

The Evolution of Library Workplaces and Workflows via Generative AI

Mohammad Hosseini and Kristi Holmes

ChatGPT was released on November 30th 2022, and very quickly popularized generative artificial intelligence (AI) to the extent that it is now seen as a mainstream technology and used by many. However, this mainstreaming and popularity has also resulted in a hype, thereby overwhelming us by a wide range of opinions and news related to its current and future applications. While we can test generative AI applications and read news about their added value, it might be hard to envision the short-, medium- and long-term impact of these tools on library operations, resources, and services.

Reflecting on how libraries and their existing workflows are evolving alongside the rise of generative AI is intriguing, yet extremely challenging due to the rapid development of the technology. This is further complicated by the variation across each library's organization, management, and use. Indeed, even two libraries in the same institution might have a different approach regarding collection management, curation, user engagement, and technology integration. These differences can be further amplified by library size, disciplinary focus (e.g., university library, medical library, law library), services offered (e.g., education and training, evaluation), communities served (e.g., students, medical trainees, researchers, faculty, public), overall approach to technology, as well as their budget. Each library is made up of several units with specific goals and responsibilities. The same technology may have a different impact on each department of a library, as well as each employee, depending on their role and background. As a result, despite various opportunities and challenges presented by generative AI, librarians should consider charting their own personal roadmap to learn about and familiarize themselves with this technology based on their unique circumstances, interests, and needs.

To understand the perspectives of different libraries and librarians on this journey, we had short conversations about generative AI with eight individuals who are involved in various roles in different libraries. Across these different libraries and roles, there were several consistent themes. Everyone noted the importance of keeping up to date about this technology and understanding opportunities and challenges. All individuals we conversed with were aware of so-called hallucination problem in generative AI (generation of nonsensical or untrue content) and biases embedded in these systems, but stressed that libraries should learn about this

* Mohammad Hosseini and Kristi Holmes are each with the Department of Preventive Medicine, Northwestern University Feinberg School of Medicine; the Institute for Artificial Intelligence in Medicine, Northwestern University Feinberg School of Medicine; and the Galter Health Sciences Library and Learning Center, Northwestern University Feinberg School of Medicine; email: mohammad.hosseini@northwestern.edu. © Mohammad Hosseini and Kristi Holmes, Attribution-NonCommercial (<http://creativecommons.org/licenses/by-nc/4.0/>) CC BY-NC

technology and use it to their benefit. When asked about current use cases, opportunities, and challenges, we heard a range of responses that reflected perspectives and experiences of each individual, closely tied to their specific context and roles. A summary of these views follows.

How Generative AI is Currently Used in Libraries and What Are Some of the Potential Opportunities and Concerns?

Those who contributed to the educational mission of libraries considered opportunities in terms of improved lesson planning and retrieving better examples or ice breakers for patrons. Nichole Novak, the Head of Reference and Instruction Services at Galvin Library at Illinois Institute of Technology and her colleagues have experimented with the free version of ChatGPT to kick-start their thought processes for lessons and help draft scripts for videos. Although they noticed that the generated content needs further refinement and editing, using ChatGPT helped them to not start from scratch. Nichole shared that she has heard AI being used for basic reference questions too. In addition to her concerns for the quality and accuracy of these AI-mediated references, she was worried about losing the human connection with patrons:

A reference interaction is a good moment for making a connection with your patron, putting a face on the library, and also an opportunity to teach them how to access a database and search for resources. A library is much more than its collection, it's a space for patrons to interact, learn and build a community. In this sense, AI lacks the welcoming presence of a librarian who knows our library.

She also raised concerns about students using generative AI to do assignments:

We didn't see a lot of it but we have had instances. The goal of completing assignments is for students to learn. When they use these tools, they don't learn about information retrieval, and how to find and use reliable sources. Until recently, students and researchers would visit websites, and based on how the landing page of a website looked (e.g., the number of pop-ups and broken links, or URL), they would get a quick sense about validity or reliability of information. But ChatGPT presents information (true or false) in a sanitized manner without letting them do the evaluation or know anything about the used sources.

