuring that AI tools are used effectively and ethically.

Moreover, faculty members play a critical role in propagating skills to the student body, which is why they need to be involved in the implementation and management of new AI services. By working collaboratively with faculty members, libraries can ensure that AI is integrated effectively into the curriculum and research workflows, and that students are prepared for a future where AI will play an increasingly important role in their academic and professional lives.

Students

During the Keenious project, students represented the largest potentially impacted user group, as well as the primary source of ethical considerations. Keenious, while considered by the authors to be a relatively benign application of AI with semantic technology, has the potential to impact student behavior in significant ways. This is particularly important considering the role that libraries play in teaching research skills to students.

One possible way that Keenious and other AI tools can impact student behavior is through their ability to encourage different modes of discovery that may be constrained by the biases of the people who develop them. For example, an AI tool may recommend certain sources of information over others, based on pre-existing biases in the data used to train the AI model. The designers of AI models may be encouraged to seek out data sources that can lead to shepherding of results based on factors including funding sources, legal jurisdiction, country of origin, and numerous other constraints. This can result in a limited and potentially biased understanding of a particular topic, which can impact the quality of a student's research.

Another important consideration is the impact of AI on data privacy. The Family Educational Rights and Privacy Act (FERPA) requires educational institutions to protect the privacy of student educational records, which includes any data pertaining to a student's educational record that is collected, stored, or shared. The use of AI in academic libraries has the potential to collect and analyze large amounts of data, which can pose a threat to student privacy if not handled properly. Therefore, libraries must secure users' privacy rights when choosing new tools and services.

Ethical considerations

Because of its reliance on AI technology, the decision to implement Keenious for a year-long trial went beyond a traditional library resource needs assessment. In the review process prior to implementation,

discussions raised several new ethical questions surrounding data privacy, transparency, and the impact on the research processes. One primary concern was how Keenious collected and utilized data provided by users. In order to assuage these concerns, the CMU team discussed this with the vendor and closely examined the tool's privacy policies and terms of use. Unlike some recommender tools, Keenious does not store personal or user-supplied data to improve its recommendations or further train its algorithms. Interaction data collected for product optimization and usage metrics is anonymized, and Keenious only collects a handful of user data fields for the purpose of account maintenance and authentication that can be easily deleted by a user if desired (Keenious 2023). The fact that the company is based in the European Union and subject to General Data Protection Regulation (GDPR) as a starting point instead of an afterthought also eased privacy concerns.

The common perception of AI tools as "black boxes" which obscure technical processes and data provenance complicated the library's comfort with endorsing Keenious as an emerging tool. This led implementation discussions to focus on various facets of transparency, including the origin of data sources used for Keenious recommendations, potential content biases in the data, and algorithmic transparency. CMU learned that Keenious uses OpenAlex for its article data, which is an open catalog of scholarly outputs with content from respected sources including ORCID, DOAJ, and PubMed. Unlike other recommenders developed by content providers with potential biases toward their own content, OpenAlex was viewed as a neutral, non-commercial data source with a commendable mission to support open source initiatives (OpenAlex n.d.). There is still work to be done to determine whether there are any gaps in subject coverage, which CMU intends to include in future assessment activities.

In discussions about algorithmic transparency, the Keenious technical team acknowledged the difficulty of sharing details about a complex technical process while balancing the user's need for insight into the resulting recommendations. A newer feature built into the tool called "Ranking Information" which shows the score of a recommended article based on shared terms' predicted meanings is a first step toward greater transparency (see Figure 1). The Keenious team expressed a genuine interest in putting transparency first when employing AI technology, and also encouraging users to actively engage with results and interrogate findings rather than use the tool to cut corners. This approach feels especially useful to libraries seeking to promote Keenious, not as a research shortcut, but as a different way to analyze research outputs.

One further concern was raised regarding the long-term roadmap of Keenious as a fairly new tool in a highly volatile and evolving information landscape. Although the tool satisfied CMU's initial criteria for user privacy and transparency, when implementing this tool we also needed to consider 1) How new features and integrations might shift the balance of the tool's emphasis on things like user privacy and transparency and 2) If there is a potential risk of Keenious being acquired by a larger entity and potentially losing some of its neutrality. These questions prompted CMU to request more information regarding the long-term roadmap for Keenious and include this step in future assessment discussions. Although roadmap documents are no

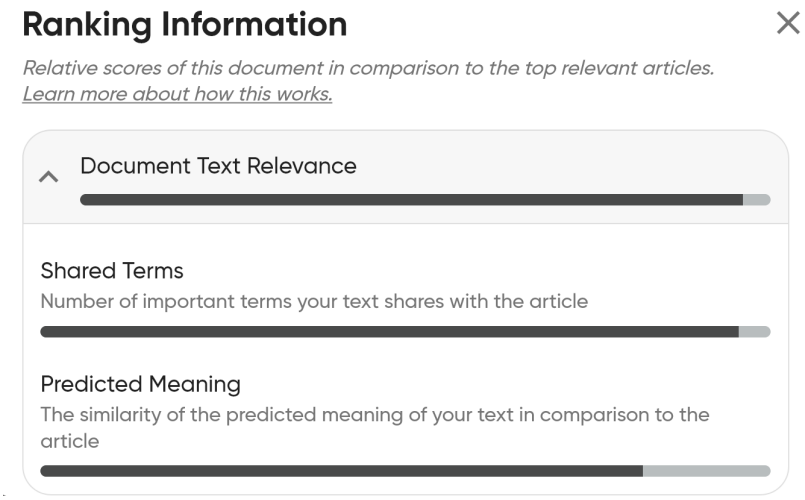

Figure 1: Ranking Information feature in Keenious showing article recommendation score

guarantee of a product's future, they are excellent resources in understanding the business objectives for newer products. Ethical concerns in this realm are a moving target, and any assessment plan for AI-based tools should strive to account for feature expansion over time.

Ethical considerations: Student research behaviors

One of the major ethical implications involved with Keenious is how the tool might change student approaches to research. As Keenious is streamlined to take away the need for the formulation of key search terms, and instead functions on a drag-and-drop or highlight-text approach for ease of searching, it is a real concern that students may change their research habits as a result of using a tool that requires less personal contemplation. Additionally, if students are drawn to Keenious because of the ease of use, we also need to consider whether they will interrogate the sources that result from a Keenious search or just assume these sources are the best for their research. In order to determine whether Keenious was detrimental to student research habits prior to its implementation, we asked the following questions:

1. Will users use Keenious responsibly? What would "irresponsible" use look like?
2. What impact will Keenious have on user behavior? The scholarly research process?
3. Will students see Keenious as a one-stop shop (like many do with Google Scholar)?
4. Will students automatically assume the articles referred are relevant and thus spend less time evaluating sources?
5. Will students trust AI more or less than current tools as a recommender of relevant articles?

These initial questions not only informed whether initial adoption of Keenious was ethical, but also helped us consider various aspects to watch out for during the implementation of Keenious–both in our user assessment stages as well as in our observations as instructors and promoters of this tool. These questions served as touchstones–allowing us to have ethical and user-focused ideas in mind as we approached implementation. While these questions don't all appear explicitly on the surveys we developed, they did inform survey creation; for example, while we may not ask students outright if they feel that they use Keenious responsibly, we created survey questions about where else they research, and whether they thought Keenious presented them with relevant articles.

In addition, as Keenious use grows and we get more research consultations or questions from users related to the tool, these questions will help us shape our informal conversations with Keenious users. By having these frameworks in mind, we can guide reference conversations with students using this tool, and ask informed questions about user perceptions of the tool's relevance, utility, and trustworthiness. In addition to helping us determine questions to ask in our formal and informal assessments, these questions helped us to shape various aspects of our approach to implementing AI-based tools in the future; as it is our hope to create a framework for ethical adoption of AI-based tools within our university, questions like these are indispensable for considering our approach to practical and theoretical concerns around AI tools in the library setting, both now and in the future.