Molly Beestrup, the Head of Research and Information Services at Galter Health Science Library at Northwestern University echoed similar views about students who have used ChatGPT to do their assignments. She has used the free version of ChatGPT to generate examples that would be understandable for students or researchers at a certain level (e.g., first-year medical students) or summarize text from a critical appraisal perspective. She emphasized that she does not use the generated content nor fully trust it before verifying it through several reliable sources because of the hallucination problem. She added that this critical attitude is absolutely necessary when working with these systems but might not always be employed:

Students and early career researchers might not know enough to understand if the generated content is actually true or not. I worry that generative AI is going to be used as an

information retrieval tool. At the moment, even a search engine like Google is better than ChatGPT in that respect.

She further highlighted that these systems could be a massive timesaver for writing, and could level the playing field for those studying and working in a second language, or those who might not have a great command of grammar or spelling. She believed that libraries should not stress out and instead, investigate how their patrons are using these tools, to understand the benefits and limitations of specific use cases. She acknowledged that given the current pace of information transmission and dissemination, “humans cannot keep up,” and this necessitates leveraging new tools:

I think Generative AI tools might have a good potential for systematic reviews in the future. We just need to figure out the best way to do this. For example, when doing systematic reviews, we ask researchers to suggest studies that they think should definitely be included in their sample. We have received several fake citations recently, which, after further investigation, were found to be generated by ChatGPT. One major concern is to have these hallucinated citations get into our papers and studies.

Another librarian who spoke with us was Amy Chatfield, an Information Services Librarian based at the University of Southern California. She also highlighted the potential of using these tools in systematic reviews, when screening for evidence, and added:

Currently, models like ChatGPT do not have access to the full text of most articles. We can train them to become better in one task but that takes a considerable amount of time and resources, without necessarily making the model useful for other tasks. Accordingly, with some tasks it would make more sense to do it ourselves instead of trying to program a model.

Even in the case of tasks that have been automated for a long time, she thought, human contributions add value to what computers can do:

For instance, we do not do so much indexing anymore because computers are perceived to be better at it. But when a computer indexes, it searches for words, characters and letters. Humans are needed to tell AI, ‘hey, epilepsy and epileptic are in the same domain, so put them together.’ This is easy, but it is much harder when you try to teach AI something like ‘here’s an article that’s talking about how to transition someone from carbamazepine to valproic acid’. The article might never mention the word epilepsy anywhere because everyone reading that article already knows these drugs are used to treat epilepsy. Since there are a lot more context clues behind the scenes than people might notice, training AI for each specific context and situation might take a lot of time and resources.

Amy has tried both free and paid versions of ChatGPT, Google’s BARD and the AI-powered version of Bing search engine. She was concerned that generative AI might be used to cut out humans from the research process, because:

There are many parts of librarianship that people consider as rote and easy and simple, that actually turn out to be not that way. Let's consider search strategy development as an example. There have been experiments to train AI to come up with a sensitive and well-done search. However, those who train AI may not know what a good search actually is. So, they come up with something that looks beautiful but is not as comprehensive as what librarians would develop using manual searches. Computer scientists who build these models do not want to believe that every single article and abstract is not similarly structured and might not contain the same information in it. Sometimes we are dealing with three or four levels of abstraction and AI cannot always be trained to handle these.

Despite these shortcomings, Amy thought that some might be happy with what AI does because completing searches by humans is tedious and difficult. She added that addressing these challenges and improving workflows that integrate human and AI capabilities require a better interplay between computer scientists and librarians. She highlighted the significance of hiring computer scientists in libraries, like those who can code or train large language models.