Ethical considerations: Teaching with Keenious

During the Keenious rollout, we recognized the importance of how to represent Keenious to our users through both our instruction sessions as well as through remote pedagogy tools like LibGuides. In order to consider responsible pedagogy with Keenious, we focused on the following factors: 1) framing the tool to our instructors, 2) framing the tool to our internal users, and 3) framing the tool to a broader community.

One of the most important things when giving instruction on any new tool (and especially tools with the potential to change student research habits) is to make sure we are framing the tool not as a one-stop-shop or a superior search tool, but as a component of the larger research ecosystem. For instance, our library instructors have highlighted Keenious in various ways, both as a search tool alongside many different databases, as well as a tool for topic/key term generation for early-stage research. But regardless of how our librarians present the tool to users, it's imperative that we both a) present library instructors with enough information to feel comfortable using the tool in their sessions and b) frame Keenious as just one component of the complex research process.

Additionally, as CMU is the first adopter of Keenious in the US, our librarians have inadvertently become representatives of Keenious to other librarians curious about implementing the tool in their own libraries. We have had numerous libraries contact our librarians about Keenious, and in these interactions our librarians have become instructors not only to our students and faculty, but to other librarians across the

country. As such, our efforts to frame this tool to our library instructors were crucial. Because we cannot be sure which of our librarians will be contacted by other libraries curious about the tool, we need to make sure our librarians have enough information to discuss the tool if necessary at conferences, in reference interactions, or in informal conversations to give peers at other institutions information as questions arise. Additionally, this makes it necessary for us to develop future strategies for situations in which we might also be early adopters of technology, and thus consulted by others about the merits and flaws of such technology.

Ethical considerations: Additional concerns

When implementing Keenious, our team was aware that we not only needed to navigate the ethical concerns of the tool itself, but also needed to consider that users in various departments might have different outlooks on AI tools. As such, we needed to implement the tool knowing that our users would bring their own preconceptions and reservations to the table when approaching both Keenious and any future AI tools we attempt to integrate into the library ecosystem.

For instance, the English department at CMU—like many other English departments globally—has voiced many concerns over tools like ChatGPT, and currently has an ongoing weekly discussion group discussing the ways that ChatGPT and other similar tools could potentially change how students and professionals approach writing as such tools become more complex. These concerns are not limited to the English department at CMU, however. In fact, questions about how to handle the new influx of AI tools in the university environment were enough to warrant the creation of an AI Tools FAQ by CMU's Teaching Excellence & Educational Innovation center (2023). Although this FAQ centers ChatGPT in the discussion, ChatGPT's association with AI tools more broadly means that our discussions with faculty will not happen independently of the associations people have with AI tools in the academic setting.

Following this implementation of Keenious, we intend to gather more information about our faculty's specific reservations about AI tools in order to better discuss Keenious with them in ways that address their particular concerns or questions. As a university is not a monolith, we must approach each department with new eyes and not assume that the questions and concerns of one department will be identical to those of another department. We must remember when meeting with departments and promoting new tools to be open to feedback, to be inquisitive about faculty and student motivations for using or not using these tools, and facilitate conversations with these groups.

Assessment

As our team implemented Keenious, we knew that a key component of analyzing its impact was assessing student, faculty, and staff interactions with the tool; as such, we needed to build in ways to gather information both on how and when users implemented this tool in their own research workflows, as well as information on how this tool potentially changed student and faculty research practices. Additionally, we wanted to

gather information to inform us about whether or not our users consider this tool to be more or less effective than similar resources.

As such, in our assessment, we decided to gather the following types of information:

1. Introductory questions (questions about how users heard about the tool, how often they use it, and what other tools they use for research)
2. Use of the resource (questions about how and when our users implement the tool)
3. Likes and dislikes (questions about efficacy of the tool/ obstacles for use)
4. User experience (questions about usability/ accessibility of the tool)
5. Future use (questions about intended continued use/ recommendations/ etc.)

We developed assessment plans targeting two user groups. First, plans for "internal users" of Keenious (see Appendix 1). These users include faculty and staff within the library as well as faculty from different departments across the university. Although we want a wide variety of internal library users to participate in the internal assessment, instruction librarians and teaching faculty are especially crucial to developing our plans to move forward with Keenious. Their input on how this tool potentially changes the research process—as well as their input on how effective this tool is in finding articles related to a particular search query—is crucial to determining both the ethics and efficacy of this tool in our university's research ecosystem.

Second, we developed assessment plans for "student users" of Keenious (see Appendix 2). Although the questions asked of the student users are somewhat similar to the questions we asked of internal users, the surveys differ slightly in asking students about the types of projects they used Keenious for (see Appendix 2), and asking internal users if they intend to use the resource in their teaching.

This assessment program will launch in the fall semester of 2023 allowing us time to promote the resources with fall classes, and to hopefully accrue some repeat users of Keenious so we can get more meaningful feedback.

When distributing surveys, we will not be targeting only Keenious users—due to privacy concerns in directly contacting users of the tool—but instead we plan to distribute our surveys both through broader listservs (i.e. department-specific, library listservs, etc.) as well as through targeted emails sent to faculty we know have investment and interest in AI conversations more broadly. Additionally, we wanted to include a mechanism for feedback in our Keenious LibGuide; as the LibGuide is one of the direct lines of communication between the library and Keenious users, it made sense to have an alternate pathway to submitting feedback there as well. Although this survey is shorter than the other two (see Appendix 3), having an additional pathway to collecting user responses gives us another potential in-road both for starting conversations with our users and for improving the tool with minimal effort on our end.

Documentation

During implementation, our team consulted several policies, best practices, and codes of ethics for additional guidance on ethical considerations relevant to AI-based technologies in libraries. Overall, Keenious adheres to the recommendations and best practices pertaining to data privacy and algorithmic transparency. The below sources were, and will likely continue to be, valuable to reference for evaluation and implementation of similar recommender tools.

- CMU Academic Integrity Policy - As the promotion of academic integrity is a core responsibility for the CMU community, the university's Academic Integrity policy is integral at the institutional level for our team to ensure all services and tools offered through the Libraries meet stated expectations (Carnegie Mellon University 2020).
- ODI Recommended Practice - Facilitated by the NISO Open Discovery Initiative Standing Committee, the ODI Recommended Practice aims to promote the adoption of conformance statements and to streamline the relationships between discovery service providers, content providers, and libraries. They offer general recommendations, as well as best practices and conformance checklists for each sector. The Keenious documentation aligns with NISO ODI recommendation for Discovery Services Providers to "explain the fundamentals of how metadata is generally utilized within the relevance algorithm (mapping metadata to indexes, weighting of indexes, etc.) and how it enhances discoverability" (Open Discovery Initiative Standing Committee 2020).
- IFLA Statement on Libraries and Artificial Intelligence - The IFLA Statement on Libraries and Artificial Intelligence additionally provides key ethical considerations for AI technologies in libraries, including privacy, bias, and transparency (IFLA Committee on Freedom of Access to Information and Freedom of Expression 2020). For example, IFLA notes the importance of libraries to know how vendors train AI systems and tools, and that transparency and explainability can be beneficial to detect and address bias. Their framework included the below recommendations for Libraries adopting AI tools, which were invaluable in our initial evaluation, procurement, and implementation of Keenious:
    - Help patrons develop digital literacies that include an understanding of how AI and algorithms work, and corresponding privacy and ethics questions
    - Ensure that any use of AI technologies in libraries should be subject to clear ethical standards and safeguard the rights of their users
    - Procure technologies that adhere to legal and ethical privacy and accessibility requirements

- IFLA Code of Ethics for Librarians and other Information Workers - The IFLA Code of Ethics for Librarians and other Information Workers was referenced (IFLA 2012), which embodies the ethical responsibilities within the library profession. This Code of Ethics serves as a set of guiding principles for providing information service in modern society, with an emphasis on social responsibility.