Bart Davis, the Head of the Collection Management and Metadata Services department at Galter Library has discussed the use of generative AI with his department but due to concerns around accuracy, they have not used this technology to create any metadata:

We need to spend time and resources to program a model but if it ends up generating inaccurate metadata, that creates a lot of additional work for catalogers. We are hesitant to jump on the bandwagon of generative AI right now to create metadata because it is not fully tested for our context, but I can see some libraries or vendors embracing it in the future. Either way, clear disclosure and labeling of AI-generated metadata, followed up by human validation are absolutely crucial.

Since Bart has some programming skills, he currently uses the free version of ChatGPT to boost his capabilities in writing code and to generate Python scripts for standard and repetitive tasks, which he can then verify. He also envisions having AI as a cataloging assistant:

Especially if they could be integrated into library management systems; suppose you create a record and then AI pops up and says 'hey, maybe consider this subject heading, or maybe update this old record based on how you have recently indexed or cataloged items.'

Using AI to enhance coding capabilities was also highlighted by Matt Carson, Senior Data Scientist and Head of the Digital Systems at Galter Library, who is responsible for Galter's library system infrastructure, as well as several different technical and data-oriented research projects. Given his management responsibilities, Matt infrequently codes and appreciates reminders about how a certain function works. He has used ChatGPT to find examples, document code, and add comments to existing code, as well as generate random data sets to test his code. Since Matt uses the paid version of ChatGPT, he can also install third-party plugins to improve the efficiency of his personal knowledge management through gathering information faster and summarizing it. He added:

I think off-the-shelf generative AI models are going to improve through these third-party plugins where developers have really focused on one aspect, and continuously improve it for better accuracy. These plugins will be an asset for librarians and researchers, and libraries should embrace these types of tools.

Matt asserted that AI tools are neither going to replace researchers nor librarians, but those who fail to learn about and incorporate these tools in their daily workflows may find themselves outperformed by those who have embraced it. Brendan Quinn, Senior Developer at Northwestern University Libraries in Evanston, Illinois relates:

Developers in my team have adopted the use of GitHub Copilot (powered by OpenAI's GPT-4) for all day-to-day programming. Using GitHub Copilot is a bit like having a coding partner beside you at all times.

Brendan shared details about their prototype which takes advantage of generative AI to develop an application for discovering information about their digital collections by chatting with a generative AI model. The prototype takes a user's question and performs a vector database query that retrieves documents based on how closely they match the question semantically. The application then sends the question and retrieved documents to the generative AI model, which provides a response grounded on the data contained in those documents. They are also experimenting with hybrid search that takes advantage of the precision of traditional keyword-based search and the semantic understanding provided by vector search capabilities. While he called generative AI the biggest game-changer he has witnessed in terms of using new technologies in libraries, he also shared concerns regarding misconceptions surrounding this technology:

I spend a lot of time explaining to folks that generative AI is neither connected to, nor works like a database – it is not like a dictionary or encyclopedia that has the data inside of it. Therefore, we should treat these more like systems that make mistakes similar to those made by humans. When I speak with a person, I do not automatically assume that everything they said is correct, factual and true. I adopt a similar approach when engaging with generative AI, take it with a pinch of salt!

These misconceptions were also echoed by Bianca Kramer, former scholarly communication and open science librarian at Utrecht University in the Netherlands, currently a consultant at Sesame Open Science. She stressed that assuming everything done by humans is the gold standard can be misleading, because we are all biased and make lots of mistakes. Therefore, thinking about how humans and technology can complement each other would be more beneficial. She was concerned that systems like ChatGPT give the impression of an authoritative source without having or claiming any authority. Given the current trend in using them, she expressed concern:

At some point, AI-generated sources will become the source. It will be interesting in the next decade to see what will happen to primary sources, and to what extent they are going to be replaced by sources generated by AI.

Jeroen Bosman, an open science and scholarly communications specialist at Utrecht University believed that the future applications of generative AI in libraries and society depend on many factors:

Besides depending on how they will develop in a technical sense, we still need to see how they will be embedded in our social and legal contexts. They can be a real timesaver for some tasks and can increase interdisciplinary relations. On the flip side, they will amplify the mainstream perspectives and could have a detrimental effect on critical thinking skills.

Jeroen recommended that libraries should adopt a balanced approach and move slowly, through small-scale implementations. While he encouraged librarians to test these systems, he also invited them to adopt a critical approach towards using them, e.g., learning how to refine and evaluate their own prompts. He added that this approach is used in Utrecht University's information literacy program and new workshops that aim to explore the impact of generative AI on the information landscape:

For centuries we collected information packaged in concrete items written by someone or a group. The onus was then on the user to extract relevant information and insights from these packages. Thanks to the internet and advancements in communication technology, we have so many of these packages to unpack and explore. This has made us realize that our existing information retrieval paradigm is extremely inefficient and generative AI might help us. However, it is very important to realize that there are different streams, approaches, visions and ideologies within societies, all of which are embedded in those separate concrete information packages. And we might lose those nuances or the ability to use a dialectical approach if we just have sort of one answer from the machine for everything.

A similar apprehension was voiced by Verónica Hoyo, Executive Director of the Network of the National Library of Medicine (NNLM) National Evaluation Center (NEC), based in Northwestern's Galter Library. The NEC is charged with assessing NNLM activities, services, and resources with special focus on understanding its impact on persons who experience health disparities or are underrepresented in biomedical research. She has used the free version of ChatGPT as a sparring partner (to improve surveys and challenge her own biases) and highlighted that while these models might get better in the future, they have inherent shortcomings:

These models are currently not representative, because the data used to train them have not been representative. Although the training data of these models have not been disclosed, given that access to technology is already unequally distributed and data from substantial sections of the population are missing in our corpus, one can infer that these models have major limitations and gaps.

Conclusion

Libraries play a central role on campus and will remain an indispensable partner and catalyst for universities' research and educational endeavors. Therefore, it is no surprise that libraries

are actively guiding the consideration and use of generative AI on their campuses. On this journey, libraries will develop strategies to responsibly leverage generative AI technology while carefully managing risks. Generative AI is evolving at a rate rarely seen before, making it difficult to anticipate every challenge and develop comprehensive and consistent policies in response. By sharing different perspectives across different roles, geographies, and types of libraries, we can learn from one another, adjust strategies and services, and inspire new partnerships and opportunities for generative AI and other technologies to support and advance the work of our libraries.

To capture a diverse spectrum of opinions and use cases of generative AI in libraries, we engaged with four individuals from Galter Health Sciences Library at Northwestern University (Chicago, Illinois), one from Northwestern University Libraries (Evanston, Illinois), another from a different university in Illinois (Illinois Institute of Technology), another from a university in California (University of Southern California), and two librarians based in Europe. This approach allowed us to start with our own context and then step beyond our local environment to obtain a more nuanced perspective. These responses not only showcased the multifaceted nature of generative AI but also highlighted the unique ways in which it is being leveraged and perceived within the library community. The breadth of insights gathered underscored the evolving narrative of generative AI in libraries.

Acknowledgement

We wish to thank the individuals who spoke with us about challenges and opportunities of using generative AI in libraries, including:

- Nichole Novak, Galvin Library, Illinois Institute of Technology
- Molly Beestrum, Galter Health Sciences Library, Northwestern University
- Amy Chatfield, Norris Library, University of Southern California
- Bart Davis, Galter Health Sciences Library, Northwestern University
- Matt Carson, Galter Health Sciences Library, Northwestern University
- Brendan Quinn, Northwestern University Libraries, Northwestern University
- Bianca Kramer, Sesame Open Science, the Netherlands
- Jeroen Bosman, Utrecht University, the Netherlands
- Verónica Hoyo, NNLM National Evaluation Center, based at Northwestern University

Funding Statement

This work was supported, in part, by the National Institutes of Health's National Center for Advancing Translational Sciences UL1TR001422 (KH, MH) and the National Library of Medicine's Network of the National Library of Medicine U24LM013751 (KH, VH, MC, MH). The funders have not played a role in the design, analysis, decision to publish, or preparation of the manuscript.