Lessons learned and future work

Although still in its early phases, the implementation of Keenious at CMU has been an eye-opening experience that has resulted in a number of takeaways and ideas for future investigation. We quickly realized that, while AI tools can in many ways be approached from an area of need similar to other library resources, the nature of the underlying technology requires additional considerations to best determine what is appropriate for library adoption. These considerations are ethically fraught, as they touch on sensitive issues including privacy, transparency, and the perceived impact on research behaviors and pedagogy. We have only scratched the surface with our ethical discussions at CMU, and much work remains to ensure that we engage with our campus population in evaluating the long-term impact tools such as Keenious have on research behavior. We must also carefully structure our assessment strategy to track changes that may be ethically concerning, such as new features or data collection activities.

This case study has also highlighted the need for new frameworks for evaluating the complete lifecycle of AI-based tools, from acquisition to implementation to ongoing assessment. Much as the NISO Open Discovery Initiative grew from the need for best practices and standards surrounding index-based discovery services, the unique nature of tools centered around AI technologies requires new standards for carefully examining product features and vendor policies prior to and following implementation.

It is also important to engage with product vendors as much as possible when ethical questions about their tools are raised. Librarians have long been experts in advocating for user privacy, and should engage with vendors in conversations about what privacy looks like in today's data-driven landscape. We had very positive experiences working with Keenious, as they were quick to answer questions, provide supporting documentation, and connect us with their technical teams to better understand their product. They also organized a workshop for librarians on generative AI in the spring of 2023 to gather more feedback on the ethical implications of AI-powered tools, and to guide future development of their product. They clearly understood our need to provide ethically sound tools for our campus community, and demonstrated a genuine interest in developing their technology responsibly.

Ultimately librarians can help to define what effective transparency means in the research and information landscape to ensure that users can engage critically with new technologies. We must acknowledge that disruptive change from AI tools has already arrived, and libraries should be proactive in preparing themselves for whatever ethical challenges lay ahead.

Data Availability

Assessment plan survey questions are available under the article Supplementary Files:

Appendix 1: Survey For Internal Assessment

Appendix 2: Survey For Student Users of Keenious

Appendix 3: LibGuide Survey

Acknowledgements

The research case study was developed as part of an IMLS-funded Responsible AI project, through grant number LG-252307-OLS-22.

Competing Interests

The authors declare that they have no competing interests.

References

Ali, Muhammad Yousuf, Salman Bin Naeem, and Rubina Bhatti. 2021. "Artificial Intelligence (AI) in Pakistani University Library Services." *Library Hi Tech News* 38 (8): 12–15. https://doi.org/10.1108/LHTN-10-2021-0065.

Andrews, James E., Heather Ward, and JungWon Yoon. 2021. "UTAUT as a Model for Understanding Intention to Adopt AI and Related Technologies among Librarians." *The Journal of Academic Librarianship* 47 (6): 102437. https://doi.org/10.1016/j.acalib.2021.102437.

Asemi, Asefeh, and Adeleh Asemi. 2018. "Artificial Intelligence(AI) Application in Library Systems in Iran: A Taxonomy Study." *Library Philosophy and Practice* 1840. https://digitalcommons.unl.edu/libphilprac/1840.

Berendt, Bettina, Özgür Karadeniz, Sercan Kıyak, Stefan Mertens, and Leen d'Haenens. 2023. "Bias, Diversity, and Challenges to Fairness in Classification and Automated Text Analysis. From Libraries to AI and Back." arXiv. https://doi.org/10.48550/arXiv.2303.07207.

Bubinger, Helen, and Jesse David Dinneen. 2021. "Actionable Approaches to Promote Ethical AI in Libraries." *Proceedings of the Association for Information Science and Technology* 58 (1): 682–684. https://doi.org/10.1002/pra2.528.

Carnegie Mellon University. n.d. "Strategic Plan 2025." Accessed March 31, 2023. https://www.cmu.edu/strategic-plan/strategic-recommendations/21st-century-library.html.

Carnegie Mellon University. 2020. "Carnegie Mellon University Policy on Academic Integrity." University Policies. Accessed March 27, 2023. https://www.cmu.edu/policies/student-and-student-life/academic-integrity.html.

Cox, Andrew M., Stephen Pinfield, and Sophie Rutter. 2019. "The Intelligent Library: Thought Leaders' Views on the Likely Impact of Artificial Intelligence on Academic Libraries." *Library Hi Tech* 37 (3): 418–435. https://doi.org/10.1108/LHT-08-2018-0105.

Duncan, Adrian St. Patrick. 2022. "The Intelligent Academic Library: Review of AI Projects & Potential for Caribbean Libraries." *Library Hi Tech News* 39 (5): 12–15. https://doi.org/10.1108/LHTN-01-2022-0014.

Eberly Center for Teaching Excellence & Education Innovation. 2023. "AI Tools (ChatGPT) FAQ." Accessed March 30, 2023. https://www.cmu.edu/teaching/technology/aitools/index.html.

Ex Libris. n.d. "bX Recommender." Accessed March 29, 2023. https://exlibrisgroup.com/products/bx-recommender.

Frederick, Jennifer, and Christine Wolff-Eisenberg. 2020. "Ithaka S+R US Library Survey 2019." Ithaka S+R. https://doi.org/10.18665/sr.312977.

Getz, Ronald J. 1991. "The Medical Library of the Future." *American Libraries* 22 (4): 340–343. http://www.jstor.org/stable/25632204.

IFLA. 2012. "IFLA Code of Ethics for Librarians and Other Information Workers (Full Version)." August 2012. https://www.ifla.org/publications/ifla-code-of-ethics-for-librarians-and-other-information-workers-full-version.

IFLA Committee on Freedom of Access to Information and Freedom of Expression. 2020. "IFLA Statement on Libraries and Artificial Intelligence." October. https://repository.ifla.org/handle/123456789/1646.

Keenious. 2023. "How Keenious Recommends Research Articles." Keenious Knowledgebase. Accessed February 27, 2023. https://help.keenious.com/article/54-how-keenious-recommends-research-articles.

McCorduck, Pamela, and Cli Cfe. 2004. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence* (2nd ed.). A K Peters/CRC Press. https://doi.org/10.1201/9780429258985.

OpenAlex. n.d. "About." OpenAlex. Accessed March 27, 2023. https://openalex.org/about.

Open Discovery Initiative Standing Committee. 2020. "NISO RP-19-2020, Open Discovery Initiative: Promoting Transparency in Discovery." NISO. https://doi.org/10.3789/niso-rp-19-2020.

Panda, Subhajit, and Rupak Chakravarty. 2022. "Adapting Intelligent Information Services in Libraries: A Case of Smart AI Chatbots." *Library Hi Tech News* 39 (1): 12–15. https://doi.org/10.1108/LHTN-11-2021-0081.

Pastva, Joelen. 2023. "Keenious: Introduction." LibGuide. https://guides.library.cmu.edu/c.php?g=1293235&p=9497310.

R-Moreno, María D., Bonifacio Castaño, David F. Barrero, and Agustín M. Hellín. 2014. "Efficient Services Management in Libraries Using AI and Wireless Techniques." *Expert Systems with Applications* 41 (17): 7904–7913. https://doi.org/10.1016/j.eswa.2014.06.047.